



User Guide

AWS Support



API Version 2026-09-18

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

AWS Support: 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

Get started with AWS Support	1
AI-enhanced troubleshooting in Support Center Console	2
Set up permissions to use AI-enhanced troubleshooting	5
Create a support interaction	10
Create a support case from a support interaction	12
View support interactions	18
Managing attachments	18
Troubleshooting	24
Virtual meetings with AWS Support	25
Join a virtual meeting	26
Security and privacy during your virtual meeting	26
Required IAM permissions for virtual meetings	27
Troubleshooting virtual meetings	27
Case management	28
Describing your problem	29
Choosing an initial severity level	29
Understanding AWS Support response times	31
Changing a support case severity level	33
Request a service quota increase	36
Legacy method: Create support cases and case management	37
Creating a support case	38
Describing your problem	40
Choosing an initial severity level	41
Understanding AWS Support response times	42
Changing a support case severity level	44
Example: Create a support case for account and billing	46
Legacy experience: Update, resolve, and reopen your cases	52
Working with AWS SDKs	58
Using AWS DevOps Agent	60
Plan eligibility and credits	60
Enabling AWS DevOps Agent	61
Resources created	62
Removing AWS DevOps Agent	64
Security	65

Getting help	65
About the Support Center Console API	66
Adding IAM policies for the Support Center Console API operations	66
Testing Support Center Console API calls	71
About the AWS Support API	72
Support case management	74
AWS Trusted Advisor	74
Endpoints	75
Support in AWS SDKs	75
AWS Support Plans	76
Features of AWS Support Plans	77
What is AWS Unified Operations	79
AWS Unified Operations pricing	79
Benefits of Unified Operations	80
Unified Operations Team	81
Unified Operations life cycle	83
Getting started with Unified Operations	86
Change AWS Support Plans	93
Related information	94
Sign up for AWS Enterprise Support using self-service subscription	95
Checking eligibility requirements	95
Subscribe to AWS Enterprise Support	96
Understanding post-enrollment steps	98
Troubleshooting	99
Changing or canceling your plan	100
Configure promotional plan expiration notifications	101
View promotional plan notifications	102
Developer, Business, and Enterprise On-Ramp end of support	102
Developer Support plan end of support	102
Business Support plan end of support	102
Enterprise On-Ramp end of support	103
AWS Trusted Advisor	104
Get started with Trusted Advisor Recommendations	106
Sign in to the Trusted Advisor console	107
View check categories	108
View specific checks	110

Filter your checks	111
Refresh check results	113
Download check results	114
Organizational view	114
Preferences	114
Get started with the Trusted Advisor API	116
Using Trusted Advisor as a web service	118
Get the list of available Trusted Advisor checks	118
Refresh the list of available Trusted Advisor checks	119
Poll a Trusted Advisor check for status changes	119
Request a Trusted Advisor check result	121
Show details of a Trusted Advisor check	122
Organizational view for AWS Trusted Advisor	123
Prerequisites	124
Enable organizational view	124
Refresh Trusted Advisor checks	125
Create organizational view reports	126
View the report summary	130
Download an organizational view report	131
Disable organizational view	136
Using IAM policies to allow access to organizational view	138
Using other AWS services to view Trusted Advisor reports	141
View Trusted Advisor checks powered by AWS Config	150
Troubleshooting	151
View your Security Hub CSPM controls in Trusted Advisor	152
Prerequisites	153
View your Security Hub CSPM findings	154
Refresh your Security Hub CSPM findings	155
Disable Security Hub CSPM from Trusted Advisor	156
Troubleshooting	156
Opt in AWS Compute Optimizer for Trusted Advisor checks	160
Related information	161
Get started with AWS Trusted Advisor Priority	161
Prerequisites	162
Enable Trusted Advisor Priority	163
View prioritized recommendations	163

Acknowledge a recommendation	166
Dismiss a recommendation	169
Resolve a recommendation	171
Reopen a recommendation	172
Download recommendation details	174
Register delegated administrators	175
Deregister delegated administrators	175
Manage Trusted Advisor Priority notifications	176
Disable Trusted Advisor Priority	177
Trusted Advisor check reference	177
Cost optimization	179
Performance	240
Security	290
Fault tolerance	338
Service limits	455
Operational Excellence	475
Change log for AWS Trusted Advisor	520
Older updates	528
New check: Amazon RDS Continuous Backup Not Enabled	529
New check: AWS CloudTrail Management Events Logging	529
Updated the Auto Scaling Group Resources check	530
Updated the IAM Access Analyzer External Access check	530
Added 1 new check	530
Updated 3 checks	530
Added 4 checks	531
Updated 3 checks	531
Added 9 new checks	531
Updated 1 Security check and added 1 Security check	532
Updated 6 Security checks	532
Updated 1 fault tolerance checks	532
Updated 9 checks	533
Removed 5 checks and added 1 check	533
Removed fault tolerance checks	534
New fault tolerance check	534
Updated fault tolerance and security checks	534
New fault tolerance check	534

Updated fault tolerance check	535
Updated security check	535
New security and performance checks	535
New security check	535
New fault tolerance and cost optimization checks	536
New fault tolerance checks	536
New checks for Amazon RDS	536
New AWS Trusted Advisor API	536
Trusted Advisor check removal	537
Integration of AWS Config checks into Trusted Advisor	537
New fault tolerance checks	537
New service limits check	538
New fault tolerance check	538
New fault tolerance and performance checks	538
New fault tolerance checks	538
New fault tolerance checks	539
Region Expansion of Amazon ECS Fault Tolerance Checks	539
New fault tolerance checks	539
New fault tolerance checks	535
Updates to the Trusted Advisor integration with AWS Security Hub CSPM	540
New fault tolerance checks for AWS Resilience Hub	536
Update to the Trusted Advisor console	541
New checks for Amazon EC2	541
Added Security Hub CSPM checks to Trusted Advisor	542
Added checks from AWS Compute Optimizer	542
Updates to the Exposed Access Keys check	542
Updated checks for AWS Direct Connect	543
AWS Security Hub CSPM controls added to the AWS Trusted Advisor console	544
New checks for Amazon EC2 and AWS Well-Architected	545
Updated check name for Amazon OpenSearch Service	545
Added checks for Amazon Elastic Block Store volume storage	546
Added checks for AWS Lambda	546
Trusted Advisor check removal	546
Updated checks for Amazon Elastic Block Store	547
Trusted Advisor check removal	548
Trusted Advisor check removal	549

AWS Support App in Slack	550
Prerequisites	551
Manage access to the AWS Support App widget	551
Manage access to the AWS Support App	553
Authorize a Slack workspace	559
Authorize multiple accounts	561
Configure a Slack channel	563
Update your Slack channel configuration	567
Create support cases in Slack	568
Reply to support cases in Slack	574
Join a live chat session with AWS Support	576
Search for support cases in Slack	582
Use your search results	584
Resolve support cases in Slack	585
Reopen support cases in Slack	586
Delete a Slack channel configuration from the AWS Support App	587
Delete a Slack workspace configuration from the AWS Support App	588
AWS Support App in Slack commands	589
Slack channel commands	589
Live chat channel commands	590
View AWS Support App correspondences in the AWS Support Center Console	590
Create AWS CloudFormation resources for the AWS Support App in Slack	591
AWS Support App and CloudFormation templates	591
Create Slack configuration resources for your organization	592
Learn more about CloudFormation	597
Create AWS Support App resources by using Terraform	597
Security	599
Data protection	600
Security for support cases	601
Support authorization	602
Concepts	603
Getting started	606
Configuring KMS keys	612
Managing support permits	616
Permit requests	621
Support actions	624

Monitoring	626
Quotas	632
Identity and access management	633
Audience	633
Authenticating with identities	634
Managing access using policies	635
How AWS Support works with IAM	636
Identity-based policy examples	638
Using service-linked roles	641
AWS managed policies	650
Manage access to AWS Support Center	732
Manage access to AWS Support Plans	740
Manage access to AWS Trusted Advisor	745
Troubleshooting	754
Incident response	756
Logging and monitoring in AWS Support and AWS Trusted Advisor	757
Compliance validation	757
Resilience	758
Infrastructure security	758
Configuration and vulnerability analysis	758
Code examples	760
Basics	761
Hello Support	761
Learn the basics	769
Actions	826
Scenarios	900
Getting started with technical support	900
Monitoring and logging for Support	910
Integrating AWS Support into EDAs	910
How EventBridge routes AWS Support events	911
AWS Support events	911
Creating event patterns	912
Support Case Update event	913
Logging AWS Support API calls with AWS CloudTrail	917
AWS Support information in CloudTrail	917
AWS Trusted Advisor information in CloudTrail logging	918

Understanding AWS Support log file entries	919
Logging AWS Support App API calls with CloudTrail	921
AWS Support App information in CloudTrail	921
Understanding AWS Support App log file entries	922
Monitoring and logging for Support Plans	927
Logging AWS Support Plans API calls with AWS CloudTrail	927
AWS Support Plans information in CloudTrail	928
Understanding AWS Support Plans log file entries	929
Logging console actions for changes to your Support plan	935
Monitoring and logging for Trusted Advisor	939
Monitoring Trusted Advisor check results with EventBridge	940
Creating CloudWatch alarms to monitor Trusted Advisor metrics	943
Prerequisites	943
CloudWatch metrics for Trusted Advisor	948
Trusted Advisor metrics and dimensions	954
Logging AWS Trusted Advisor console actions with AWS CloudTrail	955
Trusted Advisor information in CloudTrail	956
Example: Trusted Advisor Log File Entries	958
Troubleshooting resources	963
Service-specific troubleshooting	963
Document history	968
Earlier updates	1013

Getting started with AWS Support

AWS Support offers a range of plans that provide access to tools and expertise that support the success and operational health of your AWS solutions. All support plans provide round-the-clock access to customer service, AWS documentation, technical papers, and support forums. For technical support and more resources to plan, deploy, and improve your AWS environment, you can choose a support plan for your AWS use case.

Important

End of Support Notice: Developer Support will be discontinued January 1, 2027. Customers with Developer Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information, see [Business Support+ plan details](#)

End of Support Notice: Business Support will be discontinued January 1, 2027. Customers with Business Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information see, [Business Support+ plan details](#)

End of Support Notice: On January 1, 2027, AWS will discontinue Enterprise On-Ramp. Throughout 2026, Enterprise On-Ramp customers will be automatically upgraded to AWS Enterprise Support during contract renewal or in periodic batches. Customers will receive an email notification a month before their upgrade. No further action is required. Enterprise Support provides designated TAM assignment, 15-minute response times, and AWS Security Incident Response available at no additional cost, all at a lower \$5,000 minimum (reduced from \$15,000). For more information, see [AWS Enterprise Support plan details](#).

For more information, see [Developer, Business, and Enterprise On-Ramp end of support](#). Developer Support, Business Support, and Enterprise On-Ramp will remain available in the AWS GovCloud (US) Region.

Notes

- To create a Support interaction to use AI-generated troubleshooting to resolve your issue, and optionally create a support case, see [AI-enhanced troubleshooting in the Support Center Console](#).
- For more information about the different AWS Support plans, see [Compare AWS Support plans](#) and [Change AWS Support Plans](#).
- Support plans offer different response times for your support cases. See [Choosing an initial support case severity level](#) and [Understanding AWS Support response times](#).

Topics

- [AI-enhanced troubleshooting in the Support Center Console](#)
- [Virtual meetings with AWS Support](#)
- [Case management](#)
- [Request a service quota increase](#)
- [Legacy experience: Creating support cases and case management](#)
- [Using AWS Support with an AWS SDK](#)

AI-enhanced troubleshooting in the Support Center Console

AI-enhanced troubleshooting capabilities that help you resolve issues faster and more efficiently are available in supported AWS Regions. If you have an AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, then you can use the capabilities in the Support Center Console to troubleshoot technical issues as well as account and billing issues. If you have a Basic Support plan, you can use the capabilities in the Support Center Console for troubleshooting general questions and for assistance with account and billing issues. Using AI-enhanced troubleshooting streamlines the support experience by using contextual awareness and automated diagnostics to provide targeted solutions for your AWS environment.

AI-enhanced troubleshooting in the Support Center Console is supported in the following AWS Regions:

- US East (N. Virginia) Region
- US East (Ohio) Region

- Europe (Ireland) Region

Note

If you operate in an AWS Region that doesn't support AI-enhanced capabilities in the Support Center Console then you will use the legacy method of case management. For more information, see [Legacy experience: Creating support cases and case management](#).

When you access the Support Center Console, you can enter your issue description in natural language, import relevant Amazon Q conversations, receive generative AI troubleshooting guidance, and choose to create a support case with pre-populated fields, if needed.

You can provide contextual information about your environment and issue to receive personalized solutions throughout the troubleshooting process.

AI-enhanced troubleshooting in the AWS Support console provides the following key benefits:

- **Faster issue resolution:** Get immediate responses and relevant solutions as soon as you describe your problem.
- **Context preservation:** Import your previous Amazon Q conversations to maintain troubleshooting context.
- **Streamlined case creation:** Use natural language to describe issues instead of navigating multiple form fields.
- **Intelligent follow-up:** Receive relevant follow-up questions based on your specific AWS environment.
- **AI-powered guidance for account, billing, and technical issues:** AI-enhanced troubleshooting provides contextual recommendations across all supported issue types. For account and billing questions, AI-enhanced troubleshooting surfaces relevant responses, walks you through common resolutions, and helps you navigate account settings. For technical issues, it analyzes your environment context, asks targeted follow-up questions, and delivers step-by-step troubleshooting guidance informed by AWS best practices, helping you resolve issues without creating a support case.
- **Autonomous investigation for operational issues:** For operational issues such as performance degradation, elevated error rates, or post-deployment regressions, customers can invoke [AWS DevOps Agent](#) to autonomously query observability services in your account, identify root cause,

and generate a mitigation plan, streaming progress directly in your support interaction. On first use, the Support Center Console prompts you to get explicit consent to enable the DevOps Agent. To learn more, see [Using AWS DevOps Agent with your AWS Support plan](#).

For a complete list of the capabilities available in your Support plan, see [Compare AWS Support Plans](#).

Notes

- To change your support plan, see [Change AWS Support Plans](#).
- To close your account, see [Closing an account](#) in the *AWS Billing User Guide*.
- To find common troubleshooting topics for AWS services, see [Troubleshooting resources](#).
- If you're a customer of an AWS Partner that is part of the AWS Partner Network, and you use Resold Support, contact your AWS Partner directly for any billing related issues. AWS Support can't assist with non-technical issues for Resold Support, such as billing and account management. For more information, see the following topics:
 - [How AWS Partners can determine AWS Support plans in an organization](#)
 - [AWS Partner-Led Support](#)

Important

Before you open a support interaction or create a support case, check for events affecting your account resources by visiting your Health Dashboard at <https://phd.aws.amazon.com/phd/home#/>. You can verify any AWS service issues in the dashboard, though there might be a slight delay in posting incidents. If you're uncertain if there is an active incident, open a support case.

Topics

- [Set up permissions to use AI-enhanced troubleshooting](#)
- [Create a support interaction](#)
- [Create a support case from a support interaction](#)
- [View support interactions](#)

- [Managing attachments](#)
- [Troubleshooting](#)

Set up permissions to use AI-enhanced troubleshooting

To access AI-enhanced troubleshooting capabilities in Support Center, you need specific AWS Identity and Access Management permissions. This section describes the necessary IAM permissions and explains how to configure them so that you can fully use these capabilities.

AI-enhanced troubleshooting requires permissions beyond traditional support case management. The required permissions fall into four categories:

- **Support interaction permissions:** Enable the new interaction-based workflow in Support Center.
- **AI-powered classification permissions:** Allow access to AI-powered issue classification features.
- **Amazon Q integration permissions:** Enable conversation import from Amazon Q Developer.
- **AWS DevOps Agent permissions:** Enable operations investigations by DevOps Agent during a support interaction.

These permissions supplement your existing AWS Support permissions and don't replace them.

You can set up permissions for AI-enhanced troubleshooting in two ways:

[Option 1: Use the AWS managed policy \(recommended\)](#). Attach the `AWSSupportAccess` managed policy to your users or roles. This policy includes all required permissions and is automatically updated when new Support features are released.

[Option 2: Create a custom policy with minimum required permissions](#). This approach gives you more control but requires manual updates when new features are added.

Option 1: Use the AWS managed policy (recommended)

If you currently have the `AWSSupportAccess` managed policy attached, no additional permissions are required. However, to continue to use the functions included in the [Support Center Console API](#), you must add the Support Center Console operations to your IAM policies before November 2, 2026, if you don't already have them. To do this, update the AWS Support managed policy to include the `support-console:*` actions. For more information, see [Adding IAM policies for the Support Center Console API operations](#).

Option 2: Create a custom policy with minimum required permissions

You can explicitly allow-list specific actions instead of using wildcards. The following are the required permissions for support interactions, case creation, and case management:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "support:AddAttachmentsToSet",
        "support:AddCommunicationToCase",
        "support:CreateCase",
        "support:DescribeAttachment",
        "support:DescribeCaseAttributes",
        "support:DescribeCases",
        "support:DescribeCommunication",
        "support:DescribeCommunications",
        "support:DescribeCreateCaseOptions",
        "support:DescribeIssueTypes",
        "support:DescribeServices",
        "support:DescribeSeverityLevels",
        "support:DescribeSupportedLanguages",
        "support:DescribeSupportLevel",
        "support:GetInteraction",
        "support:InitiateCallForCase",
        "support:ListInteractionEntries",
        "support:ListInteractions",
        "support:InitiateChatForCase",
        "support:PutCaseAttributes",
        "support:ResolveCase",
        "support:ResolveInteraction",
        "support:SearchForCases",
        "support:StartInteraction",
        "support:UpdateInteraction",
        "support-console:GetAccountState",
        "support-console:GetAccountGovCloudEnabled",
        "support-console:GetCaseDraft",
        "support-console:CreateCaseDraft",
        "support-console>DeleteCaseDraft",
```

```
        "support-console:GetBanner",
        "support-console:DescribeDynamicHelp",
        "support-console:CreateContact"
    ],
    "Resource": "*"
}
]
```

For AWS DevOps Agent permissions required by your IAM identity, attach the following AWS managed policies:

- `AIDevOpsAgentFullAccess`. Provides full access to DevOps Agent management actions.
- `AIDevOpsAgentAccessPolicy`. Required to create the agent space.
- `AIDevOpsOperatorAppAccessPolicy`. Required to enable Operator App access.

For the full list of actions, see [DevOps Agent IAM permissions](#) in the *AWS DevOps Agent User Guide*.

Your IAM identity needs `iam:PassRole` on the `DevOpsAgentRole-AgentSpace` and `DevOpsAgentRole-WebappAdmin` roles so that the Support Center Console can pass these roles to DevOps Agent when creating them on your behalf during first-time setup. For investigations on Amazon Elastic Kubernetes Service clusters, your IAM identity also needs `eks:DescribeAccessEntry`, `eks:CreateAccessEntry`, and `eks:AssociateAccessPolicy` so that the Support Center Console can create the read-only access entry on the target cluster. For more information about access entries, see [Grant IAM users access to Kubernetes with Amazon EKS access entries](#) in the *Amazon Elastic Kubernetes Service User Guide*.

Note

Using a custom policy requires ongoing maintenance as AWS Support releases new features. For more information about the Support Center Console API operations, see [Adding IAM policies for the Support Center Console API operations](#). For more information about each of the Support API operations, see [Manage access to AWS Support Center](#).

Required permissions for Amazon Q integration

To use the Amazon Q conversation import feature in Support Center, IAM identities need permissions for the following Amazon Q Developer actions:

- `q:StartConversation`: Start a new conversation with Amazon Q.
- `q:SendMessage`: Send messages within a conversation.
- `q:GetConversation`: Retrieve conversation details. This action is required for console access.
- `q:ListConversations`: List available conversations. This action is required for console access and Support Center integration.

Amazon Q integration with Support Center Console specifically requires the `q:ListConversations` permission to display your recent conversations for import. For detailed guidance on configuring Amazon Q Developer permissions, see [Amazon Q Developer permissions reference](#) and [Manage access to Amazon Q Developer with policies](#).

Applying required permissions for support interactions

To apply permissions to your IAM users, complete the following steps:

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 **Policies**, then choose **Create policy**.
3. Choose the **JSON** tab and paste one of the policy documents mentioned in the previous sections.
4. Choose **Next: Tags**, then **Next: Review**.
5. Enter a policy name such as `SupportConsoleInteractionsAccess` and provide a description that explains the policy's purpose.
6. Choose **Create policy**.
7. Attach the policy to your IAM users, groups, or roles that need access to the Support Center.

If you have existing `AWSSupportAccess` managed policy attachments, then attach the supplementary custom policy alongside the managed policy.

Denying AI-enhanced troubleshooting operations

You can explicitly deny AI-enhanced troubleshooting operations by removing the following operations from your custom managed policy. Denying these operations prevents users from accessing the AI-powered interaction workflow and Amazon Q conversations in Support Center.

Support interaction operations

The following operations enable the interaction-based workflow in Support Center:

- `support:GetInteraction`
- `support:ListInteractions`
- `support:ResolveInteraction`
- `support:StartInteraction`
- `support:UpdateInteraction`

Amazon Q Developer conversation operations

The following operations enable Amazon Q conversations in Support Center:

- `q:StartConversation` – Start a new conversation with Amazon Q.
- `q:SendMessage` – Send messages within a conversation.
- `q:GetConversation` – Retrieve conversation details. This action is required for console access.
- `q:ListConversations` – List available conversations. This action is required for console access and Support Center integration.

The following example policy explicitly denies all AI-enhanced troubleshooting operations:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": [
        "support:GetInteraction",
        "support:ListInteractions",
        "support:ResolveInteraction",
        "support:StartInteraction",

```

```
    "support:UpdateInteraction",
    "q:StartConversation",
    "q:SendMessage",
    "q:GetConversation",
    "q:ListConversations"
  ],
  "Resource": "*"
}
```

Note

Denying these operations doesn't affect access to traditional support case management. You can continue creating and managing support cases using the standard workflow.

Create a support interaction

A support interaction is how you begin your engagement with AWS Support. You first describe your issue using natural language and receive assistance tailored to you using AI-enhanced troubleshooting. Your initial interaction might include clarifying questions, contextual solutions, and automated problem resolution, without creating a support case. These interactions might resolve issues independently or serve as the foundation for a support case looping in a human engineer, if needed.

Support interactions differ from support cases in that support cases include engagement with a Cloud Support Engineer. You can choose to automatically generate a support case based on a prior support interaction. The support case maintains all context from the initial support interaction and includes additional AI-generated insights to assist the Cloud Support Engineer with resolving your issue. This powerful combination of AI-enhanced troubleshooting and assistance from AWS Cloud Support Engineers potentially lead to faster issue resolution and reduced down time.

Notes

- You can sign in to the Support Center Console as an AWS Identity and Access Management (IAM) user. For more information, see [Manage access to AWS Support Center](#).

- If you can't sign in to the Support Center Console and create a support case, you can use the [Contact Us](#) page instead. You can use this page to get help with billing and account issues.

To start a support interaction, complete the following steps:

1. Sign in to the [AWS Support Center Console](#).

Tip

In the AWS Management Console, you can also choose the question mark icon



and then choose **Support Center**.

2. You have several options for starting your support interaction:
 - Enter details about the issues that you need assistance with. This is how you begin a new support interaction. Enter detailed information about your issue and any troubleshooting steps that you have already taken.
 - **Continue an existing support interaction:** Choose from a recent support interaction shown in the **Describe your issue or continue with** section. This section shows the two most recent support interactions. Access the **Viewing past support interactions** section to see additional past support interactions.
 - **Use a Amazon Q transcript:** Select the Amazon Q icon in the text field to see a list of recent Amazon Q conversations. The five most recent conversations from the AWS GovCloud (US-East) AWS Region are shown. Or, choose from a recent Amazon Q interaction shown in the **Describe your issue or continue with**. When you select a conversation, a summary of that conversation is generated and added to the text box. If you select an Amazon Q conversation, then you see a disclaimer regarding AWS Region and user accessibility.
3. Choose the **Send** icon in the bottom right of the text field.
4. AWS Support generative AI-powered troubleshooting analyzes your query along with your specific AWS environment. You might be prompted to provide additional information to assist with the analysis. If you see a prompt for additional information, enter the requested data, then choose **Submit**. If you don't know or don't have access to the requested information, then

you can skip this step and receive a general guidance response instead. A general guidance response isn't specific to your AWS environment.

For operational issues, such as performance degradation, elevated error rates, or post-deployment regressions, AI-enhanced troubleshooting might invoke [AWS DevOps Agent](#) to autonomously investigate your environment. DevOps Agent queries observability services in your account (Amazon CloudWatch, AWS CloudTrail, AWS X-Ray, AWS Config, and others) to identify root cause and propose a mitigation plan. If your account doesn't already have DevOps Agent configured, you are prompted to approve creation of the required resources before the investigation begins. Progress streams directly in your support interaction.

When the analysis is complete, you see a summary of the findings, along with remediation steps. For DevOps Agent investigations, the summary includes symptoms, root cause, investigation gaps, and a mitigation plan. To view the sources used in the analysis and remediation steps, choose **Sources**.

5. If you need further assistance, you can complete one of the following options:
 - **Continue with AI-assisted support:** To further refine the AI-assisted analysis and generate a new response, choose **Add more details for a better response**. Enter information in the **Additional details** field, and then choose **Submit**. Keep in mind that this option is for additional context for the original issue. If you need to enter context for a new issue, select **Start new interaction** at the top or bottom of the screen.
 - **Create a support case:** To create a support case with AWS Support, choose **Create a case**. This option starts the case creation workflow. Many of the case details are auto-populated for you based on your support interaction. You can change this information as needed. Your support interaction, including details of any resolution steps provided, are added to the support case. For details on how to create a support case, see [Create a support case from a support interaction](#).

At any time throughout the support interaction, you can use the **Thumbs up** and **Thumbs down** icons to provide feedback on your experience.

Create a support case from a support interaction

When you select **Create case** during your support interaction, a support case is created for you with many of the case details, such as the **Subject**, **Description**, **Case type** **Service**, **Category**, and **Severity level** populated for you. You can change this information as needed. Make sure that you

review this information for accuracy. For information on how to choose a severity level for your case, see [Choosing an initial support case severity level](#).

⚠ Important

Before you create a support case, check for events affecting your account resources by visiting your [Health Dashboard](#). You can verify any AWS service issues in the dashboard. There might be a slight delay in posting incidents. If you're uncertain if there is an active incident, open a support case.

After you choose **Create case**, enter or verify the following information:

1. Verify the **Subject** for this support case. The **Subject** is a brief synopsis of what your support interaction is about.
2. Verify the **Description**. Your initial inquiry appears in the **Description** field. Modify this information as needed. Make your description as detailed as possible. Include relevant resource information, along with anything else that might help us understand your issue.
3. (Optional) Choose **Choose file** to add any relevant files to your case, such as error logs or screenshots. You can attach up to 10 files. Each file can be up to 150 MB. For more information about managing attachments, see [Managing attachments](#).

ℹ Note

To attach more than 3 files or files over 5 MB, make sure that you have the correct permissions. For more information, see [Required permissions for large attachments](#).

4. For **Case type**, choose one of the following options:
 - **Account and billing**
 - **Technical**
 - **Service quotas**. You can only request certain types of service quota increases from the Support Center Console. For more information, see [Request a service quota increase](#).

 Note

If you have a Basic Support plan, then you can't create a **Technical** support case.

5. Verify the **Service**, **Category**, and **Severity**.
6. In the **Communication preference** section, indicate how you want AWS to communicate with you. You can choose one of the following options:
 - a. **Email:** Receive a response to your email.
 - b. **Phone:** Receive a phone call from a support agent. If you choose this option, enter the following information:
 - **Country or region**
 - **Phone number**
 - **(Optional) Extension**
 - c. **Chat:** Start a live chat with a support agent. If you can't connect to a chat, see [Troubleshooting](#).

(Optional) If you have an AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, you can enter up to 10 additional email addresses in the **Additional contacts**. You can enter the email addresses of people to notify when the status of the case changes. If you're signed in as an IAM user, include your email address. If you're signed in with your root account email address and password, you don't need to include your email address.

 Note

If you have the Basic Support plan, the **Additional contacts** option isn't available. However, the **Operations contact** specified in the **Alternate Contacts** section of the **My Account** page receives copies of the case correspondence, but only for the specific case types of account and billing.

7. When you're ready to submit the support case, select **Submit**. You are directed to the **Case details** page where you can see your case details, the support interaction, and the case correspondences.

Select **Case details** to view the information about your case, such as attachments, or severity level. Select **Support interactions** to see the support interactions associated with this case.

Best practices and target response times for Enterprise Support cases

Before you create a support case, check for events affecting your account resources by visiting your [Health Dashboard](#). You can verify any AWS service issues in the dashboard. There might be a slight delay in posting incidents. If you're uncertain if there is an active incident, open a support case.

When creating your support case, make sure that you choose the correct severity level and provide as much information as possible. For more information on what kind of information to provide, see the following **Enterprise Support case best practices** section.

Use the following matrix to help you identify the correct severity.

AWS ITIL Case Severity Matrix				
		Urgency		
		Low	Medium	High
Impact	Low	Low	Normal	High
	Medium	Normal	High	Urgent
	High	High	Urgent	Critical

For Critical or Urgent cases, make sure that you leave contact details so that we can reach out to you, if needed.

Note

The quickest way to get help is to use the **Chat** or **Phone** contact method.

Topics

- Enterprise Support case target response times
- Enterprise Support best practices
- Enterprise Support escalation path

Enterprise Support case target response times

The following graphic displays the target AWS Support response times for different case severity levels.

ITIL Priority	Severity Level	How to Raise		Examples for reference
Critical	Business Critical System Down		15 minutes	Major outage; loss of entire enterprise or customer base. Has potential for large revenue loss or business risk.
Urgent	Production System Down		1 hour	Business significantly impacted, where important functions or applications aren't available.
High	Production System Impaired		4 hours	Critical functions of your applications are impaired or degraded.
Normal	System Impaired		12 hours	Non-critical applications are behaving abnormally; time-sensitive development question.
Low	General Guidance		24 hours	General development/ service questions or want to submit feature request.

Enterprise Support case best practices

Submit your support case at the appropriate severity level.

- The appropriate severity level helps make sure that your case is visible across AWS and allows the TAM to manage your issue.
- Only submit one case for each event or issue.
- Make sure that you raise cases in the correct AWS account.

Provide as much detailed information as possible, answering the relevant questions provided in the following list:

- **Who:** Who did what? Who is affected? Who should be looped in? Who have you already talked to?
- **What:** What happened, exactly. What's the impact or blast radius? What have you already tried to resolve the issue?
- **When:** When did or does it happen (date, time, time zone)? When do you need an answer by?
- **Where:** AWS Region, Availability Zone, specific instance or resource IDs, and other identifiers.
- **Why:** Why are you opening this case (information, limit raise, event-analysis or RCA, outage)?
- **How:** Include information on how to reproduce the problem, how to escalate, and how to contact you.

 Note

You can change the severity of your support case after you create it. For information, see [Changing the severity level of your support case](#).

Enterprise Support business critical or production system down

- For Business-critical system down or Production system down cases use **Chat** or **Phone** and ensure there are people available at your organization to work the case.
- Clearly describe the issue including what you've tried, what you expect and context. Summarize the business impact.
- Provide (or start to capture) as many metrics, timings and symptoms as you can. More data means faster diagnosis.
- Provide a conference bridge. When opening a Business-critical system down issue, provide a conference bridge for the support teams to join to help troubleshoot the problem. As Business-critical system down issues are production down, it's best to get everyone on the phone and come to a common plan of action towards a resolution.
- Ensure that all activity related to the event is captured in the support case (update from Console, not email). The phone option is good for real-time communications; however, make sure that you record updates and outcomes in the support case.

Enterprise Support escalation path

You can follow the following step-by-step process when facing an issue.

1. **Check the Health Dashboard** if you think that the issue might be related to abnormal AWS service operations.
2. **Assess case information** as outlined in the Support Case Best Practices section above.
3. **Select the appropriate severity** and use **Phone** or **Chat** if quick response is needed.
4. **Contact your TAM** if you need additional assistance or if you are not getting response as expected.

View support interactions

Past interactions with AWS Support are saved for 10 years. You can view past interactions from your Support Dashboard by selecting the **List** icon. Then, choose the interaction that you want to view. The AWS Support interaction details appear. If you choose to create a support case from an interaction, then the interaction no longer appears in your past interactions list. The interaction now displays in the **Case details** page of the associated support case.

You can add additional details to the interaction to generate a new response. Or, you can choose to create a Support case from the interaction by choosing **Create a case** on the AWS Support interaction details screen.

Managing attachments

You can attach files to a support case when you create a case or when you reply to an existing case. You can attach up to 10 files, with each file up to 150 MB. Supported attachments include error logs, screenshots, configuration files, and other documents relevant to your issue.

The attachments table in the Support Center Console displays the following information for each file:

- **File name** – The name of the attached file.
- **Size** – The file size.
- **Status** – The current upload status of the file.
- **Action** – Available actions for the file, such as **Cancel**, **Retry**, or **Remove**.

Note

Files that are removed or not submitted with a case are automatically deleted after 24 hours.

Topics

- [Required permissions for large attachments](#)
- [Uploading attachments](#)
- [Cancelling an in-progress upload](#)
- [Retrying a failed upload](#)

- [Removing an attachment](#)
- [Downloading an attachment](#)

Required permissions for large attachments

To attach files larger than 5 MB or more than 3 files to a support case, your AWS Identity and Access Management (IAM) identity must have the correct permissions.

The following table shows the required IAM actions for each attachment operation.

Operation	Required IAM actions	Grants access to
Upload attachments	<code>support:UploadAttachment</code>	<code>GetAttachmentUploadLinks</code>
	<code>support:AddAttachmentsToSet</code>	<code>CompleteAttachmentUpload</code>
		<code>DescribeAttachmentUploadStatus</code>
Download attachments	<code>support:DownloadAttachment</code>	<code>GetAttachmentDownloadLink</code>
	<code>support:DescribeAttachment</code>	

 Important

All listed permissions are required for each operation. For example, to upload attachments, you must have both `support:UploadAttachment` and `support:AddAttachmentsToSet`.

The following example policy grants full upload and download permissions for large attachments:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "support:UploadAttachment",
        "support:AddAttachmentsToSet",
        "support:DownloadAttachment",
        "support:DescribeAttachment"
      ],
      "Resource": "*"
    }
  ]
}
```

Note

If you have the `AWSSupportAccess` managed policy attached to your IAM identity, no additional permissions are required. The managed policy includes all actions needed for large attachments.

If you don't have the required permissions, then the Support Center Console supports attaching up to 3 files at 5 MB each. An **Update your permissions** notice appears with details about the missing permissions.

Uploading attachments

You can attach files when creating a support case or when replying to an existing case. You can attach up to 10 files, with each file up to 150 MB.

Support Center Console

To upload attachments using the console

1. On the **Create case** page or in the reply section of a case, choose **Choose file**.
2. Select up to 10 files from your file picker. Each file can be up to 150 MB. Each file begins uploading immediately and the **Status** column displays a real-time progress bar during the upload.

3. When a file finishes uploading, the status changes to **Upload completed** and the **Action** column displays **Remove**.
4. After all files reach **Upload completed** status, complete the remaining case fields and then choose **Submit**.

After you submit the case, uploaded attachments appear as clickable links in the case correspondence on the **Case details** page.

AWS Support API

You can programmatically attach files to a support case or case communication using the AWS Support API. The API uses a multipart upload flow with presigned URLs.

To upload an attachment using the API

1. **Start the upload.** Call `GetAttachmentUploadLinks` with `fileName` and `fileSizeBytes`. The response returns an `uploadId`, `partSizeBytes`, `totalParts`, and a first batch of presigned upload URLs.
2. **Split and upload your file.** Divide your file into segments of `partSizeBytes` bytes by reading sequentially from the start of the file. The first segment corresponds to `partIndex` 1, the second to `partIndex` 2, and so on. The final segment may be smaller than `partSizeBytes`. PUT each segment to the presigned URL with the matching `partIndex`. Save the ETag that Amazon S3 returns for each successful PUT.
3. **Get more URLs if needed.** Each call returns up to 10 URLs. If your file has more than 10 parts, call `GetAttachmentUploadLinks` again with the same `uploadId`, passing `nextIndex` as `uploadRange.startIndex`. Repeat until `nextIndex` is null.
4. **Report completed parts.** Call `CompleteAttachmentUpload` with the `uploadId`, `partIndex`, and ETag for each part as soon as its PUT succeeds. Reporting parts immediately makes the upload resumable. If the upload is interrupted, only unreported parts need to be re-uploaded.
5. **Wait for the upload to be ready.** Poll `DescribeAttachmentUploadStatus` until `uploadStatus` is `attachment-ready`. Processing is asynchronous and may take a few minutes for larger files.
6. **Attach to a case.** Pass the `uploadId` in the `uploadIds` parameter of `CreateCase` or `AddCommunicationToCase`.

uploadId vs. attachmentId

An `uploadId` identifies an in-progress upload session and is used with the upload operations (`GetAttachmentUploadLinks`, `CompleteAttachmentUpload`, `DescribeAttachmentUploadStatus`) and as input to `CreateCase` or `AddCommunicationToCase`.

An `attachmentId` is generated after the upload is committed to a case communication. Use `GetAttachmentDownloadLink` with an `attachmentId` to download attachments. You can obtain the `attachmentId` from the `attachments` field of a `Communication` returned by `DescribeCommunications`.

Note

`DescribeAttachment` only supports attachments that are 5 MB or smaller. Calling `DescribeAttachment` for a larger attachment returns `InvalidParameterValueException`. To download attachments larger than 5 MB, use `GetAttachmentDownloadLink`.

Endpoints

Presigned URLs returned by these operations are served from the following domains:

- **Upload URLs** (`GetAttachmentUploadLinks`): `v1.uploads.attachments.support.us-east-1.amazonaws.com`
- **Download URLs** (`GetAttachmentDownloadLink`):
`v1.downloads.attachments.support.us-east-1.amazonaws.com`

If your environment restricts outbound traffic, ensure both domains are reachable.

Constraints

- Maximum file size per upload: 150 MB.
- Maximum upload URLs returned per call to `GetAttachmentUploadLinks`: 10.
- Maximum number of `uploadIds` per `CreateCase` or `AddCommunicationToCase` request: 10.

Cancelling an in-progress upload

You can cancel a file upload while it is in progress.

To cancel an upload

1. While a file is uploading, the **Status** column shows a progress bar and the **Action** column shows **Cancel**.
2. Choose **Cancel** in the **Action** column.
3. The upload is aborted and the file status changes to **Failed to upload**. You can then retry or remove the file.

Note

Any partially uploaded data is automatically cleaned up. No additional action is required.

Retrying a failed upload

If a file upload fails, you can retry the upload.

To retry a failed upload

1. In the attachments table, identify the file showing **Failed to upload** in the **Status** column.
2. Choose **Retry** in the **Action** column.
3. The upload restarts. Monitor the **Status** column for progress.

Important

You cannot submit a case while any attachment has a **Failed to upload** status. An error banner appears if you attempt to submit. You must either retry or remove the failed attachment before submitting.

Removing an attachment

You can remove an attachment that has been uploaded, has failed to upload, or that you no longer want to include with the case.

To remove an attachment

1. In the attachments table, locate the file you want to remove.
2. Choose **Remove** in the **Action** column.
3. The file is removed from the table immediately.

Note

Removed files are not immediately deleted from storage. They are automatically deleted after 24 hours.

Downloading an attachment

After a case is submitted, you can download attachments from the case correspondence.

To download an attachment

1. Navigate to the **Case details** page for the support case.
2. In the **Correspondence** section, locate the attachment link.
3. Choose the filename link. The file downloads directly to your device.

Note

You must have the required IAM permissions to download attachments. If you don't have the correct permissions, an **Access denied** error appears when you choose the attachment link.

Troubleshooting

If you have difficulty when you create or manage your support case, see the following troubleshooting information.

I want to reopen a live chat for my case

You can reply to your existing support case to open another chat window. For more information, see [Updating an existing support case](#).

I can't connect to a live chat

If you chose the **Chat** option but you can't connect to the chat window, first perform the following checks:

- Ensure that you've configured your browser to allow pop-up windows in Support Center.

Note

Review the settings for your browser. For more information, see the [Chrome Help](#) and [Firefox Support](#) websites.

- Ensure that you've configured your network so that you can use AWS Support:
 - Your network can access the `*.connect.us-east-1.amazonaws.com` endpoint.

Note

For AWS GovCloud (US), the endpoint is `*.connect-fips.us-east-1.amazonaws.com`.

- Your firewall supports web socket connections.

If you still can't connect to the chat window, contact AWS Support using email or phone contact options.

Virtual meetings with AWS Support

Virtual meetings enable you to connect with AWS Support engineers through video calls with screen sharing capabilities. This feature helps you resolve complex technical issues that require visual demonstration or real-time collaboration.

When a support engineer determines that your case requires visual assistance, they can initiate a virtual meeting. You receive a meeting invitation on your case details page in the AWS Support

Center. After you accept the invitation, you join a secure video call where you can share your screen and collaborate with the support engineer.

Virtual meetings integrate with Connect Customer and use WebRTC technology to provide secure, browser-based video conferencing without requiring additional software installation.

Virtual meetings are available in the [commercial AWS Regions](#) only.

Topics

- [Join a virtual meeting](#)
- [Security and privacy during your virtual meeting](#)
- [Required IAM permissions for virtual meetings](#)
- [Troubleshooting virtual meetings](#)

Join a virtual meeting

Virtual meetings are initiated by the AWS Support engineer assisting you with your support case. To join a virtual meeting, complete the following steps.

1. When a support engineer initiates a virtual meeting, you see a meeting invitation on the **Case details** page in your support case. To join the meeting, choose **Join virtual meeting** to accept the invitation.

Meeting invitations expire after 10 minutes. If you don't join within this time, request a new meeting from your support engineer.

2. Grant browser permissions for camera and microphone access when prompted.

The virtual meeting opens in a new window and you're connected to the support engineer. You can mute and unmute your microphone, or disconnect from the meeting, using the buttons at the bottom of the screen.

Security and privacy during your virtual meeting

Virtual meetings use the same authentication and authorization mechanisms as other AWS Support operations. The following security measures protect your meetings:

- **Case ownership validation:** You can only join meetings for cases that belong to your AWS account.

- **AWS Identity and Access Management (IAM) based access control:** You must have the appropriate IAM permissions to join virtual meetings.
- **Encrypted connections:** All meeting data is transmitted over encrypted WebRTC connections.
- **Audit logging:** All meeting activities are logged in AWS CloudTrail for compliance and auditing purposes.

Important

Virtual meetings are recorded for quality assurance and training purposes. Don't share sensitive information such as passwords or access keys during the meeting.

Required IAM permissions for virtual meetings

To join virtual meetings, your IAM user or role must have the following permission:

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

For more information about AWS Support permissions, see [Manage access to AWS Support Center](#).

Troubleshooting virtual meetings

I can't see the meeting invitation.

Verify that your support engineer has initiated the meeting. Refresh the case details page. If the invitation still doesn't appear, contact your support engineer through the case correspondence.

The meeting invitation expired.

Meeting invitations expire after 10 minutes for security reasons. Request a new meeting invitation from your support engineer.

I'm experiencing connection issues.

Check your internet connection. Ensure that your firewall or network security settings allow WebRTC traffic. Try using a different browser or network connection.

I receive an authorization error.

Verify that your IAM user or role has permissions for the `support:InitiateLiveContactForCase` action.

Case management

Note

You can revert to the legacy method of case management by choosing **Use the old experience** in the banner at the top of the Support Center console. For more information, see [Legacy experience: Creating support cases and case management](#).

In the AWS Management Console, you can create three types of customer cases in Support:

- **Account and billing** support cases are available to all AWS customers. You can get help with billing and account questions.
- **Service limit increase** requests are available to all AWS customers. For more information about the default service quotas, formerly referred to as limits, see [AWS service quotas](#) in the *AWS General Reference*.
- **Technical** support cases connect you to technical support for help with service-related technical issues and, in some cases, third-party applications. If you have Basic Support, you can't create a technical support case.

Notes

- To change your support plan, see [Change AWS Support Plans](#).
- To close your account, see [Closing an Account](#) in the *AWS Billing User Guide*.

- To find common troubleshooting topics for AWS services, see [Troubleshooting resources](#).
- If you're a customer of an AWS Partner that is part of the AWS Partner Network, and you use Resold Support, contact your AWS Partner directly for any billing related issues. AWS Support can't assist with non-technical issues for Resold Support, such as billing and account management. For more information, see the following topics:
 - [How AWS Partners can determine AWS Support plans in an organization](#)
 - [AWS Partner-Led Support](#)

Describing your problem

Make your description as detailed as possible. Include relevant resource information, along with anything else that might help us understand your issue. For example, to troubleshoot performance, include timestamps and logs. For feature requests or general guidance questions, include a description of your environment and purpose.

When you provide as much detail as possible, you increase the chances that your case can be resolved quickly.

Choosing an initial support case severity level

When creating a support case make sure that the right severity is defined and that you provide as much information as possible.

You might want to create a support case at the highest severity that your support plan allows. However, it's a best practice to choose the highest severity only for cases that can't be worked around or that directly affect production applications. For information about building your services so that losing a single resource doesn't affect your applications, see the [Building Fault-Tolerant Applications on AWS](#) technical paper.

The following table lists the severity levels, response times, and example problems.

Notes

- If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, you can reassign your support case severity level to reflect changes

to urgency and business impact. For example, you can change your support case from **System impaired** to **Production system impaired**. When you change the case severity, AWS Support receives notification and routes the case according to the new severity level. For more information, see [Changing the severity level of your support case](#).

- If you have a Basic Support plan, then you can't change the severity level for a support case after you create it. If your situation changes, work with the Support agent.
- For more information about the severity level, see the [AWS Support API Reference](#).

Severity	Severity level code	First-response time	Description and support plan
General guidance	low	24 hours	You have a general development question, or you want to request a feature. (AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan)
System impaired	normal	12 hours	Non-critical functions of your application are behaving abnormally, or you have a time-sensitive development question. (AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan)
Production system impaired	high	4 hours	Important functions of your application are impaired or degraded. (AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan)
Production system down	urgent	1 hour	Your business is significantly impacted. Important functions of your application aren't available. (AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan)
Business-critical system down	critical	• AWS Business	Your business is at risk. Critical functions of your application aren't available.

Severity	Severity level code	First-response time	Description and support plan
		Support +: Less than 30 minutes	
		• AWS Enterprise Support: Less than 15 minutes • AWS Unified Operations: 5 minutes from an Incident Manager or Engineer	

Understanding AWS Support response times

AWS Support makes every reasonable effort to respond to your initial request within the indicated timeframe. For information about the scope of support for each Support plan, see [AWS Support features](#).

If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, you have round-the-clock access for technical support.

 Note

If you choose Japanese as your preferred contact language for support cases, support in Japanese may be available as follows:

- If you need customer service for non-technical support cases, support in Japanese is available during business hours in Japan defined as 09:00 AM to 06:00 PM Japan Standard Time (GMT+9), excluding holidays and weekends.
- If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, technical support is available all day, every day in Japanese.

If you choose Chinese as your preferred contact language for support cases, support in Chinese might be available as follows:

- If you need customer service for non-technical support cases, support in Chinese is available 09:00 AM to 06:00 PM (GMT+8), excluding holidays and weekends.
- If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, technical support is available all day, every day in Chinese.

If you choose Korean as your preferred contact language for support cases, support in Korean may be available as follows:

- If you need customer service for non-technical support cases, support in Korean is available during business hours in Korea defined as 09:00 AM to 06:00 PM Korean Standard Time (GMT+9), excluding holidays and weekends.
- If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, technical support is available all day, every day in Korean.

If you choose Turkish as your preferred contact language for support cases, support in Turkish may be available as follows:

- If you need customer service for non-technical support cases, support in Turkish is available during business hours in Türkiye defined as 09:00 AM to 06:00 PM (GMT+3), excluding holidays and weekends.

- If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, technical support is available all day, every day in English.

Changing the severity level of your support case

If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, you can reassign your support case severity level to reflect changes to urgency and business impact. For example, you can change your support case from **System impaired** to **Production system impaired**. When you change the case severity, AWS Support receives notification and attends to the case according to the new severity level.

Note

Japanese (JP) account or billing, Service Quota Increase Request (SQIR), and Turkish (TR) account or billing cases created in these languages have a default severity and can't be changed.

To change the severity of a support case, complete the following steps:

1. Sign in to the [AWS Support Center Console](#).

Tip

In the AWS Management Console, you can also choose the question mark icon



and then choose **Support Center**.

2. Select the case that you want to change the severity level for.
3. In **Case details**, choose the pencil icon next to the **Severity** field, as shown in the following example.

Case details

Subject		Status	Unassigned
Case ID		Severity	General guidance 
Created	2025-04-24T08:49:10.815Z	Category	-
Case type	Technical	Language	Portuguese

4. For **Severity**, choose the new severity level from the following options:
 - General guidance
 - System impaired
 - Production system impaired
 - Production system down
 - Business-critical system down
5. For **Reason for case severity change**, choose from the available options for why you're changing the case severity.
6. (Optional) For **Tell us more**, enter additional information about this change.
7. Do one of the following:
 - If you're lowering the support case severity, or if you're raising it from **General guidance** to **System impaired** or **Production system impaired**, choose **Update**.
 - If you're raising the severity to **Production system down** or **Business-critical system down**, use one of the options in the **Contact methods** section to engage with AWS Support, and then choose **Update**. The following example shows the options available in the **Contact methods** section.

Change case severity ✕

Severity

Production system down ▼

Reason for case severity change

- ▼

Tell us more - optional

Tell us more about the reason for the escalation

Contact methods [Info](#)

For high severity issues we recommend using our live contact mechanisms

Web ☐
We'll respond by email and Support Center.

Chat(Recommended) ☒
Chat online with a representative.

Phone(Recommended) ☐
We'll call you back at your number.

[Cancel](#) [Update](#)

Note

- If you upgrade your support case severity to **Production system down** or **Business-critical system down**, you must wait 60 minutes before you can change the severity again.
- If your support case is currently set to **Business-critical system down**, you're prompted to initiate live contact with AWS Support instead of assigning a higher severity.
- If you're raising your support case severity level after already raising it at least once, you might encounter a waiting period. For example, if you change the severity from **System impaired** to **Production system impaired** at 6:00 AM, then your support case falls under the 4-hour first-response time for the **Production system impaired** severity level. In this scenario, you can upgrade the severity level again at 10:00 AM, after the 4-hour window. For a list of first-response times for each severity level, see the table in [Understanding AWS Support response times](#).

Request a service quota increase

You can request increases to your service quotas (formerly referred to as limits) to support your workload requirements.

Use the Service Quotas service to request increases directly for your services. For more information, see the following documentation:

- [What is Service Quotas?](#) in the *Service Quotas User Guide*
- [Requesting a quota increase](#) in the *Service Quotas User Guide*

At this time, Service Quotas doesn't support service quotas for all AWS services in all AWS Regions. If your AWS service or AWS Region isn't available in the [Service Quotas console](#), complete the following steps to create a support case to request the quota increase:

1. Sign in to the [AWS Support Center Console](#).

Tip

In the AWS Management Console, you can also choose the question mark icon



and then choose **Support Center**.

2. On the **Support interactions** page, enter details about this service limit increase. When prompted, choose **Create a case**. Many of the Support case fields will be pre-populated with the text that you entered during your interaction. You can edit these fields as needed. For additional details on creating a support interaction, see [Create a support interaction](#).
3. For **Case type**, select **Service quotas**.
4. For **Service**, select **Service Limit increase**.
5. For **Category**, select the type of increase that you're requesting from the list. Only service limit increase requests available in Support Center are listed here. For other types of service limit requests, see [Requesting a quota increase](#) in the *Service Quotas User Guide*.
6. (Optional) From the **Preferred contact language** drop down, select the language that you want AWS Support to use when corresponding with you.
7. For **Region**, select the AWS Region where you're requesting the increase.

 Note

AWS Region selection isn't available if you selected **General** as the **Category**.

8. (Optional) To request multiple limit increases, choose **Add another limit** and then choose another AWS Region.
9. Enter a **Description** for this service quota increase or multiple increases. You can attach files, if necessary.
10. Choose **Next step: Solve now or contact us**.
11. For **Contact options**, choose one of the following options:
 - **Web** – Receive a reply in Support Center.
 - **Chat** – Start a live chat with a support agent. If you can't connect to a chat, see [Troubleshooting](#).
 - **Phone** – Receive a phone call from a support agent. If you choose this option, enter the following information:
 - **Country/Region**
 - **Phone number**
 - **(Optional) Extension**
12. When you're ready to submit the support case, select **Submit**. You are directed to the **Case details** page where you can see your case details, the support interaction, and the case correspondences.

Select **Case details** to view the information about your case, such as attachments, or severity level. Select **Support interactions** to see the support interactions associated with this case.

Legacy experience: Creating support cases and case management

In the AWS Management Console, you can create three types of customer cases in Support:

- **Account and billing** support cases are available to all AWS customers. You can get help with billing and account questions.

- **Service limit increase** requests are available to all AWS customers. For more information about the default service quotas, formerly referred to as limits, see [AWS service quotas](#) in the *AWS General Reference*.
- **Technical** support cases connect you to technical support for help with service-related technical issues and, in some cases, third-party applications. If you have Basic Support, you can't create a technical support case.

Notes

- To change your support plan, see [Change AWS Support Plans](#).
- To close your account, see [Closing an Account](#) in the *AWS Billing User Guide*.
- To find common troubleshooting topics for AWS services, see [Troubleshooting resources](#).
- If you're a customer of an AWS Partner that is part of the AWS Partner Network, and you use Resold Support, contact your AWS Partner directly for any billing related issues. AWS Support can't assist with non-technical issues for Resold Support, such as billing and account management. For more information, see the following topics:
 - [How AWS Partners can determine AWS Support plans in an organization](#)
 - [AWS Partner-Led Support](#)

Creating a support case

You can create a support case in the Support Center of the AWS Management Console.

Notes

- You can sign in to Support Center as an AWS Identity and Access Management (IAM) user. For more information, see [Manage access to AWS Support Center](#).
- If you can't sign in to Support Center and create a support case, you can use the [Contact Us](#) page instead. You can use this page to get help with billing and account issues.

To create a support case

1. Sign in to the [AWS Support Center Console](#).

Tip

In the AWS Management Console, you can also choose the question mark icon



and then choose **Support Center**.

2. Choose **Create case**.
3. Choose one of the following options:
 - **Account and billing**
 - **Technical**
 - For service quota increases, choose **Looking for service quota increases?** and then follow the instructions for [Request a service quota increase](#).
4. Choose the **Service**, **Category**, and **Severity**.

Tip

You can use the recommended solutions that appear for commonly asked questions.

5. Choose **Next step: Additional information**
6. On the **Additional information** page, for **Subject**, enter a title about your issue.
7. For **Description**, follow the prompts to describe your case, such as the following:
 - Error messages that you received
 - Troubleshooting steps that you followed
 - How you're accessing the service:
 - AWS Management Console
 - AWS Command Line Interface (AWS CLI)
 - API operations
8. (Optional) Choose **Attach files** to add any relevant files to your case, such as error logs or screenshots. You can attach up to three files. Each file can be up to 5 MB.
9. Choose **Next step: Solve now or contact us**.
10. On the **Contact us** page, choose your preferred language.
11. Choose your preferred contact method. You can choose one of the following options:

- a. **Web** – Receive a reply in Support Center.
- b. **Chat** – Start a live chat with a support agent. If you can't connect to a chat, see [Troubleshooting](#).
- c. **Phone** – Receive a phone call from a support agent. If you choose this option, enter the following information:
 - **Country or region**
 - **Phone number**
 - **(Optional) Extension**

Notes

- The contact options that appear depend on the type of case and your support plan.
- You can choose **Discard draft** to clear your support case draft.

12. (Optional) If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, the **Additional contacts** option appears. You can enter the email addresses of people to notify when the status of the case changes. If you're signed in as an IAM user, include your email address. If you're signed in with your root account email address and password, you don't need to include your email address

Note

If you have the Basic Support plan, the **Additional contacts** option isn't available. However, the **Operations** contact specified in the **Alternate Contacts** section of the [My Account](#) page receives copies of the case correspondence, but only for the specific case types of account and billing, and technical.

13. Review your case details and then choose **Submit**. Your case ID number and summary appear.

Describing your problem

Make your description as detailed as possible. Include relevant resource information, along with anything else that might help us understand your issue. For example, to troubleshoot performance, include timestamps and logs. For feature requests or general guidance questions, include a

description of your environment and purpose. In all cases, follow the **Description Guidance** that appears on your case submission form.

When you provide as much detail as possible, you increase the chances that your case can be resolved quickly.

Choosing an initial support case severity level

You might want to create a support case at the highest severity that your support plan allows. However, it's a best practice to choose the highest severity only for cases that can't be worked around or that directly affect production applications. For information about building your services so that losing a single resource doesn't affect your applications, see the [Building Fault-Tolerant Applications on AWS](#) technical paper.

The following table lists the severity levels, response times, and example problems.

 **Notes**

- If you have Enterprise Support or an Enterprise On-Ramp plan, you can reassign your support case severity level to reflect changes to urgency and business impact. For example, you can change your support case from **System impaired** to **Production system impaired**. When you change the case severity, AWS Support receives notification and routes the case according to the new severity level. For more information, see [Changing the severity level of your support case](#).
- If you don't have Enterprise support or an Enterprise On-Ramp plan, then you can't change the severity level for a support case after you create it. If your situation changes, work with the Support agent for your support case.
- For more information about the severity level, see the [AWS Support API Reference](#).

Severity	Severity level code	First-response time	Description and support plan
General guidance	low	24 hours	You have a general development question, or you want to request a feature. (*Develop

Severity	Severity level code	First-response time	Description and support plan
			er, AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan)
System impaired	normal	12 hours	Non-critical functions of your application are behaving abnormally, or you have a time-sensitive development question. (*Developer, AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan)
Production system impaired	high	4 hours	Important functions of your application are impaired or degraded. (AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan)
Production system down	urgent	1 hour	Your business is significantly impacted. Important functions of your application aren't available. (AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan)
Business-critical system down	critical	15 minutes	Your business is at risk. Critical functions of your application aren't available (Enterprise Support plan). Note that this is 30 minutes for the Enterprise On-Ramp Support plan.

Understanding AWS Support response times

AWS Support makes every reasonable effort to respond to your initial request within the indicated timeframe. For information about the scope of support for each Support plan, see [AWS Support features](#).

If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, you have round-the-clock access for technical support. *For Developer Support, response targets for support cases are calculated in business hours. Business hours are generally defined as 08:00 to

18:00 in the customer country, excluding holidays and weekends. These times can vary in countries with multiple time zones. The customer country information appears in the **Contact Information** section of the [My Account](#) page in the AWS Management Console.

Note

If you choose Japanese as your preferred contact language for support cases, support in Japanese may be available as follows:

- If you need customer service for non-technical support cases, or if you have a Developer Support plan and need technical support, support in Japanese is available during business hours in Japan defined as 09:00 AM to 06:00 PM Japan Standard Time (GMT+9), excluding holidays and weekends.
- If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, technical support is available all day, every day in Japanese.

If you choose Chinese as your preferred contact language for support cases, support in Chinese may be available as follows:

- If you need customer service for non-technical support cases, support in Chinese is available 09:00 AM to 06:00 PM (GMT+8), excluding holidays and weekends.
- If you have a Developer Support plan, technical support in Chinese is available during business hours generally defined as 8:00 AM to 6:00 PM in your country as set in [My Account](#), excluding holidays and weekends. These times may vary in countries with multiple time zones.
- If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, technical support is available all day, every day in Chinese.

If you choose Korean as your preferred contact language for support cases, support in Korean may be available as follows:

- If you need customer service for non-technical support cases, support in Korean is available during business hours in Korea defined as 09:00 AM to 06:00 PM Korean Standard Time (GMT+9), excluding holidays and weekends.
- If you have a Developer Support plan, technical support in Korean is available during business hours generally defined as 8:00 AM to 6:00 PM in your country as set in [My](#)

[Account](#), excluding holidays and weekends. These times may vary in countries with multiple time zones.

- If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, technical support is available all day, every day in Korean.

Changing the severity level of your support case

If you have Enterprise Support or an Enterprise On-Ramp plan, you can reassign your support case severity level to reflect changes to urgency and business impact. For example, you can change your support case from **System impaired** to **Production system impaired**. When you change the case severity, AWS Support receives notification and attends to the case according to the new severity level.

Note

Japanese (JP) account or billing, Service Quota Increase Request (SQIR), and Turkish (TR) account or billing cases created in these languages have a default severity and can't be changed.

To change the severity of a support case, complete the following steps:

1. Sign in to the [AWS Support Center Console](#).

Tip

In the AWS Management Console, you can also choose the question mark icon



and then choose **Support Center**.

2. Select the case that you want to change the severity level for.
3. In **Case details**, choose the pencil icon next to the **Severity** field, as shown in the following example.

Case details

Subject		Status	Unassigned
Case ID		Severity	General guidance 
Created	2025-04-24T08:49:10.815Z	Category	-
Case type	Technical	Language	Portuguese

4. For **Severity**, choose the new severity level from the following options:
 - General guidance
 - System impaired
 - Production system impaired
 - Production system down
 - Business-critical system down
5. For **Reason for case severity change**, choose from the available options for why you're changing the case severity.
6. (Optional) For **Tell us more**, enter additional information about this change.
7. Do one of the following:
 - If you're lowering the support case severity, or if you're raising it from **General guidance** to **System impaired** or **Production system impaired**, choose **Update**.
 - If you're raising the severity to **Production system down** or **Business-critical system down**, use one of the options in the **Contact methods** section to engage with AWS Support, and then choose **Update**. The following example shows the options available in the **Contact methods** section.

Change case severity

Severity

Production system down

Reason for case severity change

-

Tell us more - optional

Tell us more about the reason for the escalation

Contact methods

Info

For high severity issues we recommend using our live contact mechanisms

Web

We'll respond by email and Support Center.

Chat(Recommended)

Chat online with a representative.

Phone(Recommended)

We'll call you back at your number.

Cancel

Update

Note

- If you upgrade your support case severity to **Production system down** or **Business-critical system down**, you must wait 60 minutes before you can change the severity again.
- If your support case is currently set to **Business-critical system down**, you're prompted to initiate live contact with AWS Support instead of assigning a higher severity.
- If you're raising your support case severity level after already raising it at least once, you might encounter a waiting period. For example, if you change the severity from **System impaired** to **Production system impaired** at 6:00 AM, then your support case falls under the 4-hour first-response time for the **Production system impaired** severity level. In this scenario, you can upgrade the severity level again at 10:00 AM, after the 4-hour window. For a list of first-response times for each severity level, see the table in [Understanding AWS Support response times](#).

Example: Create a support case for account and billing

The following example is a support case for a billing and account issue.

Example: Create a support case for account and billing

API Version 2026-09-18 46



Hello!

We're here to help.

Account: 123456789012 • Support plan: Basic • [Change](#) 

How can we help?

Choose the related issue for your case.

1

☒ Account and billing

[Looking for Service limit increase?](#)

☐ Technical

2

Service

Billing ▼

3

Category

Other Billing Questions ▼

4

Severity [Info](#)

General question ▼

1. **Create case** – Choose the type of case to create. In this example, the case type is **Account and billing**.

 Note

If you have the Basic Support plan, you can't create a technical support case.

2. **Service** – If your question affects multiple services, choose the service that's most applicable.
3. **Category** – Choose the category that best fits your use case. When you choose a category, links to information that might resolve your problem appear below.
4. **Severity** – Customers with a paid support plan can choose the **General guidance** (1-day response time) or **System impaired** (12-hour response time) severity level. Customers with a Business Support plan can also choose **Production system impaired** (4-hour response) or **Production system down** (1-hour response). Customers with an Enterprise On-Ramp or Enterprise Support plan can choose **Business-critical system down** (15-minute response for Enterprise Support and 30-minute response for Enterprise On-Ramp).

Response times are for first response from AWS Support. These response times don't apply to subsequent responses. For third-party issues, response times can be longer, depending on the availability of skilled personnel. For more information, see [Choosing an initial support case severity level](#).

 Note

Based on your category choice, you might be prompted for more information.

After you specify the case type and classification, you can specify the description and how you want to be contacted.

Additional information

Describe your issue

✔ Case draft saved

1 Subject

I have an issue with my bill

Maximum 250 characters (222 remaining)

Description

Don't share any sensitive information in case correspondences, such as credentials, credit cards, signed URLs, or personally identifiable information.

[Learn more](#) 

2

I found a charge on my bill for unused resources.

Maximum 5000 characters (4951 remaining)

3

 **Attach files**

Up to 3 attachments, each less than 5MB



Description Guidance

Provide a detailed description of your issue. If you have a question about a charge, provide the date, amount, or any other details about the charge.

Cancel

Previous

Next step: Solve now or contact us

1. **Subject** – Enter a title that briefly describes your issue.

2. **Description** – Describe your support case. This is the most important information that you provide to Support. For some service and category combinations, a prompt appears with related information. Use these links to help resolve your issue. For more information, see [Describing your problem](#).
3. **Attachments** – Attach screenshots and other files that can help support agents resolve your case faster. You can attach up to 10 files. Each file can be up to 150 MB.

After you add your case details, you can choose how you want to be contacted.

How can we help?
[Account and billing](#), [Billing](#),
[Dispute a Charge](#), [General ...](#)

Additional information
[I have an issue in my account](#)

Solve now or contact us

Account: 123456789012 • Support plan: Basic • [Change](#)

Hello! We're here to help.

✔ Case draft saved

Solve now or contact us

Solve now | **Contact us**

Preferred contact language

English ▲

🔍 |

English ✓

中文

한국어

日本語

☐ Phone
We'll call you back at your number.

☐ Chat
Chat online with a representative.

Cancel Previous **Submit**

1. **Preferred contact language** – Choose your preferred language. Currently you can choose Chinese, English, Japanese, or Korean. The customized contact options in your preferred language will be shown by your support plan.
2. Choose a contact method. The contact options that appear depend on the type of case and your support plan.
 - If you choose **Web**, you can read and respond to the case progress in Support Center.
 - Choose **Chat** or **Phone**. If you choose **Phone**, you're prompted for a callback number.
3. Choose **Submit** when your information is complete and you're ready to create the case.

Note

If you choose Japanese as your preferred contact language for support cases, support in Japanese may be available as follows:

- If you need customer service for non-technical support cases, or if you have a Developer Support plan and need technical support, support in Japanese is available during business hours in Japan defined as 09:00 AM to 06:00 PM Japan Standard Time (GMT+9), excluding holidays and weekends.
- If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, technical support is available 24/7 in Japanese.

If you choose Chinese as your preferred contact language for support cases, support in Chinese may be available as follows:

- If you need customer service for non-technical support cases, support in Chinese is available 09:00 AM to 06:00 PM (GMT+8), excluding holidays and weekends.
- If you have a Developer Support plan, technical support in Chinese is available during business hours generally defined as 8:00 AM to 6:00 PM in your country as set in [My Account](#), excluding holidays and weekends. These times may vary in countries with multiple time zones.
- If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, technical support is available 24/7 in Chinese.

If you choose Korean as your preferred contact language for support cases, support in Korean may be available as follows:

- If you need customer service for non-technical support cases, support in Korean is available during business hours in Korea defined as 09:00 AM to 06:00 PM Korean Standard Time (GMT+9), excluding holidays and weekends.
- If you have a Developer Support plan, technical support in Korean is available during business hours generally defined as 8:00 AM to 6:00 PM in your country as set in [My Account](#), excluding holidays and weekends. These times may vary in countries with multiple time zones.

- If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, technical support is available 24/7 in Korean.

Legacy experience: Updating, resolving, and reopening your case

After you create your support case, you can monitor the status of your case in Support Center. A new case begins in the **Unassigned** state. When a support agent begins work on a case, the status changes to **Work in Progress**. The support agent might respond to your case to ask for more information (**Pending Customer Action**) or to let you know that the case is being investigated (**Pending Amazon Action**).

When your case is updated, you receive an email with the correspondence and a link to the case in Support Center. Use the link in the email message to navigate to the support case. You can't respond to case correspondences by email.

Notes

- You must sign in to the AWS account that submitted the support case. If you sign in as an AWS Identity and Access Management (IAM) user, you must have the required permissions to view support cases. For more information, see [Manage access to AWS Support Center](#).
- If you don't respond to the case within a few days, AWS Support resolves the case automatically.
- Support cases that have been in the resolved state for more than 14 days can't be reopened. If you have a similar issue that is related to the resolved case, you can create a related case. For more information, see [Creating a related case](#).

Topics

- [Updating an existing support case](#)
- [Resolving a support case](#)
- [Reopening a resolved case](#)
- [Creating a related case](#)
- [Case history](#)

Updating an existing support case

You can update your case to provide more information for the support agent. For example, you can reply to correspondences, start another live chat, add additional email recipients, and so on.

Note

If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, you can reassign your support case severity level to reflect changes to urgency and business impact. If you don't have one of these support plans, then you can't update the severity of your case. For more information, see [Choosing an initial support case severity level](#) and [Changing the severity level of your support case](#).

To update an existing support case

1. Sign in to the [AWS Support Center Console](#).

Tip

In the AWS Management Console, you can also choose the question mark icon



and then choose **Support Center**.

2. Under **Open support cases**, choose the **Subject** of the support case.
3. Choose **Reply**. In the **Correspondence** section, you can also make any of the following changes:
 - Provide information that the support agent requested
 - Upload file attachments
 - Change your preferred contact method
 - Add email addresses to receive case updates
4. Choose **Submit**.

Tip

If you closed a chat window and you want to start another live chat, add a **Reply** to your support case, choose **Chat**, and then choose **Submit**. A new pop-up chat window opens.

Resolving a support case

When you're satisfied with the response or your problem is solved, you can resolve the case in Support Center.

To resolve a support case

1. Sign in to the [AWS Support Center Console](#).

Tip

In the AWS Management Console, you can also choose the question mark icon



and then choose **Support Center**.

2. Under **Open support cases**, choose the **Subject** of the support case that you want to resolve.
3. (Optional) Choose **Reply** and in the **Correspondence** section, enter why you're resolving the case, and then choose **Submit**. For example, you can enter information about how you fixed the issue yourself in case you need this information for future reference.
4. Choose **Resolve case**.
5. In the dialog box, choose **Ok** to resolve the case.

Note

If AWS Support resolved your case for you, you can use the feedback link to provide more information about your experience with AWS Support.

Example: Feedback links

The following screenshot shows the feedback links in the correspondence of a case in Support Center.

Please let us know if we helped resolve your issue:

If YES, click here:

<https://console.aws.amazon.com/support/feedback?eventId=1234567890&language=en&questionnaireId=Support-HMD-Yes> 

If NO, click here:

<https://console.aws.amazon.com/support/feedback?eventId=1234567890&language=en&questionnaireId=Support-HMD-No> 

Reopening a resolved case

If you're experiencing the same issue again, you can reopen the original case. Provide details about when the issue occurred again and what troubleshooting steps that you tried. Include any related case numbers so that the support agent can refer to previous correspondences.

Notes

- You can reopen your support case up to 14 days from when your issue was resolved. However, you can't reopen a case that has been inactive for more than 14 days. You can create a new case or a related case. For more information, see [Creating a related case](#).
- If you reopen an existing case that has different information than your current issue, the support agent might ask you to create a new case.

To reopen a resolved case

1. Sign in to the [AWS Support Center Console](#).

Tip

In the AWS Management Console, you can also choose the question mark icon



and then choose **Support Center**.

2. Choose **View all cases** and then choose the **Subject** or the **Case ID** of the support case that you want to reopen.
3. Choose **Reopen case**.
4. Under **Correspondence**, for **Reply**, enter the case details.
5. (Optional) Choose **Choose files** to attach files to your case. You can attach up to 3 files.
6. For **Contact methods**, choose one of the following options:
 - **Web** – Get notified by email and the Support Center.
 - **Chat** – Chat online with a support agent.
 - **Phone** – Receive a phone call from a support agent.
7. (Optional) For **Additional contacts**, enter email addresses for other people that you want to receive case correspondences.
8. Review your case details and choose **Submit**.

Creating a related case

After 14 days of inactivity, you can't reopen a resolved case. If you have a similar issue that is related to the resolved case, you can create a related case. This related case will include a link to the previously resolved case, so that the support agent can review the previous case details and correspondences. If you're experiencing a different issue, we recommend that you create a new case.

To create a related case

1. Sign in to the [AWS Support Center Console](#).

Tip


In the AWS Management Console, you can also choose the question mark icon



and then choose **Support Center**.

2. Choose **View all cases** and then choose the **Subject** or the **Case ID** of the support case that you want to reopen.
3. Choose **Reopen case**.

4. In the dialog box, choose **Create related case**. The previous case information will be automatically added to your related case. If you have a different issue, choose **Create new case**.



5. Follow the same steps to create your case. See [Creating a support case](#).

Note

By default, your related case has the same **Type**, **Category**, and **Severity** of the previous case. You can update the case details as needed.

6. Review your case details and choose **Submit**.

After you create your case, the previous case appears in the **Related cases** section, such as in the following example.

Case ID 234567891 [Info](#) Resolve case

Case details

Subject Same issue is happening for my Amazon EC2 instances	Status Unassigned
Case ID 234567891	Severity General question
Created 2021-04-21T20:30:23.945Z	Category General Info and Getting Started
Case type Account	Additional contacts johndoe@example.com
Opened by janedoe@example.com	

Related cases

Subject	Case ID
Problem with EC2 instances	1234567890

Correspondence Reply

Jane Doe Wed Apr 21 2021 13:30:23 GMT-0700 (Pacific Daylight Time)	I keep getting an error for my EC2 instances. What do you recommend that I do to fix it?
--	--

Case history

You can view case history information up to 24 months after you create a case.

Using AWS Support with an AWS SDK

AWS software development kits (SDKs) are available for many popular programming languages. Each SDK provides an API, code examples, and documentation that make it easier for developers to build applications in their preferred language.

SDK documentation	Code examples
AWS SDK for C++	AWS SDK for C++ code examples
AWS CLI	AWS CLI code examples
AWS SDK for Go	AWS SDK for Go code examples
AWS SDK for Java	AWS SDK for Java code examples
AWS SDK for JavaScript	AWS SDK for JavaScript code examples
AWS SDK for Kotlin	AWS SDK for Kotlin code examples
AWS SDK for .NET	AWS SDK for .NET code examples
AWS SDK for PHP	AWS SDK for PHP code examples
AWS Tools for PowerShell	AWS Tools for PowerShell code examples
AWS SDK for Python v4	AWS SDK for Python v4 code examples
AWS SDK for Python (Boto3)	AWS SDK for Python (Boto3) code examples
AWS SDK for Ruby	AWS SDK for Ruby code examples
AWS SDK for Rust	AWS SDK for Rust code examples
AWS SDK for SAP ABAP	AWS SDK for SAP ABAP code examples
AWS SDK for Swift	AWS SDK for Swift code examples

Example availability

Can't find what you need? Request a code example by using the **Provide feedback** link at the bottom of this page.

Using AWS DevOps Agent with your AWS Support plan

AWS DevOps Agent is an AI-powered operations teammate that investigates incidents, recommends preventive improvements, and handles on-demand SRE tasks across your AWS account. If your account is on an eligible Support plan, you can enable AWS DevOps Agent directly from the Support Center Console. For an overview of what the agent does, see [About AWS DevOps Agent](#) in the *AWS DevOps Agent User Guide*.

Topics

- [Plan eligibility and credits](#)
- [Enabling AWS DevOps Agent from your AWS Support Center Console](#)
- [Resources created for AWS DevOps Agent activated from AWS Support](#)
- [Removing AWS DevOps Agent](#)
- [Security for AWS DevOps Agent activated from AWS Support](#)
- [Getting help with AWS DevOps Agent activated from AWS Support](#)

Plan eligibility and credits

You can enable AWS DevOps Agent from the Support Center Console if your account is on one of the following plans:

- AWS Business Support+
- AWS Enterprise Support
- AWS Unified Operations

Each plan includes a monthly allotment of AWS DevOps Agent usage credits. New activations include a 2-month free credit period. After that, your plan provides ongoing monthly credits toward AWS DevOps Agent usage. To view your current allotment and usage, open the Billing and Cost Management console at <https://console.aws.amazon.com/costmanagement/home#/credits>.

Enabling AWS DevOps Agent from your AWS Support Center Console

Enabling AWS DevOps Agent from the Support Center Console creates an agent space and the AWS Identity and Access Management (IAM) resources required by the agent and its operator web app. Activation runs in the US East (N. Virginia) (us-east-1) Region.

To enable AWS DevOps Agent

1. Sign in to the [AWS Support Center Console](#).
2. Navigate to your Support plan management page.
3. On the **AWS DevOps Agent** card, choose **Enable**.
4. In the activation dialog, review the description, choose **Enable DevOps Agent**, and then choose **Save**.

Activation includes the following four phases:

- Setting up monitoring permissions
- Creating the agent space
- Setting up operator access
- Connecting the agent to your AWS account

When activation completes, you can open the AWS DevOps Agent web app from the success message link, or anytime from the AWS Management Console.

For a detailed explanation of what the agent does in your account, what data it stores, and the IAM access it has, see [About AWS DevOps Agent](#) in the *AWS DevOps Agent User Guide*. For an overview of the agent's security posture, see [AWS DevOps Agent Security](#) in the *AWS DevOps Agent User Guide*.

Resources created for AWS DevOps Agent activated from AWS Support

Activation from the Support Center Console creates the following resources in us-east-1. Replace **ACCOUNT_ID** with your 12-digit AWS account ID. The role suffix is a 12-character identifier derived from the agent space.

Resources created when you enable AWS DevOps Agent from the Support Center Console

AWS service	Resource type	Resource name	Trust scope	Permissions granted
AWS DevOps Agent	Agent space	DevOpsAgentSpace	Not applicable	Container for the account association, operator web app configuration, and data the agent generates while it operates.
AWS Identity and Access Management (IAM)	Role	DevOpsAgentRole-AgentSpace- suffix	Trusted by aidevops.amazonaws.com with aws:SourceAccount and aws:SourceArn conditions that scope the role to agent spaces in your own account (confused-deputy protection).	Grants the agent the read-only investigation permissions across AWS services that it needs to investigate resources in your account. Permissions come from the AWS-managed AIDevOpsAgentAccessPolicy attached at activation time. For the full list, see AIDevOpsAgentAccessPolicy in

AWS service	Resource type	Resource name	Trust scope	Permissions granted
				the <i>AWS DevOps Agent User Guide</i> . The customer-managed AIDevOpsAllowAwsSupportActionsPolicy - <i>suffix</i> policy is also attached.
AWS Identity and Access Management (IAM)	Role	DevOpsAgentRole-WebappAdmin - <i>suffix</i>	Trust policy scoped to a specific agent space, so only that agent space's operator web app can assume it.	Grants the operator web app the permissions it needs for chat, journal, recommendations, and Support integration. Permissions come from the AWS-managed AIDevOpsOperatorAppAccessPolicy . For the full list, see AIDevOpsOperatorAppAccessPolicy in the <i>AWS DevOps Agent User Guide</i> .

AWS service	Resource type	Resource name	Trust scope	Permissions granted
AWS Identity and Access Management (IAM)	Customer-managed policy	AIDevOpsAllowAwsSupportActionsPolicy - <i>suffix</i>	Attached to the DevOpsAgentRole-AgentSpace- <i>suffix</i> role.	Grants <code>iam:CreateServiceLinkedRole</code> , scoped to the AWS Resource Explorer service-linked role ARN (<code>arn:aws:iam:: <i>ACCOUNT_ID</i> :role/aws-service-role/resource-explorer-2.amazonaws.com/AWSServiceRoleForResourceExplorer</code>). This permission allows the agent to create the AWS Resource Explorer service-linked role on your behalf if it doesn't already exist, so the agent can use AWS Resource Explorer for topology discovery.

The Support Center Console activation doesn't create resources in any other AWS Region.

Removing AWS DevOps Agent

To remove the resources that were created when you enabled AWS DevOps Agent from the Support Center Console, see [Deleting an agent space](#) in the *AWS DevOps Agent User Guide*. The

procedures cover deleting the agent space, the two IAM roles, and any customer-managed policies that start with AIDevOps. The procedure also documents how to re-enable AWS DevOps Agent later and what to watch out for when you do.

 Note

The AWS-managed policies attached to the IAM roles (AIDevOpsAgentAccessPolicy and AIDevOpsOperatorAppAccessPolicy) are detached but not deleted, because these policies are owned by AWS.

When you follow that procedure, note that your resources have the names listed in [Resources created for AWS DevOps Agent activated from AWS Support](#), and they are all in us-east-1.

Security for AWS DevOps Agent activated from AWS Support

AWS DevOps Agent provides the following security controls:

- Agent spaces are the primary security boundary. Each agent space is isolated to a single AWS account.
- Data is encrypted at rest with AWS-managed keys and encrypted in transit.
- Agent activity is captured in an immutable agent journal and in AWS CloudTrail (CloudTrail).
- AWS DevOps Agent enforces account-boundary, limited-write, and prompt-injection protections.

For the full security posture, including regional processing, integration security, network connectivity, and the shared responsibility model, see [AWS DevOps Agent Security](#) in the *AWS DevOps Agent User Guide*.

Getting help with AWS DevOps Agent activated from AWS Support

For questions about your Support plan, your AWS DevOps Agent usage credits, or to report issues with activation or removal that you cannot resolve through the AWS DevOps Agent User Guide, [contact AWS Support](#).

About the Support Center Console API

The Support Center Console API enhances your experience with the Support Center Console. Examples of the functionality provided by the Support Center Console API include the following:

- The ability for you to create and update a draft of your support case
- The ability for the Support Center Console to display the current status of your account
- The ability for the Support Center Console to display dynamic help for the selected service and category

For a complete list of actions provided by the Support Center Console API, see the table in [Adding IAM policies for the Support Center Console API operations](#).

Important

To continue to use the functions included in the Support Center Console API, you must add the Support Center Console operations to your AWS Identity and Access Management policies before November 16, 2026. After you create the IAM policies, update the AWS Support managed policy to include the `support-console:*` actions. For more information, see [Adding IAM policies for the Support Center Console API operations](#).

Topics

- [Adding IAM policies for the Support Center Console API operations](#)
- [Testing Support Center Console API calls](#)

Adding IAM policies for the Support Center Console API operations

Before November 16, 2026, you must create AWS Identity and Access Management policies for the Support Center Console API operations. If you don't create these policies by November 16, 2026, you will receive `AccessDenied` errors.

To add these operations to your IAM policies, see [Create IAM policies \(console\)](#) in the *AWS Identity and Access Management User Guide*.

The following table summarizes the console operations.

 Note

These operations are for the console only. They're not available for use in the AWS SDK or the AWS CLI.

Operation	Access level	Description
GetAccountState	READ	Grants permission for the console to show the current account state.
GetAccountGovCloudEnabled	READ	Grants permission to determine if your account is GovCloud enabled.
GetCaseDraft	READ	Grants permission for the console to show the case draft that you previously created.
CreateCaseDraft	WRITE	Grants permission to create or update a case draft for the given case type.
DeleteCaseDraft	WRITE	Grants permission to delete a case draft for the given case type.
GetBanner	READ	Grants permission for the console to show the Support banner displayed during customer impacting events.
DescribeDynamicHelp	READ	Grants permission for the console to show dynamic help

Operation	Access level	Description
		resources for the selected service and category.

Operation	Access level	Description
CreateContact	WRITE	Grants permission for the console to create an authenticated contact for the selected contact type, such as the Request to remove email sending limitations, Report Abusive Activity from AWS Resources, and Simulated Event Submissions forms. The following are examples of some of the forms that use this permission. Request to remove email sending limitations In order to maintain the quality of Amazon EC2 addresses for sending email, we enforce default limits on the amount of email that can be sent from EC2 accounts. If you wish to send larger amounts of email from EC2, you can apply to have these limits removed from your account by filling out this form. You can now set and remove your own DNS for EC2 resources. Documentation can be found here. Request information Email address Use case description A valid use case description should consist of the following: * A description of how you plan to send email from EC2 * A statement indicating how you intend to ensure this account is not implicated in sending of unwanted mail. Please be advised that the sending of unwanted mail (spam) from AWS resources is a violation of the AWS Acceptable Use Policy and may result in escalated enforcement by AWS. Elastic IPs Information info IPs Information Elastic IP address - optional Reverse DNS record - optional Enter your Elastic IP address Enter your reverse DNS record Remove item Add item You can add 50 rows (max). By submitting this form, you agree to the Terms and Conditions Submit Simulated Event Submissions Form Complete the form below for Simulated Event type testing of Amazon AWS Services. Customer information Email Address Provide us with your email address for correspondence Details of the source(s) simulating the event type Account ID of Source Add an account ID or NA, if not applicable Non-AWS Source IP addresses Add one or more IP addresses or NA, if not applicable Note: Use commas to separate multiple IP addresses AWS Regions List all the AWS regions in the scope of testing Note: Use commas to separate multiple regions (i.e., us-east-1, us-west-2) Details of the destination(s) to which the event is being simulated Account ID of Destination Add an account ID AWS Resources List all the AWS resources in the scope of testing Note: Use commas to separate multiple AWS resources Non-AWS Destination IP addresses Add one or more IP addresses or NA, if not applicable Note: Use commas to separate multiple IP addresses AWS Regions List all the AWS regions in the scope of testing

Operation	Access level	Description
CheckSubscription	READ	Grants permission for the console to verify if your account has access to the selected product.
GetQuestionnaire	READ	Grants permission for the console to show the customer feedback questionnaire.
SaveFeedback	WRITE	Grants permission to save questionnaire feedback.

 Note

If you have a custom VPN configuration, make sure that you configure your VPN to correctly forward your client IP address to the Support Center Console API endpoint. When using a VPN with AWS Identity and Access Management policies that include [aws:SourceIp](#) conditions, the client IP address specified in your IAM policy must be forwarded to the API endpoint, not the VPN's IP address. If the VPN forwards its own IP address instead of the client IP address, authorization might fail because the IP address doesn't match the `aws:SourceIp` condition in your IAM policy. The following table provides the Support Center Console API endpoints by AWS Region.

AWS Region	Support Center Console API endpoint
<code>https://api.us-east-1.prod.support-console.support.aws.dev</code>	US East (N. Virginia)
<code>https://api.us-west-2.prod.support-console.support.aws.dev</code>	US West (Oregon)

AWS Region	Support Center Console API endpoint
<code>https://api.eu-west-1.prod. support-console.support.aws .dev</code>	Europe (Ireland)

Testing Support Center Console API calls

To validate that your IAM policies are correctly configured for Support Center Console API operations, open the [AWS Support Center Console](#) to generate recent API calls.

To check for missing IAM permissions, complete the following steps:

1. Sign in to the AWS Management Console and open the CloudTrail console at <https://console.aws.amazon.com/cloudtrail>.
2. Check the AWS Region dropdown to make sure that you're in the **US East (N. Virginia)** Region.
3. In the navigation pane, choose **Event history**.
4. Filter by event source **support-console.amazonaws.com**.
5. Match the event names to the list of `support-console:*` operations in [Adding IAM policies for the Support Center Console API operations](#) (for example, `GetAccountState`).
6. Open the matching events and check for an `additionalEventData` field containing an `authZHeader` entry. If present, your IAM policy is missing the permission listed in that entry.
7. Add the specific `support-console` permission to your IAM policy. You can grant access to all operations using `support-console:*`, or select individual operations for fine-grained control. For the full list of operations, see [Adding IAM policies for the Support Center Console API operations](#).
8. To verify the fix, revisit the AWS Support Center Console to generate new API calls, then repeat steps 1–6. Make sure that the new events no longer contain an `additionalEventData` field.

About the AWS Support API

Important

End of Support Notice: Developer Support will be discontinued January 1, 2027. Customers with Developer Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information, see [Business Support+ plan details](#)

End of Support Notice: Business Support will be discontinued January 1, 2027. Customers with Business Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information see, [Business Support+ plan details](#)

End of Support Notice: On January 1, 2027, AWS will discontinue Enterprise On-Ramp. Throughout 2026, Enterprise On-Ramp customers will be automatically upgraded to AWS Enterprise Support during contract renewal or in periodic batches. Customers will receive an email notification a month before their upgrade. No further action is required. Enterprise Support provides designated TAM assignment, 15-minute response times, and AWS Security Incident Response available at no additional cost, all at a lower \$5,000 minimum (reduced from \$15,000). For more information, see [AWS Enterprise Support plan details](#).

For more information, see [Developer, Business, and Enterprise On-Ramp end of support](#). Developer Support, Business Support, and Enterprise On-Ramp will remain available in the AWS GovCloud (US) Region.

The AWS Support API provides access to some of the features in the [AWS Support Center](#).

The API provides two different groups of operations:

- [Support case management](#) operations to manage the entire life cycle of your AWS support cases, from creating a case to resolving it
- [AWS Trusted Advisor](#) operations to access [AWS Trusted Advisor](#) checks

 Note

You must have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan to use the AWS Support API. For more information, see [Support](#).

 Email authentication

When corresponding with AWS Support through email, make sure that your email domain is configured with the following authentication standards:

- **SPF** (Sender Policy Framework) – Authorizes which mail servers can send email on behalf of your domain.
- **DKIM** (DomainKeys Identified Mail) – Adds a cryptographic signature to verify that your emails haven't been altered in transit.
- **DMARC** (Domain-based Message Authentication, Reporting & Conformance) – Tells receiving servers how to handle emails that fail SPF or DKIM checks.

Emails that can't be verified might have reduced functionality, such as the inability to add new recipients to your case through CC.

For guidance on configuring these controls, see [Email authentication methods](#) and [Complying with DMARC](#) in the *Amazon Simple Email Service Developer Guide*. If you are unsure whether your domain is configured correctly, contact your email administrator.

For more information about the operations and data types provided by Support, see the [AWS Support API Reference](#).

Topics

- [Support case management](#)
- [AWS Trusted Advisor](#)
- [Endpoints](#)
- [Support in AWS SDKs](#)

Support case management

You can use the API to perform the following tasks:

- Open a support case
- Get a list and detailed information about recent support cases
- Filter your search for support cases by dates and case identifiers, including resolved cases
- Add communications and file attachments to your cases, and add the email recipients for case correspondences. You can attach up to three files. Each file can be up to 5 MB
- Resolve your cases

The AWS Support API supports CloudTrail logging for support case management operations. For more information, see [Logging AWS Support API calls with AWS CloudTrail](#).

For code examples that demonstrate how to manage the entire life cycle of a support case, see [Code examples for Support using AWS SDKs..](#)

AWS Trusted Advisor

You can use the Trusted Advisor operations to perform the following tasks:

- Get the names and identifiers for the Trusted Advisor checks
- Request that a Trusted Advisor check be run against your AWS account and resources
- Get summaries and detailed information for your Trusted Advisor check results
- Refresh your Trusted Advisor checks
- Get the status of each Trusted Advisor check

The AWS Support API supports CloudTrail logging for Trusted Advisor operations. For more information, see [AWS Trusted Advisor information in CloudTrail logging](#).

You can use Amazon CloudWatch Events to monitor for changes to your check results for Trusted Advisor. For more information, see [Monitoring AWS Trusted Advisor check results with Amazon EventBridge](#).

For example Java code that demonstrates how to use the Trusted Advisor operations, see [Using Trusted Advisor as a web service](#).

Endpoints

Support is a global service. This means that any endpoint that you use will update your support cases in the Support Center Console.

For example, if you use the US East (N. Virginia) endpoint to create a case, you can use the US West (Oregon) or Europe (Ireland) endpoint to add a correspondence to the same case.

You can use the following endpoints for the Support API:

- US East (N. Virginia) – <https://support.us-east-1.amazonaws.com>
- US West (Oregon) – <https://support.us-west-2.amazonaws.com>
- Europe (Ireland) – <https://support.eu-west-1.amazonaws.com>

Important

- If you call the [CreateCase](#) operation to create test support cases, then we recommend that you include a subject line, such as **TEST CASE-Please ignore**. After you're done with your test support case, call the [ResolveCase](#) operation to resolve it.
- To call the AWS Trusted Advisor operations in the AWS Support API, you must use the US East (N. Virginia) endpoint. Currently, the US West (Oregon) and Europe (Ireland) endpoints don't support the Trusted Advisor operations.

For more information about AWS endpoints, see [AWS Support endpoints and quotas](#) in the *Amazon Web Services General Reference*.

Support in AWS SDKs

The AWS Command Line Interface (AWS CLI), and the AWS Software Development Kits (SDKs) include support for the Support API.

For a list of languages that support the AWS Support API, choose an operation name, such as [CreateCase](#), and in the [See Also](#) section, choose your preferred language.

AWS Support Plans

AWS offers the following AWS Support Plans for you to choose from based on your business needs.

- Basic
- AWS Business Support+
- AWS Enterprise Support
- AWS Unified Operations

Important

End of Support Notice: Developer Support will be discontinued January 1, 2027. Customers with Developer Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information, see [Business Support+ plan details](#)

End of Support Notice: Business Support will be discontinued January 1, 2027. Customers with Business Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information see, [Business Support+ plan details](#)

End of Support Notice: On January 1, 2027, AWS will discontinue Enterprise On-Ramp. Throughout 2026, Enterprise On-Ramp customers will be automatically upgraded to AWS Enterprise Support during contract renewal or in periodic batches. Customers will receive an email notification a month before their upgrade. No further action is required. Enterprise Support provides designated TAM assignment, 15-minute response times, and AWS Security Incident Response available at no additional cost, all at a lower \$5,000 minimum (reduced from \$15,000). For more information, see [AWS Enterprise Support plan details](#).

For more information, see [Developer, Business, and Enterprise On-Ramp end of support](#). Developer Support, Business Support, and Enterprise On-Ramp will remain available in the AWS GovCloud (US) Region.

Topics

- [Features of AWS Support Plans](#)
- [What is AWS Unified Operations](#)
- [Change AWS Support Plans](#)
- [Sign up for AWS Enterprise Support using self-service subscription](#)
- [Configure promotional plan expiration notifications](#)
- [Developer, Business, and Enterprise On-Ramp end of support](#)

Features of AWS Support Plans

Basic Support offers assistance for account and billing questions and service quota increases. The other plans offer a number of technical support cases with pay-by-the-month pricing and no long-term contracts.

All AWS customers automatically have 24x7 access to these features of Basic Support:

- One-on-one responses to account and billing questions
- Support forums
- Service health checks
- Documentation, technical papers, and best practice guides

Customers with a Business Support+ plan have access to these additional features:

- Real time and contextual responses through generative AI.
- 24x7 phone, web, and chat access to Cloud Support Engineers.
- Less than 30 minutes human engagement for business critical system down cases.
- More than 500 checks with [AWS Trusted Advisor](#).
- The AWS Support API to interact with Support Center and Trusted Advisor. You can use the AWS Support API to automate support case management and Trusted Advisor operations.
- Third-party software support – Help with Amazon Elastic Compute Cloud (Amazon EC2) instance operating systems and configuration. Also, help with the performance of the most popular third-party software components on AWS.
- Supports an unlimited number of support cases that can be opened by any user with [permissions](#).

Customers with an Enterprise Support plan have all of the benefits available with the AWS Business Support+, plan plus the following additional benefits:

- A designated Technical Account Manager (TAM).
- Up to 15 minute production-critical case response.
- Access to [AWS Security Incident Response](#), a service that helps you quickly prepare for, respond to, and receive guidance to help recover from security incidents. This includes incidents like account takeovers, data breaches, and ransomware attacks.
- Strategic business reviews with AWS experts.
- Access to AWS Countdown event management. AWS Countdown provides specialized TAM-led support to help you succeed during critical business events.

Customers with an AWS Unified Operations plan have access to the following benefits. For details about the AWS Unified Operations plan, including onboarding details and getting started information, see [What is AWS Unified Operations](#).

- Application architecture guidance – Consultative guidance on how services fit together to meet your specific use case, workload, or application.
- Infrastructure event management – Short-term engagement with AWS Support to get a deep understanding of your use case. After analysis, provide architectural and scaling guidance for an event.
- Designated AWS Technical Account Manager (TAM) and Domain Specialist Engineer (DSE) for your specific use cases and applications.
- One recurring subscription to [AWS Countdown Premium](#).
- Critical workload reviews and personalized runbooks and operational procedures.

For more information about features and pricing for each support plan, see [AWS Support](#) and [Compare AWS Support plans](#). Some features, such as 24x7 phone and chat support, aren't available in all languages.

 Note

If you work with an AWS partner and want to learn more about Partner-led Support, see [AWS Partner-Led Support](#)

What is AWS Unified Operations

AWS Unified Operations combines proven expertise with AI-powered insights to help reduce operational and security risks, resolve issues faster, and help you architect more resilient cloud solutions from the start—accelerating your most critical cloud initiatives. Designated AWS specialists act as an extension of your team through your preferred collaboration channels—conducting workload reviews, providing strategic guidance, and optimizing performance through deep contextual knowledge. With 24/7 security and performance monitoring, we detect and mitigate incidents early while reducing alert volume. A five-minute, context-aware response for critical incidents further shortens resolution time and helps maintain peak operational performance.

AWS Unified Operations pricing

AWS Unified Operations pricing is based on your specific requirements and workload complexity. For detailed pricing information, see [AWS Support Plan Pricing](#).

Contents

- [Benefits of Unified Operations](#)
- [Unified Operations Team](#)
 - [Technical Account Manager](#)
 - [Domain Specialist Engineer](#)
 - [Senior Billing and Account Specialist](#)
 - [Incident Management Engineer](#)
 - [Migrations and Events Engineer](#)
 - [AWS Customer Incident Response Team](#)
 - [Specialist Support Engineer](#)
- [Unified Operations life cycle](#)
 - [Unified Operations pre-onboarding](#)
 - [Unified Operations onboarding](#)
 - [Unified Operations Pre-event or migration planning](#)
 - [Unified Operations Event migration or cutover](#)
 - [Unified Operations Post Go-live or event](#)
 - [Unified Operations Workload incident management](#)

- [Unified Operations Post incident](#)
- [Unified Operations Continuous improvement](#)
- [Getting started with Unified Operations](#)
 - [Unified Operations Getting started: Prerequisites](#)
 - [Unified Operations Getting started: Onboard critical alarms to rapid incident management](#)
 - [Unified Operations Getting started: How to request 5-minute incident response](#)
 - [Unified Operations Getting started: Plan for domain coverage](#)
 - [Unified Operations Getting started: Onboard your account to proactive security incident management](#)
 - [Unified Operations Getting started: AWS expectations from you](#)
 - [Unified Operations Getting started: What you can expect from AWS](#)

Benefits of Unified Operations

Unified Operations offers several key benefits.

- **Designated AWS experts:** Your AWS team includes designated Technical Account Managers (TAMs), AWS Domain Specialist Engineer, and Senior Billing and Account Specialist. Your AWS team integrates with your collaboration tools like Slack or Microsoft Teams.
- **Deep technical guidance:** Your AWS team helps you strengthen your resiliency through application-specific deep dives, readiness assessments, guided testing support, and custom runbooks tailored to your environment. Your AWS team provides workload focused financial management and aligns costs strategically with your business objectives.
- **Migrations, events, and launch support:** Accelerate your critical cloud migrations and business events with designated AWS engineers who help mitigate risks proactively, guiding you from planning through execution. Real-time assistance for scheduled events, increasing migration velocity and successful launches.
- **24x7 Proactive workload monitoring:** Onboard application and infrastructure level alarms from your existing AWS and third-party tools to detect early warning signs and drive proactive mitigation with your team.
- **Engage AWS incident managers within 5 mins for business-critical system down issues:** Receive proactive support from Incident Management Engineers within 5 minutes of alarms, workload alerts, or business-critical system down issues.

- **Context-specific case response:** Access to context-aware support engineers to drive incident resolution. Runbooks and incident response playbooks customized for your workflows, streamlining problem diagnosis and resolution for your specific business needs.
- **Security guidance and support:** Proactive monitoring of security events with automated triage and investigation, along with 24/7 access to the AWS Customer Incident Response Team to prepare for, respond to, and recover from security events in your AWS environment.

Unified Operations Team

AWS Unified Operations brings together a designated team of specialized experts who work together to support your cloud journey. Each role is carefully designed to provide comprehensive coverage across technical, financial, operational, and security dimensions of your cloud environment. From strategic guidance to technical support, from financial optimization to rapid incident response, these experts act as an extension of your team, available 24/7 to ensure your cloud operations run efficiently and securely.

Contents

- [Technical Account Manager](#)
- [Domain Specialist Engineer](#)
- [Senior Billing and Account Specialist](#)
- [Incident Management Engineer](#)
- [Migrations and Events Engineer](#)
- [AWS Customer Incident Response Team](#)
- [Specialist Support Engineer](#)

Technical Account Manager

Your designated cloud strategist, the Technical Account Manager (TAM) orchestrates the overall engagement and drives business outcomes. TAMs lead strategic planning, conduct quarterly business reviews, provide resiliency guidance, and coordinate across specialized teams to ensure your cloud initiatives succeed. They serve as your primary point of contact for escalations and strategic decisions.

Domain Specialist Engineer

A Domain Specialist Engineer (DSE) is a technical expert who deeply understands your specific workload architecture and AWS services. DSEs conduct critical workload reviews, create technical documentation, develop troubleshooting guides, and provide ongoing technical consultation. They analyze incidents to prevent recurrence and maintain global knowledge sharing to make sure that there is consistent support quality.

Senior Billing and Account Specialist

Your Senior Billing and Account Specialist (SBAS) is a designated financial optimization expert who helps balance performance with cost. They manage cost optimization strategies, oversee reserved instance and savings plan portfolios, conduct financial business reviews, and provide detailed spend analytics to maximize your cloud investment efficiency.

Incident Management Engineer

Incident Management Engineers (IME) are rapid response specialists who coordinate critical incident resolution. IMEs provide 5-minute response times, orchestrate technical teams during incidents, manage stakeholder communications. During active incidents, IMEs conduct real-time assessments of incident handling and response effectiveness, upon request they document the sequence of events, decisions made, and immediate outcomes. They observe and evaluate the execution of response protocols, team coordination, and the application of existing playbooks while the incident is still unfolding.

Migrations and Events Engineer

On-demand experts who guide critical transitions like launches and migrations. They validate architectures, create detailed execution plans, provide real-time monitoring during events, and conduct post-event analysis to capture learnings and optimize future operations.

AWS Customer Incident Response Team

The AWS Customer Incident Response Team (CIRT) are security experts providing 24/7 specialized assistance for security events. They monitor security findings, provide guided response within minutes of detection, and enhance your security operations capabilities through expert investigation support and best practices guidance.

Specialist Support Engineer

Specialist Support Engineers (SSE) are highly experienced technical experts using advanced contextual tools to deliver precise support solutions. They leverage AI-powered systems and deep technical knowledge to quickly understand your environment and resolve complex technical challenges.

Unified Operations life cycle

The Unified Operations Support plan contains different phases, from pre-onboarding, through continuous improvement, that help you get the most out of your cloud environment. This topic covers key points of each phase.

Contents

- [Unified Operations pre-onboarding](#)
- [Unified Operations onboarding](#)
- [Unified Operations Pre-event or migration planning](#)
- [Unified Operations Event migration or cutover](#)
- [Unified Operations Post Go-live or event](#)
- [Unified Operations Workload incident management](#)
- [Unified Operations Post incident](#)
- [Unified Operations Continuous improvement](#)

Unified Operations pre-onboarding

During the pre-onboarding phases, AWS gathers information from you needed to start the onboarding process, including the following information:

Environment discovery and technical validation

- Understanding of your workload architecture and your key AWS services.
- Your future planning needs, such as migration or events.
- Pre-requisites, such as a list of your AWS accounts and AWS Regions.
- Your specific business needs.

Unified Operations onboarding

Onboarding kickoff

- Meet the team (customer and AWS allocated resources).

Onboarding workshop

- Identify critical workloads.
- Conduct deep architecture review of the workloads.
- Review roles and responsibilities (RACI review – AWS and your roles and responsibilities).
- Review Incident and Change management workflows (ITSM).
- Communication protocol - tooling and processes, escalation path, on-call schedules
- Identify and define the critical alarms (in Amazon CloudWatch, your 3rd party APM or custom monitoring tool).
- Build runbooks for critical alarms.

Service onboarding

- Your AWS account onboarding to AWS Security Incident Response.
- Critical alarm onboarding to rapid Incident Management.

Unified Operations Pre-event or migration planning

Pre-event or migration planning in AWS Unified Operations includes the following key elements.

- **Context gathering** : Workload discovery and architecture.
- **Operational Readiness Review (ORR)**: Systematic assessment against AWS best practices and runbooks to identify potential issues before they impact the event.
- **Risk assessments**: Identify, list potential risks, and mitigation plans.
- **Review Events**: Secure the cut-over or migration event support.

Unified Operations Event migration or cutover

Event migration or cutover in Unified Operations includes the following key elements.

- **Real-time bridge to resolution:** AWS experts join your communication channels to monitor events and resolve issues during critical business periods.
- **24/7 comprehensive support:** Round-the-clock expert assistance with immediate guidance for resource scaling needs. .
- **AWS engineering:** Work directly with engineers who understand your business, delivering tailored solutions.

Unified Operations Post Go-live or event

The post go-live or post event process in Unified Operations includes the following key elements:

- Spin-down engagement and event-specific resources.
- Conduct event reviews.
- Update runbooks and documentation based on learnings.
- Perform retrospectives to identify areas for improvement.

Unified Operations Workload incident management

Workload incident management includes the following key elements:

- AWS Support case creation.
- Engagement with you and an AWS Incident Management Engineer (IME) for context-aware incident management.
- Joining or creation of incident bridge call.
- Monitoring of AWS service health and large-scale event (LSE).
- Rapid recovery of critical applications or workloads.

Unified Operations Post incident

Post-incident analysis for critical incidents is conducted by the assigned Domain Specialist Engineer (DSE). DSE takes a broader analytical approach after resolution, conducting comprehensive root cause analysis to identify any gaps in processes or tooling. The DSE transforms insights into actionable improvements by updating response playbooks, recommending preventive measures, and suggesting architectural enhancements to help prevent similar incidents in the future.

Unified Operations Continuous improvement

Continuous improvement includes the following key elements:

- Update Critical Workload Reviews, proposing AWS service-specific recommendations and resilience guidance .
- Review old and new cases on a continuous basis, and provide troubleshooting guides for identified technical issues.
- Apply lessons and oversee implementations and testing.
- Discuss technical issues, configurations, past, and upcoming projects and milestones.
- Review new feature implementation plans with risk assessments and performance optimizations.
- Conduct Monthly or Quarterly Business Review (MBR / QBR).

Getting started with Unified Operations

This topic discusses the steps to onboard to AWS Unified Operations.

Contents

- [Unified Operations Getting started: Prerequisites](#)
- [Unified Operations Getting started: Onboard critical alarms to rapid incident management](#)
- [Unified Operations Getting started: How to request 5-minute incident response](#)
- [Unified Operations Getting started: Plan for domain coverage](#)
- [Unified Operations Getting started: Onboard your account to proactive security incident management](#)
- [Unified Operations Getting started: AWS expectations from you](#)
- [Unified Operations Getting started: What you can expect from AWS](#)

Unified Operations Getting started: Prerequisites

The following items are required to onboard to AWS Unified Operations

A signed AWS Unified Operations contract. For more information, contact your AWS sales representative.

- Identified business needs, such as migration, modernization, events, target uptime, and so on.

- A list of your workloads.
- A list of your AWS accounts and associated AWS Regions.
- Identified stakeholders across Application, Architecture, Operations, and Security teams.

Unified Operations Getting started: Onboard critical alarms to rapid incident management

To help quickly notify you of critical incidents, complete the following steps to onboard your alarms to AWS Incident Detection and Response

1. Define and configure your critical alarms for rapid incident management. For detailed information, see [Define and configure alarms in Incident Detection and Response](#) in the *Incident Detection and Response User Guide*.
 - a. For steps to set up alarms using Amazon CloudWatch, see [Define and configure alarms in Incident Detection and Response](#) in the *Incident Detection and Response User Guide*. For AWS recommendations on critical alarm types for various AWS services, see [Incident Detection and Response \(IDR\)](#). Contact your AWS Unified Operations team if you want AWS to automate the creation of critical AWS alarms for your tagged AWS resources.
 - b. To redirect or ingest critical alarms from 3rd party APM tools with [direct Amazon EventBridge integration](#), such as DataDog, NewRelic, and so on, see [Ingest alarms from APMs that have direct integration with Amazon EventBridge](#) in the *AWS Incident Detection and Response User Guide*. You must deploy a set of AWS resources (AWS Lambda and Amazon EventBridge event bus rules) to transform and redirect your alarm (event) to AWS Incident Detection and Response. Your AWS Unified Operations team can help provide the CloudFormation template to install these resources.
 - c. Redirect or ingest critical alarms from your custom monitoring tool through a 3rd party APM tool that doesn't have direct integration with Amazon EventBridge. For more information, see [Use webhooks to ingest alarms from APMs without direct integration with Amazon EventBridge](#) in the *AWS Incident Detection and Response User Guide*. You must deploy a set of AWS resources (API Gateway AWS Lambda functions, and Amazon EventBridge event bus rules) to transform and redirect your alarm (event) to AWS Incident Detection and Response. Your AWS Unified Operations team can help provide the CloudFormation template to install these resources.
2. Provide workload architecture details, point of contact information and runbook information on mitigation actions for critical alarms. To do this, complete the following steps:

- Download and complete the [AWS Incident Detection and Response Workload onboarding questionnaire](#) for each critical workload or application and the [Alarm ingestion questionnaire](#) related to each unique workload.

The information in these questionnaires helps the AWS team develop an incident remediation runbook. This runbook enables appropriate actions to be taken to quickly troubleshoot and remediate critical alarms before they cause business downtime. For examples and sample information, see [Workload onboarding and alarm ingestion questionnaires in AWS Incident Detection and Response](#).

3. Provide access to onboard your critical alarms to AWS Incident Detection and Response
 - a. Deploy the `AWSServiceRoleForHealth_EventProcessor` service-linked role (SLR) in your AWS account running the critical workload to be monitored by the AWS incident management team. For more information, see [Provision access for alert ingestion to AWS Incident Detection and Response](#).

 Note

To assist you with onboarding of large AWS accounts, AWS can provide you with a AWS Command Line Interface script to fast track the provisioning of this SLR.

- b. (Optional) If your alarms are in Amazon CloudWatch, make sure that the AWS Identity and Access Management user or role that's used for alarm testing (before go-live) has the `cloudwatch:SetAlarmState` IAM permission in your AWS account that's running the critical workload. This is needed for alarm testing (gameday) post onboarding. For more information, see [Test onboarded workloads in AWS Incident Detection and Response](#).
4. Create a AWS Support case to subscribe a workload for rapid incident management. Note that your AWS account is automatically enabled for inbound rapid incident management, which means you can raise a case to the Unified Operations Incident Detection and Response queue through the Support Center Console, the AWS Command Line Interface, or the AWS SDK for quick action. For AWS to proactively monitor and create incidents with an outbound AWS Support case, create an AWS Support case for your critical workload. To do this, complete the following steps:
 - a. Sign in to the [AWS Support Center Console](#), select **Create case**, and then select **Technical support**.
 - b. For **Service** select **Incident Detection and Response**.

- c. For **Category** select **Onboard new workload**.
- d. For **Severity** select **General guidance**.
- e. Attached the Workload and Alarm questionnaires that you completed in the previous step.

Unified Operations Getting started: How to request 5-minute incident response

AWS Unified Operations offers 5-minute incident response for your critical incidents. To request a 5-minute inbound response you can [create a support case from a support interaction](#) or use the [legacy support case creation method](#). When you create your case, make sure that you enter the following information to ensure that your case receives a response within 5 minutes:

Note

Use this method regardless of whether you use proactive monitoring or not, or when incidents occur before alarm setup is complete.

1. For **Case type**, choose **Technical**.
2. For **Service**, choose the affected AWS service.
3. For **Category**, choose the option that best matches your issue.
4. For **Severity**, choose **Business-critical system down**.
5. In the **Description**, include the following information
 - a. Technical information
 - Workload name
 - Affected AWS Resource ARN(s)
 - b. Business information
 - Description of impact to the business
 - (Optional) Customer bridge details

 Note

Resolution effectiveness depends on available context. Alarm onboarded workloads benefit from immediate access to comprehensive workload information, while non-onboarded workloads begin with service health data and customer-provided information.

Unified Operations Getting started: Plan for domain coverage

AWS Unified Operations provides specialized expertise through a domain-based coverage approach. Each domain is supported by a team of AWS Domain Specialists who provide the following services:

- **Specialized expertise** aligned to your specific technology areas.
- **Continuous coverage** with availability through your preferred collaboration tools (Slack or Microsoft Teams) during business days.
- **Proactive guidance** on architecture, best practices, and optimization opportunities.
- **Enhanced incident response** through deep domain knowledge and workload familiarity.
- **Consistent experience** maintained by a coordinated team rather than individuals.

This approach to domain coverage enables AWS specialists to maintain deep familiarity with your critical workloads while providing comprehensive support across your technology stack.

To select the domains, organizations maintain decision authority from a choice of 23 AWS Domains and consider the following factors in their decision:

- Primary AWS services running critical workloads
- Critical AWS service dependencies (such as Amazon EC2, Amazon EKS, or Amazon RDS)
- Major upcoming events requiring 24x7 support coverage (migrations, launches) planned within 3-6 months

This information, combined with guidance from your Technical Account Manager, enables precise alignment of domain expertise with your specific organizational needs, helping you maintain optimal support for mission-critical workloads.

Unified Operations Getting started: Onboard your account to proactive security incident management

Unified Operations entitles you to AWS Security Incident Response to help you quickly prepare for, respond to, and recover from security incidents, such as account takeovers, data breaches, and ransomware attacks. AWS Security Incident Response triages findings, escalates events, and manages critical cases, while also providing access to the AWS Customer Incident Response Team (CIRT) to investigate impacted resources. This access helps you to effectively mitigate and resolve security incidents, minimizing the impact on your operations. To onboard to this service feature, complete the following steps:

1. Create a centralized AWS account for AWS Security Incident Response. This AWS account will be used to configure all other AWS accounts that you want monitored, to manage your incident response team, and to create and view security events. We recommend that you align this account with the account that you use for other security services such as Amazon GuardDuty and AWS Security Hub CSPM. You can use an [AWS Organizations](#) management account, or an AWS Organizations delegated administrator account as the Security Incident Response membership account. For more information, see [Select a membership account](#) in the *AWS Security Incident Response User Guide*.
 - a. Choose basic membership details. For more information, see [Setup membership details](#) in the *AWS Security Incident Response User Guide*.
 - b. Choose how you want to associate accounts with AWS Organizations. For more information, see [Associate accounts with AWS Organizations](#) in the *AWS Security Incident Response User Guide*.
 - c. (Optional) You can optionally enable proactive response and alert triaging workflow to enable within your organization to monitor and investigate alerts generated from Amazon GuardDuty and AWS Security Hub CSPM integrations. For more information, see [Setup proactive response and alert triaging workflows](#) in the *AWS Security Incident Response User Guide*.
2. (Optional) Enable the proactive containment of a potential security incident. AWS can perform containment actions to quickly mitigate impact, such as isolating compromised hosts or rotating credentials. To turn on this feature, you must first grant the necessary permissions to the service. To do this, deploy an [Step Functions StackSet](#).

Unified Operations Getting started: AWS expectations from you

For Unified Operations to deliver maximum value, we recommend the following collaborative approach:

Team engagement

- Identify subject matter experts from your team to collaborate with AWS engineers during onboarding and ongoing engagement.
- Participate in initial discovery calls and subsequent meetings to share architecture details and operational requirements.
- Establish regular touchpoints to review architecture updates or workload changes.

Operational integration

- Configure critical alarms in your account to enable effective incident management.
- Implement recommended action items provided by AWS specialists,
- Participate in gameday exercises to validate incident response processes.

This collaborative framework helps you maximize the value of Unified Operations, achieve your uptime goals, mitigate operational risks, and receive comprehensive support for your mission-critical workloads.

Unified Operations Getting started: What you can expect from AWS

When you onboard to Unified Operations, you can expect the following from AWS.

- Provide a team of designated AWS experts with deep technical expertise in the your workload domain and services.
- Offer proactive guidance, ongoing optimization, and continuous improvement recommendations to enhance workload performance and resiliency and accelerate path to migrations and modernization.
- Help provide rapid incident response, with context-aware engineers engaged within 5 minutes after you raise a business-critical support case.

Note

Resolution effectiveness depends on available context. Alarm onboarded workloads benefit from immediate access to comprehensive workload information, while non-onboarded workloads begin with service health data and customer-provided information.

- Offer comprehensive support throughout the application lifecycle, from design and migration to production launch and long-term operations.
- Proactively monitor security threats with auto-triaging, reducing false positives, and raising incidents for potential security incidents.
- Assist in trouble and joint mitigation of AWS or your identified security incident.

Change AWS Support Plans

You can use the AWS Support Plans console to change your support plan for your AWS account. To change your support plan, you must have AWS Identity and Access Management(IAM) permissions. For more information, see [Manage access to AWS Support Plans](#) and [AWS managed policies for AWS Support Plans](#).

To change your support plan

1. Sign in to the AWS Support Plans console at <https://console.aws.amazon.com/support/plans/home>.
2. (Optional) To compare support plans, on the **AWS Support Plans** page, choose **Compare all Support plans and features**.
3. (Optional) To view estimated costs for a support plan, choose **Pricing calculator**. In the **Pricing calculator**, select a support tier, enter an estimate of how much you expect to spend with AWS each month, and then choose **Calculate**.
4. To downgrade your Enterprise Support plan, reach out to your Technical Account Manager (TAM).

To downgrade your Business Support+ plan, on the [Manage Support Plans](#) page, choose **Review downgrade** in the Basic Support plan section.

To upgrade to an AWS Enterprise Support plan, see [Sign up for AWS Enterprise Support using self-service subscription](#). To upgrade to an AWS Unified Operations plan, choose **Contact sales**.

To upgrade to an AWS Business Support+ plan from Basic Support, complete the following steps:

- a. Choose **Get started** in the AWS Business Support+ section.
- b. If you are onboarded to AWS Organizations and have the all-feature mode enabled, you can subscribe your entire organization to Business Support+. For information about all-feature mode, see [Enabling all features for an organization with AWS Organizations](#) in the *AWS Organizations User Guide*. To enroll at the organization level, select the **My organization** radio button. Review the [plan details and pricing](#), and then check the box to agree to the subscription terms. Only your organization's management account can subscribe your entire organization.

Or, to enroll at the account level, select the **My account** radio button, then check the box to agree to the subscription terms.

- c. Choose **Confirm upgrade** to complete your AWS Business Support+ plan subscription.

 Note

After you subscribe your entire organization to AWS Business Support+, any new accounts created within or invited to that organization are automatically subscribed to AWS Business Support+. If an account that was previously subscribed to AWS Business Support+ at the organization-level leaves the organization, that account is downgraded to the Basic support plan.

 Note

If you sign up for a paid support plan, you're responsible for a minimum one month subscription of AWS Support. For more information, see the [AWS Support FAQs](#).

Related information

For more information about AWS Support Plans, see the [AWS Support FAQs](#). You can also choose **Contact us** from the Support Plans console.

To close your account, see [Closing an Account](#) in the *AWS Billing User Guide*.

Sign up for AWS Enterprise Support using self-service subscription

With [AWS Enterprise Support](#), you get a designated Technical Account Manager (TAM), 24x7 access to Cloud Support Engineers by phone, chat, and web, and a 15-minute response time when a business-critical system is down. You also get proactive risk detection, 24x7 security monitoring with Amazon GuardDuty and AWS Security Hub CSPM, architectural reviews, cost optimization workshops, AI/ML guidance, AWS Well-Architected Agent access, monthly AWS DevOps Agent credits, AWS Countdown, and AWS Security Incident Response at no additional fee. For a full comparison, see [Compare AWS Support Plans](#).

Subscribe directly from the AWS Support Plans console without contacting sales. Enrollment covers every member account in your organization, including accounts you add later. It activates immediately for most eligible accounts and has no long-term commitment. If an account that was enrolled at the organization level later leaves the organization, that account is downgraded to the Basic support plan. For current pricing, see [AWS Support pricing](#).

For instructions, see [Subscribe to AWS Enterprise Support](#). Your TAM reaches out within the first few business days after enrollment to schedule onboarding.

Checking eligibility requirements

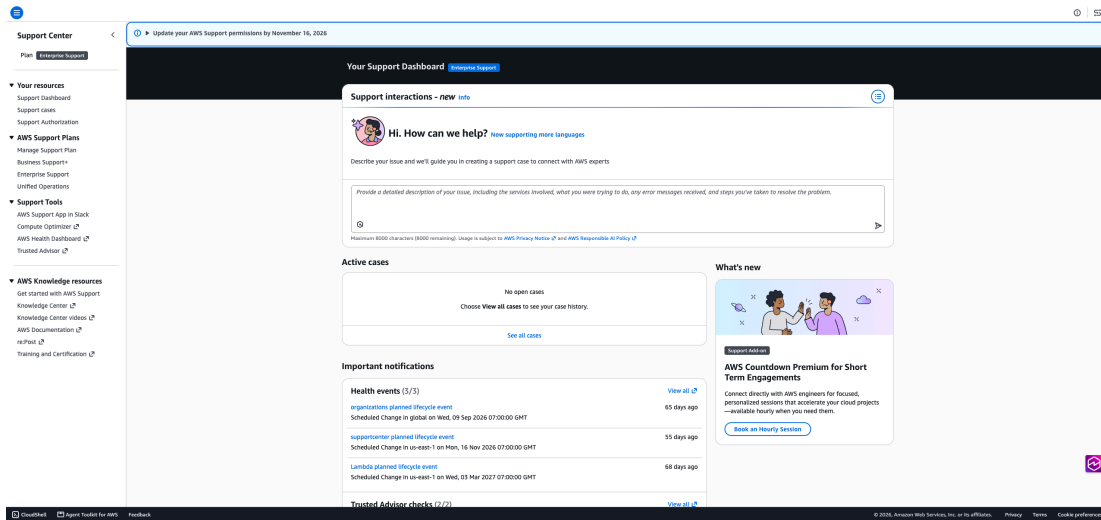
Self-service subscription is available to AWS customers who meet all of the following requirements:

- You have a single-payer AWS Organization.
- You are not currently on an Enterprise Support, Enterprise On-Ramp, or Unified Operations plan.
- Your account is not associated with a reseller or AWS Partner account.
- You are the root user of the management account, or an AWS Identity and Access Management (IAM) user with support purchase permissions.

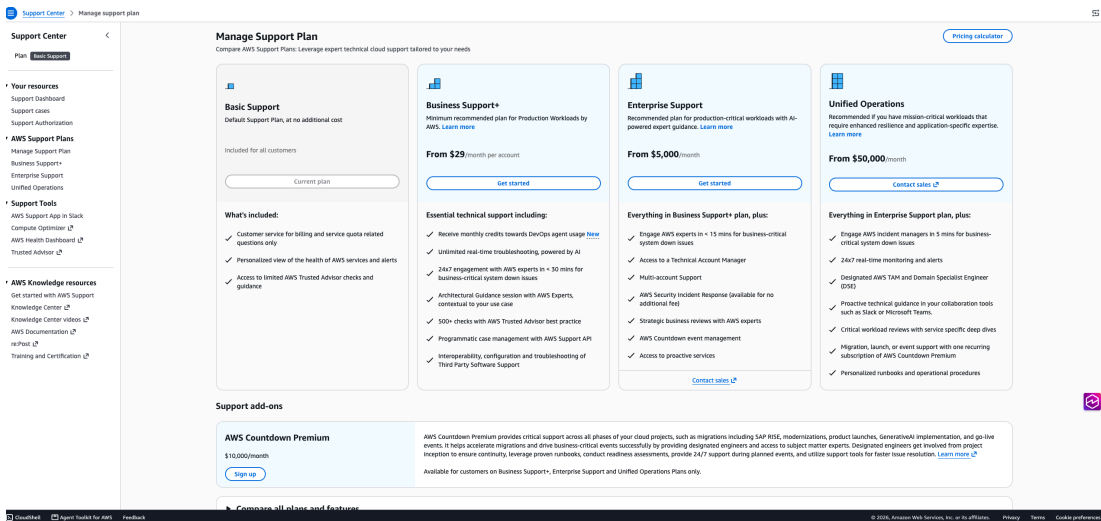
Subscribe to AWS Enterprise Support

To sign up for AWS Enterprise Support

1. Open the [AWS Support console](#) and choose **Manage Support Plan** under **AWS Support Plans** in the left navigation pane.



2. On the **Manage Support Plan** page, find the **Enterprise Support** column and choose **Get started**.



3. Review the Enterprise Support benefits and pricing, then choose **Next** to check your account's eligibility.

Support Center > Manage support plan > Review upgrade

Confirm Enterprise Support plan subscription

No long-term commitment. No risk. Just a 30-day minimum to start seeing results.

You're about to unlock these benefits

- AWS Well-Architected Agent**
Optimize your AWS applications and architecture across cost, security, performance, and reliability, with recommendations ranked against your declared business goals and get automation-ready fixes.
- Resolves and proactively prevents incidents by DevOps agent**
Quickly resolve issues as they arise, proactively prevent future incidents, and unlock untapped insights across your operations and workloads. Each month, you get 75% of your AWS Enterprise Support spend as credits for AWS DevOps Agent.
- Designated TAM**
Your Technical Account Manager (TAM) works with you to define and refine strategic outcomes on AWS across Well-Architected pillars. Your team will participate in planning workshops, define strategic outcomes, and identify the critical proactive AWS engagements to support your cloud journey.
- Advisory Services**
Your Advanced plan includes Well-Architected reviews and comprehensive engagements that evaluate your applications and operations and plan for improvements.
- AWS Security Incident Response**
Get help when it matters most with automated monitoring and investigation, accelerated communication and coordination, and direct 24/7 access to the AWS Customer Incident Response Team to help you to prepare for, respond to, and recover quickly from security events.

► See more benefits

► Calculate pricing

Cancel Next

© 2024, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

- If your account is eligible, you proceed to the confirmation page. If validation fails, the page lists each reason. See [Troubleshooting](#) to resolve eligibility failures.
- On the confirmation page, review what to expect and the subscription terms, select the checkbox to accept the terms, and then choose **Submit request**.

Support Center > Manage support plan > Review upgrade

Confirm Enterprise Support plan subscription

No long-term commitment. No risk. Just a 30-day minimum to start seeing results.

What to expect

Review below to understand what happens after you submit your request.

Email notification - You'll receive an email confirmation once your account is successfully enrolled in AWS Enterprise Support.

When your benefits begin - Once subscribed to Enterprise Support, you'll have immediate access to all features including 24/7 technical support. Your Technical Account Manager (TAM) will reach out to you to begin working with your team.

Billing start date - Billing starts when your account is subscribed to Enterprise Support. Your charges will appear on your monthly AWS bill.

☒ By checking this box, you agree to the terms and conditions as described below:

- Pricing: From \$5,000 per month. For detailed pricing, refer to [pricing details](#).
- Enrollment: Subscribing through this console enrolls all member accounts within your AWS Organization. For selective enrollment, [Contact Sales](#) for enrollment assistance.
- Payment terms: You'll be charged to your saved payment method.
- Commitment: I am committing to a 30-day subscription.
- Cancellation: You may terminate AWS Enterprise Support at any time after 30 days of enrollment by contacting your account manager.
- Auto-renewal: Subscription will be auto-renewed after every 30 days.
- Service Linked Roles: AWS Support Plans will have the necessary permissions to update your support agreement in your management account.

Cancel Previous **Submit request**

© 2024, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

- After the activation request completes, the console displays the Enterprise Support plan page with a banner confirming that your enrollment is complete. You also receive a subscription confirmation email.

Your current plan

AWS Enterprise Support

Help your business operate and grow at scale with expert strategic guidance



Welcome to Enterprise Support plan!

Support Plan for your AWS Organization has been updated to Enterprise Support. You now have access to expert guidance and priority assistance to maximize your experience on AWS.

A designated Technical Account Manager (TAM) will be assigned to your organization within the first few business days. Your TAM will reach out to introduce themselves and schedule an onboarding session.

Help us improve. [Share feedback on your sign up experience.](#)

Here's what you can do next



Create case

Your Enterprise Support plan includes 24x7 access to Cloud Support Engineers via phone, chat, and web. You can have an unlimited number of contacts that can open an unlimited amount of cases.

[Create case](#)



AWS Well-Architected Agent

Optimize your AWS applications and architecture across cost, security, performance, and reliability, with recommendations ranked against your declared business goals and get automation-ready fixes.

[Get started](#)



AWS DevOps Agent

Drive operational excellence with DevOps agent that resolves and proactively prevents incidents. You get first 2 months free, then 75% of Enterprise Support charges as credits towards the agent. [See credits](#).

Agent status

Enabled

[Manage your AWS DevOps Agent](#)

Feature benefits



Designated TAM

Your Enterprise Support plan includes a designated Technical Account Manager (TAM) who brings cloud expertise to accelerate your strategic outcomes. Your TAM provides architectural and operational guidance tailored to your cloud journey, from migrations to generative AI innovations. [Learn about TAM engagement](#)



Advisory Services

Your Enterprise Support plan includes Well-Architected Reviews and comprehensive engagements that evaluate your applications and operations and plan for improvements.



AWS Security Incident Response

Get help when it matters most with automated monitoring and investigation, accelerated communication and coordination, and direct 24/7 access to the AWS Customer Incident Response Team to help you to prepare for, respond to, and recover quickly from security events.



AWS Countdown

Execute your critical launches, migrations, and peak traffic events with confidence through AWS Countdown. Get comprehensive planning and proven playbooks from your TAM and AWS experts, enabling your teams to innovate faster while supporting successful outcomes every step of the way.



Cloud Expertise

Your Enterprise Support plan offers engagements designed to increase your team's expertise in AWS services, including AWS GameDays, AWS Immersion Days, and AWS Workshops and Hackathons.



Continuous Improvement

Get structured reviews, tools, and programs that help you optimize AWS services and improve your usage continuously over time as your environment evolves.



Access to Cloud Experts

Our Cloud Support Engineering, Trust & Safety, and Customer Service teams have experience solving millions of cases across hundreds of thousands of AWS customers. Support Engineers help you address issues affecting your environment, and advise on how to avoid future issues.



Enhanced Technical Support

Get 24x7 access to Cloud Support Engineers via phone, chat, and web. Open unlimited cases with unlimited contacts. Response times: less than 24 hours for general guidance, 12 hours for system impaired, 4 hours for production system impaired, 1 hour for production system down, and 15 mins for production-critical system down.



Billing and Account Management

Our billing and account experts specialize in working with enterprise accounts. They quickly assist with billing and account inquiries and help implement best practices so you can focus on running your business.

[► Compare all plans and features](#)

Understanding post-enrollment steps

After your account is enrolled:

- You receive an email confirmation. Billing starts when your account is subscribed and appears on your monthly AWS bill. For details, see [AWS Support pricing](#).
- AWS assigns a designated Technical Account Manager (TAM) to your team within the first few business days. Your TAM reaches out to introduce themselves and schedule an onboarding session.

Troubleshooting

Eligibility failures

Management account required

You're signed in as a member account of an AWS Organization, or as a standalone account that isn't part of an organization. Only the organization's management account can submit an Enterprise Support subscription request. Sign out, sign in with your management account, and reopen the Support Plans console. If your account isn't part of an organization, create one by following the steps in [Creating an organization](#) in the *AWS Organizations User Guide*.

Payer account required

The management account for your organization isn't the payer account, and only the payer account can submit the request. Sign in with your payer account and try again. If your organization uses a separate payer, [contact sales](#) for assisted enrollment.

All Features mode required

Your organization is in consolidated billing mode, and Enterprise Support requires AWS Organizations all features mode. Enable all features for your organization by following the steps in [Enabling all features](#) in the *AWS Organizations User Guide*, then return to the console and try again.

Email address not supported

The email address on the account isn't on a company domain, and support plan enrollment requires a company domain. Update the account's contact email to a company domain address, or [create a case](#) in Support Center for help.

Partner or reseller account not supported

Your account is associated with an AWS Partner or reseller, and self-service enrollment is available to direct AWS customers only. Contact your AWS Partner or reseller to add Enterprise Support to your account.

Support plan request already exists

An Enterprise Support activation is already in progress for your organization. Wait for the in-flight request to complete. To check its status or cancel it, contact your account manager or [create a case](#) in Support Center.

You don't have the necessary IAM permissions

Your IAM principal is missing the permissions needed to read or change the support plan. Ask your AWS administrator to grant the required permissions, or sign in as the management account root user. For details, see [Manage access to AWS Support Plans](#).

Unable to verify your account eligibility

The console couldn't complete the eligibility check, usually because of a transient error. Reload the page and try again. If the problem persists, [create a case](#) in Support Center.

Enrollment failures after you submit your request

Taking longer than expected

Activation didn't complete within the expected window, and your request might still be processing. Check your support plan in the Support Plans console before resubmitting.

Unable to complete enrollment

Your enrollment request was rejected and Enterprise Support was not activated. You were not charged.

If enrollment fails after you submit your request, you have the following options:

- **Contact sales** – In the [Support Plans console](#), choose **Contact sales**. Complete the form with your full name, email address, and phone number. A representative responds within 24 hours.
- **Create a support case** – In [Support Center](#), create a case referencing your Enterprise Support enrollment attempt.
- **Share feedback** – If you're already subscribed to Enterprise Support, go to the [Support Plans console](#) and choose **Share feedback on your sign up experience** in the banner to tell us about your enrollment.

Changing or canceling your plan

To change or cancel your AWS Enterprise Support plan, contact your designated TAM or [your AWS sales representative](#).

Configure promotional plan expiration notifications

You can use [AWS User Notifications](#) to configure notifications to inform you when your support plan's promotional period is ending. You can subscribe to receive notifications by email, in the AWS Console Mobile Application, or in other chat channels of your choice.

Configure promotional support plan expiration notifications

1. Open User Notifications in the [AWS Management Console](#):
 - a. Choose the bell icon in the top navigation bar.
 - b. Choose **Notification center**.
 - c. In the navigation pane, choose **Notification configuration**.
 - d. Choose **Create notification configuration**.
 - e. Select at least one **Configuration hub**. For more information, see [Storing, processing, and replicating notifications using notification hubs in AWS User Notifications](#).
2. For **Event Rule**, enter the following information:
 - For **AWS service name**, enter **Support Plans**.
 - For **Event type**, enter **Support Plan Promotion Expiration**.
 - For **Regions**, select the source AWS Regions where you want to receive notifications. For this option, choose US East (N. Virginia), US East (Ohio), US West (N. California), and US West (Oregon).
3. Configure aggregation settings to reduce notification frequency. We recommend that you set aggregation to **Receive within 5 minutes**.
4. Configure the delivery channels where you want to receive notifications. If you don't select a delivery channel, you can view notifications by selecting the bell icon in the AWS Management Console navigation bar.

For detailed instructions on creating user-configured notifications, see [Step 1: creating a notification configuration in the AWS User Notifications User Guide](#).

View promotional plan notifications

Your notifications are delivered to the delivery channel that you chose during configuration. You can also view notifications by choosing the bell icon in the console navigation bar. The bell icon shows a red badge when new notifications are available.

For more information on viewing notifications, see [Step 2: Viewing notifications](#) in the *AWS User Notifications User Guide*.

Developer, Business, and Enterprise On-Ramp end of support

For information on transitioning to a new plan, see the following information specific to your plan and the [AWS Support Frequently Asked Questions](#).

Note

Developer Support, Business Support, and Enterprise On-Ramp will remain available in the AWS GovCloud (US) Region.

Developer Support plan end of support

End of Support Notice: Developer Support will be discontinued January 1, 2027. Customers with Developer Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information see, [Business Support+ plan details](#).

Business Support plan end of support

End of Support Notice: Business Support will be discontinued January 1, 2027. Customers with Business Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information see, [Business Support+ plan details](#)

Enterprise On-Ramp end of support

End of Support Notice: On January 1, 2027, AWS will discontinue Enterprise On-Ramp. Throughout 2026, Enterprise On-Ramp customers will be automatically upgraded to AWS Enterprise Support during contract renewal or in periodic batches. Customers will receive an email notification a month before their upgrade. No further action is required. Enterprise Support provides designated TAM assignment, 15-minute response times, and AWS Security Incident Response available at no additional cost, all at a lower \$5,000 minimum (reduced from \$15,000). For more information, see [AWS Enterprise Support plan details](#).

AWS Trusted Advisor

Important

End of Support Notice: Developer Support will be discontinued January 1, 2027. Customers with Developer Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information, see [Business Support+ plan details](#)

End of Support Notice: Business Support will be discontinued January 1, 2027. Customers with Business Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information see, [Business Support+ plan details](#)

End of Support Notice: On January 1, 2027, AWS will discontinue Enterprise On-Ramp. Throughout 2026, Enterprise On-Ramp customers will be automatically upgraded to AWS Enterprise Support during contract renewal or in periodic batches. Customers will receive an email notification a month before their upgrade. No further action is required. Enterprise Support provides designated TAM assignment, 15-minute response times, and AWS Security Incident Response available at no additional cost, all at a lower \$5,000 minimum (reduced from \$15,000). For more information, see [AWS Enterprise Support plan details](#).

For more information, see [Developer, Business, and Enterprise On-Ramp end of support](#). Developer Support, Business Support, and Enterprise On-Ramp will remain available in the AWS GovCloud (US) Region.

Trusted Advisor draws upon best practices learned from serving hundreds of thousands of AWS customers. Trusted Advisor inspects your AWS environment, and then makes recommendations when opportunities exist to save money, improve system availability and performance, or help close security gaps.

AWS Trusted Advisor checks are available to customers with an AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan.

If you have a Basic Support plan, you can use the Trusted Advisor console to access all checks in the Service Limits category and [selected checks](#) in the Security and Fault tolerance categories. Automatic check updates aren't available in the Basic Support plan. You must manually refresh Trusted Advisor checks in the Security category. To manually refresh a check, do the following:

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. Select the **Refresh** button on the check that you want to refresh.

If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, you can use the Trusted Advisor console and the [AWS Trusted Advisor API](#) to access all Trusted Advisor checks. You also can use Amazon CloudWatch Events to monitor the status of Trusted Advisor checks. For more information, see [Monitoring AWS Trusted Advisor check results with Amazon EventBridge](#).

You can access Trusted Advisor in the AWS Management Console. For more information about controlling access to the Trusted Advisor console, see [Manage access to AWS Trusted Advisor](#).

For more information, see [Trusted Advisor](#).

Topics

- [Get started with Trusted Advisor Recommendations](#)
- [Get started with the Trusted Advisor API](#)
- [Using Trusted Advisor as a web service](#)
- [Organizational view for AWS Trusted Advisor](#)
- [View AWS Trusted Advisor checks powered by AWS Config](#)
- [Viewing AWS Security Hub CSPM controls in AWS Trusted Advisor](#)
- [Opt in AWS Compute Optimizer for Trusted Advisor checks](#)
- [Get started with AWS Trusted Advisor Priority](#)
- [AWS Trusted Advisor check reference](#)
- [Change log for AWS Trusted Advisor](#)

Get started with Trusted Advisor Recommendations

Important

End of Support Notice: Developer Support will be discontinued January 1, 2027. Customers with Developer Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information, see [Business Support+ plan details](#)

End of Support Notice: Business Support will be discontinued January 1, 2027. Customers with Business Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information see, [Business Support+ plan details](#)

End of Support Notice: On January 1, 2027, AWS will discontinue Enterprise On-Ramp. Throughout 2026, Enterprise On-Ramp customers will be automatically upgraded to AWS Enterprise Support during contract renewal or in periodic batches. Customers will receive an email notification a month before their upgrade. No further action is required. Enterprise Support provides designated TAM assignment, 15-minute response times, and AWS Security Incident Response available at no additional cost, all at a lower \$5,000 minimum (reduced from \$15,000). For more information, see [AWS Enterprise Support plan details](#).

For more information, see [Developer, Business, and Enterprise On-Ramp end of support](#). Developer Support, Business Support, and Enterprise On-Ramp will remain available in the AWS GovCloud (US) Region.

You can use the Trusted Advisor Recommendations page of the Trusted Advisor console to review check results for your AWS account and then follow the recommended steps to fix any issues. For example, Trusted Advisor might recommend that you delete unused resources to reduce your monthly bill, such as an Amazon Elastic Compute Cloud (Amazon EC2) instance.

You can also use the AWS Trusted Advisor API to perform operations on your Trusted Advisor checks. For more information, see the [AWS Trusted Advisor API Reference](#)

Topics

- [Sign in to the Trusted Advisor console](#)
- [View check categories](#)
- [View specific checks](#)
- [Filter your checks](#)
- [Refresh check results](#)
- [Download check results](#)
- [Organizational view](#)
- [Preferences](#)

Sign in to the Trusted Advisor console

You can view the checks and the status of each check in the Trusted Advisor console.

Note

You must have AWS Identity and Access Management (IAM) permissions to access the Trusted Advisor console. For more information, see [Manage access to AWS Trusted Advisor](#).

To sign in to the Trusted Advisor console

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. On the **Trusted Advisor Recommendations** page, view the summary for each check category:
 - **Action recommended (red)** – Trusted Advisor recommends an action for the check. For example, a check that detects a security issue for your IAM resources might recommend urgent steps.
 - **Investigation recommended (yellow)** – Trusted Advisor detects a possible issue for the check. For example, a check that reaches a quota for a resource might recommend ways to delete unused resources.
 - **Checks with excluded items (gray)** – The number of checks that have excluded items, such as resources that you want a check to ignore. For example, this might be Amazon EC2 instances that you don't want the check to evaluate.

3. You can do the following on the **Trusted Advisor Recommendations** page:

- To refresh all checks in your account, choose **Refresh all checks**.
- To create an .xls file that includes all check results, choose **Download all checks**.
- Under **Checks summary**, choose a check category, such as **Security**, to view the results.
- Under **Potential monthly savings**, you can view how much you can save for your account and the cost optimization checks for recommendations.
- Under **Recent changes**, you can view changes to check statuses within the last 30 days. Choose a check name to view the latest results for that check or choose the arrow icon to view the next page.

Example: Trusted Advisor Recommendations

The following example shows a summary of the check results for an AWS account.

The screenshot displays the 'Trusted Advisor Recommendations' page. At the top, there are buttons for 'Refresh all checks' and 'Download all checks'. Below this, a brief description states: 'Use this page to get an overview of the check results in your AWS account. Choose a check name or category to view the recommended actions or potential issues that Trusted Advisor has identified. Each check provides more information about how to address any issues. You can also download a summary of all check results. [Learn more](#)'.

The main content is divided into two sections:

Checks summary

Action recommended		Investigation recommended		Checks with excluded items	
Info		Info		Info	
Security	30	Fault tolerance	29	Security	11
Performance	1	Performance	9	Cost optimization	11
Fault tolerance	9	Operational Excellence	12	Service limits	1
Cost optimization	1	Cost optimization	14	Performance	2
Service limits	1	Security	63	Fault tolerance	3

Potential monthly savings

\$7,082.26

Trusted Advisor has identified 18 cost optimization checks that can save you money. For example, you might have unused resources in your AWS account that can be deleted. Choose a cost optimization check to view the recommendations.

[View all cost optimization checks](#)

View check categories

You can view the check descriptions and results for the following check categories:

- **Cost optimization** – Recommendations that can potentially save you money. These checks highlight unused resources and opportunities to reduce your bill.
- **Performance** – Recommendations that can improve the speed and responsiveness of your applications.
- **Security** – Recommendations for security settings that can make your AWS solution more secure.
- **Fault tolerance** – Recommendations that help increase the resiliency of your AWS solution. These checks highlight redundancy shortfalls and overused resources.

- **Service limits** – Checks the usage for your account and whether your account approaches or exceeds the limit (also known as quotas) for AWS services and resources.
- **Operational Excellence** – Recommendations to help you operate your AWS environment effectively, and at scale.

To view check categories

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. In the navigation pane, choose the check category.
3. On the category page, view the summary for each check category:
 - **Action recommended (red)** – Trusted Advisor recommends an action for the check.
 - **Investigation recommended (yellow)** – Trusted Advisor detects a possible issue for the check.
 - **No problems detected (green)** – Trusted Advisor doesn't detect an issue for the check.
 - **Excluded items (gray)** – The number of checks that have excluded items, such as resources that you want a check to ignore.
4. For each check, choose the refresh icon
()
to refresh this check.
5. Choose the download icon
()
to create an .xls file that includes the results for this check.

Example: Cost optimization category

The following example shows 10 (green) checks that don't have any issues.

Cost optimization

Refresh all checks

Download all checks

Choose a check name to see recommendations for ways to help save money for your AWS account. Trusted Advisor might recommend that you delete unused and idle resources, or use reserved capacity.

Overview

Potential monthly savings

\$7,082.26

1

Action recommended

Info

14

Investigation recommended

Info

10

No problems detected

Info

11

Checks with excluded items

Info

Cost optimization checks

Filter by tag key

Learn more about using tags

Tag Key

Tag Value

Reset

Apply filter

Search by keyword

Info

Source

View

Filter checks

All sources

All checks

< 1 2 >

Amazon Comprehend Underutilized Endpoints

Checks the throughput configuration of your endpoints.

Last updated: 2 hours ago

View specific checks

Expand a check to view the full check description, your affected resources, any recommended steps, and links to more information.

To view a specific check

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. In the navigation pane, choose a check category.
3. Choose the check name to view the description and the following details:
 - **Alert Criteria** – Describes the threshold when a check will change status.
 - **Recommended Action** – Describes the recommended actions for this check.
 - **Additional Resources** – Lists related AWS documentation.
 - A table that lists the affected items in your account. You can include or exclude these items from check results.
4. (Optional) To exclude items so that they don't appear in check results:
 - a. Select an item and choose **Exclude & Refresh**.
 - b. To view all excluded items, choose **Excluded items**.
5. (Optional) To include items so that the check evaluates them again:

- a.

Choose **Excluded items**, select an item, and then choose **Include & Refresh**.

b.

To view all included items, choose **Included items**.
6.

Choose the settings icon
()
In the **Preferences** dialog box, you can specify the number of items or the properties to display, and then choose **Confirm**.

Example: Cost optimization check

The following **Low Utilization Amazon EC2 Instances** check lists the affected instances in the account. This check identifies 38 Amazon EC2 instances that have low usage and recommends that you stop or terminate the resources.

▼  Low Utilization Amazon EC2 Instances

Last updated: 14 hours ago  

Checks the Amazon Elastic Compute Cloud (Amazon EC2) instances that were running at any time during the last 14 days and alerts you if the daily CPU utilization was 10% or less and network I/O was 5 MB or less on 4 or more days. Running instances generate hourly usage charges. Although some scenarios can result in low utilization by design, you can often lower your costs by managing the number and size of your instances.

Estimated monthly savings are calculated by using the current usage rate for On-Demand Instances and the estimated number of days the instance might be underutilized. Actual savings will vary if you are using Reserved Instances or Spot Instances, or if the instance is not running for a full day. To get daily utilization data, download the report for this check.

Alert Criteria

Yellow: An instance had 10% or less daily average CPU utilization and 5 MB or less network I/O on at least 4 of the previous 14 days.

Recommended Action

Consider stopping or terminating instances that have low utilization, or scale the number of instances by using Auto Scaling. For more information, see [Stop and Start Your Instance](#), [Terminate Your Instance](#), and [What is Auto Scaling?](#)

Additional Resources

[Monitoring Amazon EC2](#)
[Instance Metadata and User Data](#)
[Amazon CloudWatch Developer Guide](#)
[Auto Scaling Developer Guide](#)

Low Utilization Amazon EC2 Instances (38)

Exclude & Refresh

Included items ▼

38 of 39 Amazon EC2 instances have low average daily utilization. Monthly savings of up to \$713.23 might be available by minimizing underutilized instances. 1 items have been excluded.

< 1 2 > 

Region/AZ ▼	Instance ID ▼	Instance Name	Instance Type ▼	Estimated Monthly Savings ▼	CPU Utilization 14-Day Average ▼
ca-central-1b	i-0f818268643c7ae32		t2.micro	\$9.22	0.1%
ca-central-1a	i-06c233a11aa626588		t2.micro	\$9.22	0.1%

Filter your checks

On the check category pages, you can specify which check results that you want to view. For example, you might filter by checks that have detected errors in your account so that you can investigate urgent issues first.

Filter your checks

API Version 2026-09-18 111

If you have checks that evaluate items in your account, such as AWS resources, you can use tag filters to only show items that have the specified tag.

To filter your checks

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. In the navigation pane or the **Trusted Advisor Recommendations** page, choose the check category.
3. For **Search by keyword**, enter a keyword from the check name or description to filter your results.
4. For the **View** list, specify which checks to view:
 - **All checks** – List all checks for this category.
 - **Action recommended** – List checks that recommend that you take action. These checks are highlighted in red.
 - **Investigation recommended** – List checks that recommend that you take possible action. These checks are highlighted in yellow.
 - **No problems detected** – List checks that don't have any issues. These checks are highlighted in green.
 - **Checks with excluded items** – List checks that you specified to exclude items from the check results.
5. If you added tags to your AWS resources, such as Amazon EC2 instances or AWS CloudTrail trails, you can filter your results so that the checks only show items that have the specified tag.

For **Filter by tag**, enter a tag key and value, and then choose **Apply filter**.

6. In the table for the check, the check results only show items that have the specified key and value.
7. To clear the filter by tags, choose **Reset**.

Related information

For more information about tagging for Trusted Advisor, see the following topics:

- [AWS Support enables tagging capabilities for Trusted Advisor](#)
- [Tagging AWS resources](#) in the *AWS General Reference*

Refresh check results

You can refresh checks to get the latest results for your account. If you have a Developer or Basic Support plan, you can sign in to the Trusted Advisor console to refresh the checks. If you have an AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, Trusted Advisor automatically refreshes the checks in your account on a weekly basis.

To refresh Trusted Advisor checks

1. Navigate to the AWS Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor>.
2. On the **Trusted Advisor Recommendations** or a check category page, choose **Refresh all checks**.

You can also refresh specific checks in the following ways:

- Choose the refresh icon



for an individual check.

- Use the [RefreshTrustedAdvisorCheck](#) API operation.

Notes

- Trusted Advisor automatically refreshes some checks several times a day, such as the **AWS Well-Architected high risk issues for reliability** check. It might take a few hours for changes to appear in your account. For these automatically refreshed checks, you can't choose the refresh icon  to manually refresh your results.
- If you enabled AWS Security Hub CSPM for your account, you can't use the Trusted Advisor console to refresh Security Hub CSPM controls. For more information, see [Refresh your Security Hub CSPM findings](#).

Download check results

You can download check results to get an overview of Trusted Advisor in your account. You can download results for all checks or a specific check.

To download check results from Trusted Advisor Recommendations

1. Navigate to the AWS Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor>.
 - To download all check results, in the **Trusted Advisor Recommendations** or a check category page, choose **Download all checks**.
 - To download a check result for a specific check, choose the check name, and then choose the download icon
().
2. Save or open the .xls file. The file contains the same summary information from the Trusted Advisor console, such as the check name, description, status, affected resources, and so on.

Organizational view

You can set up the organizational view feature to create a report for all member accounts in your AWS organization. For more information, see [Organizational view for AWS Trusted Advisor](#).

Preferences

On the **Manage Trusted Advisor** page, you can [disable Trusted Advisor](#).

On the **Notifications** page, you can configure your weekly email messages for the check summary. See [Set up notification preferences](#).

On the **Your organization** page, you can enable or disable trusted access with AWS Organizations. This is required for the [Organizational view for AWS Trusted Advisor](#) feature and [Trusted Advisor Priority](#).

Set up notification preferences

Specify who can receive the weekly Trusted Advisor email messages for check results and the language. You receive an email notification about your check summary for Trusted Advisor Recommendations once a week.

The email notifications for Trusted Advisor Recommendations don't include results for Trusted Advisor Priority. For more information, see [Manage Trusted Advisor Priority notifications](#).

To set up notification preferences

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. In the navigation pane, under **Preferences**, choose **Notifications**.
3. For **Recommendations**, select whom to notify for your check results. You can add and remove contacts from the [Account Settings](#) page in the AWS Billing and Cost Management console.
4. For **Language**, choose the language for the email message.
5. Choose **Save your preferences**.

Set up organizational view

If you set up your account with AWS Organizations, you can create reports for all member accounts in your organization. For more information, see [Organizational view for AWS Trusted Advisor](#).

Disable Trusted Advisor

When you disable this service, Trusted Advisor won't perform any checks on your account. Anyone who tries to access the Trusted Advisor console or use the API operations will receive an access denied error message.

To disable Trusted Advisor

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. In the navigation pane, under **Preferences**, choose **Manage Trusted Advisor**.
3. Under **Trusted Advisor**, turn off **Enabled**. This action disables Trusted Advisor for all checks in your account.
4. You can then manually delete the [AWSServiceRoleForTrustedAdvisor](#) from your account. For more information, see [Deleting a service-linked role for Trusted Advisor](#).

Related information

For more information about Trusted Advisor, see the following topics:

- [How do I start using Trusted Advisor?](#)
- [AWS Trusted Advisor check reference](#)

Get started with the Trusted Advisor API

Important

End of Support Notice: Developer Support will be discontinued January 1, 2027. Customers with Developer Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information, see [Business Support+ plan details](#)

End of Support Notice: Business Support will be discontinued January 1, 2027. Customers with Business Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information see, [Business Support+ plan details](#)

End of Support Notice: On January 1, 2027, AWS will discontinue Enterprise On-Ramp. Throughout 2026, Enterprise On-Ramp customers will be automatically upgraded to AWS Enterprise Support during contract renewal or in periodic batches. Customers will receive an email notification a month before their upgrade. No further action is required. Enterprise Support provides designated TAM assignment, 15-minute response times, and AWS Security Incident Response available at no additional cost, all at a lower \$5,000 minimum (reduced from \$15,000). For more information, see [AWS Enterprise Support plan details](#).

For more information, see [Developer, Business, and Enterprise On-Ramp end of support](#). Developer Support, Business Support, and Enterprise On-Ramp will remain available in the AWS GovCloud (US) Region.

The AWS Trusted Advisor API Reference is intended for programmers that need detailed information about the Trusted Advisor API operations and data types. This API provides access to Trusted Advisor recommendations for your account or all the accounts within your AWS Organization. The Trusted Advisor API uses HTTP methods that returns results in JSON format.

Note

- You must have an AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan to use the Trusted Advisor API.
- If you call the AWS Trusted Advisor API from an account that doesn't have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, then you receive an Access Denied exception. For more information about changing your support plan, [see AWS Support](#).

You can use the AWS Trusted Advisor API to get a list of checks and their descriptions, recommendations, and resources for recommendations. You can also update the lifecycle of recommendations. To manage recommendations, use the following API operations:

- Use the [ListChecks](#), [ListRecommendations](#), [GetRecommendation](#), and [ListRecommendationResources](#) API operations to view recommendations and corresponding accounts and resources.
- Use The [UpdateRecommendationLifecycle](#) API operation to update the lifecycle of a recommendation that's managed by Trusted Advisor Priority.
- Use The [BatchUpdateRecommendationResourceExclusion](#) API operation to include or exclude one or more resources from your Trusted Advisor results.
- The [ListOrganizationRecommendations](#), [GetOrganizationRecommendation](#), [ListOrganizationRecommendationResources](#), [ListOrganizationRecommendationAccounts](#), and [UpdateOrganizationRecommendationLifecycle](#) API calls support only recommendations that are managed by Trusted Advisor Priority. These recommendations are also referred to as prioritized recommendations. You can view and manage your prioritized recommendations from a management or delegated admin account if you have activated Trusted Advisor Priority. If Priority isn't activated, then you receive an Access Denied exception when you make requests.

For more information, [see AWS Trusted Advisor in the AWS Support User Guide](#).

For authentication of requests, [see the Signature Version 4 Signing Process](#).

Using Trusted Advisor as a web service

The AWS Support service enables you to write applications that interact with [AWS Trusted Advisor](#). This topic shows you how to get a list of Trusted Advisor checks, refresh one of them, and then get the detailed results from the check. These tasks are demonstrated in Java. For information about support for other languages, see [Tools for Amazon Web Services](#).

Topics

- [Get the list of available Trusted Advisor checks](#)
- [Refresh the list of available Trusted Advisor checks](#)
- [Poll a Trusted Advisor check for status changes](#)
- [Request a Trusted Advisor check result](#)
- [Show details of a Trusted Advisor check](#)

Get the list of available Trusted Advisor checks

The following Java code snippet creates an instance of an Support client that you can use to call all Trusted Advisor API operations. Next, the code gets the list of Trusted Advisor checks and their corresponding CheckId values by calling the [DescribeTrustedAdvisorChecks](#) API operation. You can use this information to build user interfaces that enable users to select the check they want to run or refresh.

```
private static AWSSupport createClient()
{
    return AWSSupportClientBuilder.defaultClient();
}
// Get the List of Available Trusted Advisor Checks
public static void getTAChecks() {
    // Possible language parameters: "en" (English), "ja" (Japanese), "fr" (French),
    "zh" (Chinese)
    DescribeTrustedAdvisorChecksRequest request = new
    DescribeTrustedAdvisorChecksRequest().withLanguage("en");
    DescribeTrustedAdvisorChecksResult result =
    createClient().describeTrustedAdvisorChecks(request);
    for (TrustedAdvisorCheckDescription description : result.getChecks()) {
        // Do something with check description.
        System.out.println(description.getId());
        System.out.println(description.getName());
    }
}
```

```
}  
}
```

Refresh the list of available Trusted Advisor checks

The following Java code snippet creates an instance of an Support client that you can use to refresh Trusted Advisor data.

```
// Refresh a Trusted Advisor Check  
// Note: Some checks are refreshed automatically, and they cannot be refreshed by using  
// this operation.  
// Specifying the check ID of a check that is automatically refreshed causes an  
// InvalidParameterValue error.  
public static void refreshTACheck(final String checkId) {  
    RefreshTrustedAdvisorCheckRequest request = new  
    RefreshTrustedAdvisorCheckRequest().withCheckId(checkId);  
    RefreshTrustedAdvisorCheckResult result =  
    createClient().refreshTrustedAdvisorCheck(request);  
    System.out.println("CheckId: " + result.getStatus().getCheckId());  
    System.out.println("Milliseconds until refreshable: " +  
    result.getStatus().getMillisUntilNextRefreshable());  
    System.out.println("Refresh Status: " + result.getStatus().getStatus());  
}
```

Poll a Trusted Advisor check for status changes

After you submit the request to run a Trusted Advisor check to generate the latest status data, you use the [DescribeTrustedAdvisorCheckRefreshStatuses](#) API operation to request the progress of the check's run, and when new data is ready for the check.

The following Java code snippet gets the status of the check requested in the following section, using the value corresponding in the CheckId variable. In addition, the code demonstrates several other uses of the Trusted Advisor service:

1. You can call `getMillisUntilNextRefreshable` by traversing the objects contained in the `DescribeTrustedAdvisorCheckRefreshStatusesResult` instance. You can use the value returned to test whether you want your code to proceed with refreshing the check.
2. If `timeUntilRefreshable` equals zero, you can request a refresh of the check.
3. Using the status returned, you can continue to poll for status changes; the code snippet sets the polling interval to a recommended ten seconds. If the status is either `enqueued` or

`in_progress`, the loop returns and requests another status. If the call returns successful, the loop terminates.

4. Finally, the code returns an instance of a `DescribeTrustedAdvisorCheckResultResult` data type that you can use to traverse the information produced by the check.

Note: Use a single refresh request before polling for the status of the request.

```
// Retrieves TA refresh statuses. Multiple checkId's can be submitted.
public static List<TrustedAdvisorCheckRefreshStatus> getTARefreshStatus(final String...
    checkIds) {
    DescribeTrustedAdvisorCheckRefreshStatusesRequest request =
        new
DescribeTrustedAdvisorCheckRefreshStatusesRequest().withCheckIds(checkIds);
    DescribeTrustedAdvisorCheckRefreshStatusesResult result =
        createClient().describeTrustedAdvisorCheckRefreshStatuses(request);
    return result.getStatuses();
}
// Retrieves a TA check status, and checks to see if it has finished processing.
public static boolean isTACheckStatusInTerminalState(final String checkId) {
    // Since we only submitted one checkId to getTARefreshStatus, just retrieve the
    // only element in the list.
    TrustedAdvisorCheckRefreshStatus status = getTARefreshStatus(checkId).get(0);
    // Valid statuses are:
    // 1. "none", the check has never been refreshed before.
    // 2. "enqueued", the check is waiting to be processed.
    // 3. "processing", the check is in the midst of being processed.
    // 4. "success", the check has succeeded and finished processing - refresh data is
    // available.
    // 5. "abandoned", the check has failed to process.
    return status.getStatus().equals("abandoned") ||
status.getStatus().equals("success");
}
// Enqueues a Trusted Advisor check refresh. Periodically polls the check refresh
// status for completion.
public static TrustedAdvisorCheckResult getFreshTACheckResult(final String checkId)
    throws InterruptedException {
    refreshTACheck(checkId);
    while(!isTACheckStatusInTerminalState(checkId)) {
        Thread.sleep(10000);
    }
    return getTACheckResult(checkId);
}
```

```
// Retrieves fresh TA check data whenever possible.
// Note: Some checks are refreshed automatically, and they cannot be refreshed by using
// this operation. This method
// is only functional for checks that can be refreshed using the
// RefreshTrustedAdvisorCheck operation.
public static void pollForTACheckResultChanges(final String checkId) throws
    InterruptedException {
    String checkResultStatus = null;
    do {
        TrustedAdvisorCheckResult result = getFreshTACheckResult(checkId);
        if (checkResultStatus != null && !checkResultStatus.equals(result.getStatus()))
        {
            break;
        }
        checkResultStatus = result.getStatus();
        // The rule refresh has completed, but due to throttling rules the checks may
        not be refreshed again
        // for a short period of time.
        // Since we only submitted one checkId to getTARefreshStatus, just retrieve the
        only element in the list.
        TrustedAdvisorCheckRefreshStatus refreshStatus =
            getTARefreshStatus(checkId).get(0);
        Thread.sleep(refreshStatus.getMillisUntilNextRefreshable());
    } while(true);
    // Signal that a TA check has changed check result status here.
}
```

Request a Trusted Advisor check result

After you select the check for the detailed results that you want, you submit a request by using the [DescribeTrustedAdvisorCheckResult](#) API operation.

Tip

The names and descriptions for Trusted Advisor checks are subject to change. We recommend that you specify the check ID in your code to uniquely identify a check. You can use the [DescribeTrustedAdvisorChecks](#) API operation to get the check ID.

The following Java code snippet uses the `DescribeTrustedAdvisorChecksResult` instance referenced by the variable `result`, which was obtained in the preceding code snippet. Rather

than defining a check interactively through a user interface. After you submit the request to run the snippet submits a request for the first check in the list to be run by specifying an index value of 0 in each `result.getChecks().get(0)` call. Next, the code defines an instance of `DescribeTrustedAdvisorCheckResultRequest`, which it passes to an instance of `DescribeTrustedAdvisorCheckResultResult` called `checkResult`. You can use the member structures of this data type to view the results of the check.

```
// Request a Trusted Advisor Check Result
public static TrustedAdvisorCheckResult getTACheckResult(final String checkId) {
    DescribeTrustedAdvisorCheckResultRequest request = new
    DescribeTrustedAdvisorCheckResultRequest()
        // Possible language parameters: "en" (English), "ja" (Japanese),
        "fr" (French), "zh" (Chinese)
        .withLanguage("en")
        .withCheckId(checkId);
    DescribeTrustedAdvisorCheckResultResult requestResult =
    createClient().describeTrustedAdvisorCheckResult(request);
    return requestResult.getResult();
}
```

Note: Requesting a Trusted Advisor Check Result doesn't generate updated results data.

Show details of a Trusted Advisor check

The following Java code snippet iterates over the `DescribeTrustedAdvisorCheckResultResult` instance returned in the previous section to get a list of resources flagged by the Trusted Advisor check.

```
// Show ResourceIds for flagged resources.
for (TrustedAdvisorResourceDetail flaggedResource :
    result1.getResult().getFlaggedResources())
{
    System.out.println(
        "The resource for this ResourceID has been flagged: " +
        flaggedResource.getResourceId());
}
```


Organizational view for AWS Trusted Advisor

Important

End of Support Notice: Developer Support will be discontinued January 1, 2027. Customers with Developer Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information, see [Business Support+ plan details](#)

End of Support Notice: Business Support will be discontinued January 1, 2027. Customers with Business Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information see, [Business Support+ plan details](#)

End of Support Notice: On January 1, 2027, AWS will discontinue Enterprise On-Ramp. Throughout 2026, Enterprise On-Ramp customers will be automatically upgraded to AWS Enterprise Support during contract renewal or in periodic batches. Customers will receive an email notification a month before their upgrade. No further action is required. Enterprise Support provides designated TAM assignment, 15-minute response times, and AWS Security Incident Response available at no additional cost, all at a lower \$5,000 minimum (reduced from \$15,000). For more information, see [AWS Enterprise Support plan details](#).

For more information, see [Developer, Business, and Enterprise On-Ramp end of support](#). Developer Support, Business Support, and Enterprise On-Ramp will remain available in the AWS GovCloud (US) Region.

Organizational view lets you view Trusted Advisor checks for all accounts in your [AWS Organizations](#). After you enable this feature, you can create reports to aggregate the check results for all member accounts in your organization. The report includes a summary of check results and information about affected resources for each account. For example, you can use the reports to identify which accounts in your organization are using AWS Identity and Access Management (IAM) with the IAM Use check or whether you have recommended actions for Amazon Simple Storage Service (Amazon S3) buckets with the Amazon S3 Bucket Permissions check.

Topics

- [Prerequisites](#)
- [Enable organizational view](#)
- [Refresh Trusted Advisor checks](#)
- [Create organizational view reports](#)
- [View the report summary](#)
- [Download an organizational view report](#)
- [Disable organizational view](#)
- [Using IAM policies to allow access to organizational view](#)
- [Using other AWS services to view Trusted Advisor reports](#)

Prerequisites

You must meet the following requirements to enable organizational view:

- Your accounts must be members of an [AWS Organizations](#).
- Your organization must have all features enabled for Organizations. For more information, see [Enabling all features in your organization](#) in the *AWS Organizations User Guide*.
- The management account in your organization must have an AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan. You can find your support plan from the AWS Support Center or from the [Support plans](#) page. See [Compare AWS Support plans](#).
- You must sign in as a user in the [management account](#) (or [assumed equivalent role](#)). Whether you sign in as an IAM user or an IAM role, you must have a policy with the required permissions. See [Using IAM policies to allow access to organizational view](#).

Enable organizational view

After you meet the prerequisites, follow these steps to enable organizational view. After you enable this feature, the following happens:

- Trusted Advisor is enabled as a *trusted service* in your organization. For more information, see [Enabling trusted access with other AWS services](#) in the *AWS Organizations User Guide*.
- The `AWSServiceRoleForTrustedAdvisorReporting` service-linked-role is created for you in the management account in your organization. This role includes the permissions that Trusted

Advisor needs to call Organizations on your behalf. This service-linked role is locked, and you can't delete it manually. For more information, see [Using service-linked roles for Trusted Advisor](#).

You enable organizational view from the Trusted Advisor console.

To enable organizational view

1. Sign in as an administrator in the organization's management account and open the AWS Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor>.
2. In the navigation pane, under **Preferences**, choose **Your organization**.
3. Under **Enable trusted access with AWS Organizations**, turn on **Enabled**.

Note

Enabling organizational view for the management account doesn't provide the same checks for all member accounts. For example, if your member accounts all have Basic Support, those accounts won't have the same checks available as your management account. The AWS Support plan determines which Trusted Advisor checks are available for an account.

Refresh Trusted Advisor checks

Before you create a report for your organization, we recommend that you refresh the statuses of your Trusted Advisor checks. You can download a report without refreshing your Trusted Advisor checks, but your report might not have the latest information.

If you have an AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, Trusted Advisor automatically refreshes the checks in your account on a weekly basis.

Note

If you have accounts in your organization that have a Developer or Basic support plan, a user for those accounts must sign in to the Trusted Advisor console to refresh the checks. You can't refresh checks for all accounts from the organization's management account.

To refresh Trusted Advisor checks

1. Navigate to the AWS Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor>.
2. On the **Trusted Advisor Recommendations** page, choose the **Refresh all checks**. This refreshes all checks in your account.

You can also refresh specific checks in the following ways:

- Use the [RefreshTrustedAdvisorCheck](#) API operation.
- Choose the refresh icon



for an individual check.

Create organizational view reports

After you enable organizational view, you can create reports so that you can view Trusted Advisor check results for your organization.

You can create up to 50 reports. If you create reports beyond this quota, Trusted Advisor deletes the earliest report. You can't recover deleted reports.

To create organizational view reports

1. Sign in to the organization's management account and open the AWS Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor>.
2. In the navigation pane, choose **Organizational View**.
3. Choose **Create report**.
4. By default, the report includes all AWS Regions, check categories, checks, and resource statuses. On the **Create report** page, you can use the filter options to customize your report. For example, you can clear the **All** option for **Region**, and then specify the individual Regions to include in the report.
 - a. Enter a **Name** for the report.
 - b. For **Format**, choose **JSON** or **CSV**.
 - c. For **Region**, specify the AWS Regions or choose **All**.

- d. For **Check category**, choose the check category or choose **All**.
- e. For **Checks**, choose the specific checks for that category or choose **All**.

 Note

The **Check category** filter overrides the **Checks** filter. For example, if you choose the **Security** category and then choose a specific check name, your report includes all check results for that category. To create a report for only specific checks, keep the default **All** value for **Check category** and then choose your check names.

- f. For **Resource status**, choose the status to filter, such as **Warning**, or choose **All**.
5. For **AWS Organizations**, select the organizational units (OUs) to include in your report. For more information about OUs, see [Managing organizational units](#) in the *AWS Organizations User Guide*.
6. Choose **Create report**.

Example: Create report filter options

The following example creates a JSON report for the following:

- Three AWS Regions
- All **Security** and **Performance** checks

Report filters

Choose the filter options for your report.

Report name

Trusted-Advisor-report

The report name can be up to 100 characters and can't start with a hyphen. Valid characters: A-Z, a-z, 0-9, and - (hyphen)

Format

JSON

Region

Choose Region

us-east-1 X us-east-2 X us-west-1 X

Check category

Choose category

Security X Performance X

Checks

Choose check names

Resource status

Choose status

All X

In the following example, the report includes the **support-team** organizational unit and one AWS account that are part of the organization.

AWS organization

You can select the organizational units (OUs) and individual AWS accounts to include in your report.

Organizational structure

▼ ☒  **Root**
r-xa9c

▶ ☐  **instance-management**
ou-xa9c-example1

▼ ☒  **support-team**
ou-xa9c-example2

☒  **Jane Doe**
111122223333 | janedoe@example.com

☒  **Mateo Jackson**
444455556666 | mateojackson@example.com

▶ ☐  **security-team**
ou-xa9c-example3

☒  **Ana Carolina Silva**
777788889999 | anacarolinasilva@example.com

Notes

- The amount of time it takes to create the report depends on the number of accounts in the organization and the number of resources in each account.
- You can't create more than one report at a time unless the current report has been running for more than 6 hours.
- Refresh the page if you don't see the report appear on the page.


View the report summary

After the report is ready, you can view the report summary from the Trusted Advisor console. This lets you quickly view the summary of your check results across your organization.

To view the report summary

1. Sign in to the organization's management account and open the AWS Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor>.
2. In the navigation pane, choose **Organizational View**.
3. Choose the report name.
4. On the **Summary** page, view the check statuses for each category. You can also choose **Download report**.

Example: Report summary for an organization



Download an organizational view report

After your report is ready, download it from the Trusted Advisor console. The report is a .zip file that contains three files:

- `summary.json` – Contains a summary of the check results for each check category.
- `schema.json` – Contains the schema for the specified checks in the report.
- A resources file (.json or .csv) – Contains detailed information about the check statuses for resources in your organization.

To download an organizational view report

1. Sign in to the organization's management account and open the AWS Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor>.
2. In the navigation pane, choose **Organizational View**.

The **Organizational View** page displays the available reports to download.

3. Select a report, choose **Download report**, and then save the file. You can only download one report at a time.

Organizational View

With AWS organizations, you can create reports for check results across all AWS accounts within an organization. This provides you a centralized view for all AWS Trusted Advisor checks. You can also view and download reports on this page. Use this report to identify issues and take action for accounts in your organization. [Learn more](#) .

Reports (50)

Create reportDownload report

	Report name	Date generated	Status	Format
☐	all-regions-check-report	June 15, 2021 18:43:42	Success	JSON
☐	json-us-east-1-region-only	June 14, 2021 20:54:29	Success	JSON
☐	security-checks-only-all-accounts	June 10, 2021 03:33:59	Success	JSON

4. Unzip the file.
5. Use a text editor to open the .json file or a spreadsheet application to open the .csv file.

Note

You might receive multiple files if your report is 5 MB or larger.

Example: summary.json file

The `summary.json` file shows the number of accounts in the organization and the statuses of the checks in each category.

Trusted Advisor uses the following color code for check results:

- Green – Trusted Advisor doesn't detect an issue for the check.
- Yellow – Trusted Advisor detects a possible issue for the check.
- Red – Trusted Advisor detects an error and recommends an action for the check.
- Blue – Trusted Advisor can't determine the status of the check.

In the following example, two checks are Red, one is Green, and one is Yellow.

```
{
  "numAccounts": 3,
  "filtersApplied": {
    "accountIds": ["123456789012", "111122223333", "111111111111"],
    "checkIds": "All",
    "categories": [
      "security",
      "performance"
    ],
    "statuses": "All",
    "regions": [
      "us-west-1",
      "us-west-2",
      "us-east-1"
    ],
    "organizationalUnitIds": [
      "ou-xa9c-EXAMPLE1",
      "ou-xa9c-EXAMPLE2"
    ]
  },
  "categoryStatusMap": {
    "security": {
      "statusMap": {
        "ERROR": {
          "name": "Red",
          "count": 2
        },
        "OK": {
          "name": "Green",
          "count": 1
        },
        "WARN": {
          "name": "Yellow",
          "count": 1
        }
      }
    }
  }
}
```

```

        }
      },
      "name": "Security"
    }
  },
  "accountStatusMap": {
    "123456789012": {
      "security": {
        "statusMap": {
          "ERROR": {
            "name": "Red",
            "count": 2
          },
          "OK": {
            "name": "Green",
            "count": 1
          },
          "WARN": {
            "name": "Yellow",
            "count": 1
          }
        }
      },
      "name": "Security"
    }
  }
}

```

Example: schema.json file

The `schema.json` file includes the schema for the checks in the report. The following example includes the IDs and properties for the IAM Password Policy (Yw2K9puPzl) and IAM Key Rotation (DqdJqYeRm5) checks.

```

{
  "Yw2K9puPzl": [
    "Password Policy",
    "Uppercase",
    "Lowercase",
    "Number",
    "Non-alphanumeric",
    "Status",
    "Reason"
  ]
}

```

```

    ],
    "DqdJqYeRm5": [
        "Status",
        "IAM User",
        "Access Key",
        "Key Last Rotated",
        "Reason"
    ],
    ...
}

```

Example: resources.csv file

The `resources.csv` file includes information about resources in the organization. This example shows some of the data columns that appear in the report, such as the following:

- Account ID of the affected account
- The Trusted Advisor check ID
- The resource ID
- Timestamp of the report
- The full name of the Trusted Advisor check
- The Trusted Advisor check category
- The account ID of the parent organizational unit (OU) or root

AccountId	CheckId	ResourceId	TimeStamp	CheckName	Category
1.11122E+11	Qch7DwouX1	LnW14f1M40NMjMMLvY5	1.58983E+12	Low Utilization Amazon EC2 Instances	Cost Optimizing
1.11122E+11	HCP4007JGY	dJrQZXw36ZdswBeo9WUL	1.58983E+12	Security Groups - Specific Ports Unrestricted	Security
1.11122E+11	HCP4007JGY	1hzakmTbWd5UmAM_a0L	1.58983E+12	Security Groups - Specific Ports Unrestricted	Security
4.44456E+11	1iG5NDGVre	dJrQZXw36ZdswBeo9WUL	1.58983E+12	Security Groups - Unrestricted Access	Security
4.44456E+11	1iG5NDGVre	1hzakmTbWd5UmAM_a0L	1.58983E+12	Security Groups - Unrestricted Access	Security
4.44456E+11	Pfx0RwqBli	vioZmlba45kf2JWle_W0j5	1.58983E+12	Amazon S3 Bucket Permissions	Security
4.44456E+11	Pfx0RwqBli	wAvASS3YOWy6WWxIBHf	1.58983E+12	Amazon S3 Bucket Permissions	Security
1.23457E+11	Pfx0RwqBli	Llc4zRaUSiIGRSImqaMa5V	1.58983E+12	Amazon S3 Bucket Permissions	Security
1.23457E+11	Pfx0RwqBli	gWB27TMXof2evYzMSYBg	1.58983E+12	Amazon S3 Bucket Permissions	Security
7.77789E+11	Pfx0RwqBli	M3LBsF0e15Cl9Mxppapcx	1.58983E+12	Amazon S3 Bucket Permissions	Security
7.77789E+11	Yw2K9puPzl	47DEQpj8HBSa-_TlMw-5Jc	1.58983E+12	IAM Password Policy	Security
7.77789E+11	H7lgTzjTYb	1xHQ5ovV8bS0H1Z-t7Kbik	1.58983E+12	Amazon EBS Snapshots	Fault Tolerance
7.77789E+11	wuy7G1zxql	10F6p6VAF0F-MuL6Dc-dl1	1.58983E+12	Amazon EC2 Availability Zone Balance	Fault Tolerance

The resources file only contains entries if a check result exists at the resource level. You might not see checks in the report for the following reasons:

- Some checks, such as **MFA on Root Account**, don't have resources and won't appear in the report. Checks without resources appear in the `summary.json` file instead.
- Some checks only show resources if they are Red or Yellow. If all resources are Green, they might not appear in your report.
- If an account isn't enabled for a service that requires the check, the check might not appear in the report. For example, if you're not using Amazon Elastic Compute Cloud Reserved Instances in your organization, the Amazon EC2 Reserved Instance Lease Expiration check won't appear in your report.
- The account hasn't refreshed check results. This might happen when users with a Basic or Developer support plan sign in to the Trusted Advisor console for the first time. If you have an AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, it can take up to one week from account sign up for users to see check results. For more information, see [Refresh Trusted Advisor checks](#).
- If only the organization's management account enabled recommendations for checks, the report won't include resources for other accounts in the organization.

For the resources file, you can use common software such as Microsoft Excel to open the `.csv` file format. You can use the `.csv` file for one-time analysis of all checks across all accounts in your organization. If you want to use your report with an application, you can download the report as a `.json` file instead.

The `.json` file format provides more flexibility than the `.csv` file format for advanced use cases such as aggregation and advanced analytics with multiple datasets. For example, you can use a SQL interface with an AWS service such as Amazon Athena to run queries on your reports. You can also use Amazon Quick to create dashboards and visualize your data. For more information, see [Using other AWS services to view Trusted Advisor reports](#).

Disable organizational view

Follow this procedure to disable organizational view. You must sign in to the organization's management account or assume a role with the required permissions to disable this feature. You can't disable this feature from another account in the organization.


After you disable this feature, the following happens:

- Trusted Advisor is removed as a trusted service in Organizations.

- The `AWSServiceRoleForTrustedAdvisorReporting` service-linked role is unlocked in the organization's management account. This means you can delete it manually, if needed.
- You can't create, view, or download reports for your organization. To access previously created reports, you must reenable organizational view from the Trusted Advisor console. See [Enable organizational view](#).

To disable organizational view for Trusted Advisor

1. Sign in to the organization's management account and open the AWS Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor>.
2. In the navigation pane, choose **Preferences**.
3. Under **Organizational View**, choose **Disable organizational view**.



After you disable organizational view, Trusted Advisor no longer aggregates checks from other AWS accounts in your organization. However, the `AWSServiceRoleForTrustedAdvisorReporting` service-linked role remains on the organization's management account until you delete it through the IAM console, IAM API, or AWS Command Line Interface (AWS CLI). For more information, see [Deleting a service-linked role](#) in the *IAM User Guide*.

Note

You can use other AWS services to query and visualize your data for organizational view reports. For more information, see the following resources:

- [View AWS Trusted Advisor recommendations at scale with AWS Organizations](#) in the *AWS Management & Governance Blog*

- [Using other AWS services to view Trusted Advisor reports](#)

Using IAM policies to allow access to organizational view

You can use the following AWS Identity and Access Management (IAM) policies to allow users or roles in your account access to organizational view in AWS Trusted Advisor.

Example: Full access to organizational view

The following policy allows full access to the organizational view feature. A user with these permissions can do the following:

- Enable and disable organizational view
- Create, view, and download reports

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ReadStatement",
      "Effect": "Allow",
      "Action": [
        "organizations:ListAccountsForParent",
        "organizations:ListAccounts",
        "organizations:ListRoots",
        "organizations:DescribeOrganization",
        "organizations:ListOrganizationalUnitsForParent",
        "organizations:ListAWSServiceAccessForOrganization",
        "trustedadvisor:DescribeAccount",
        "trustedadvisor:DescribeChecks",
        "trustedadvisor:DescribeCheckSummaries",
        "trustedadvisor:DescribeAccountAccess",
        "trustedadvisor:DescribeOrganization",
        "trustedadvisor:DescribeReports",
        "trustedadvisor:DescribeServiceMetadata",
        "trustedadvisor:DescribeOrganizationAccounts",
        "trustedadvisor:ListAccountsForParent",

```



```

        "trustedadvisor:ListRoots",
        "trustedadvisor:ListOrganizationalUnitsForParent"
    ],
    "Resource": "*"
},
{
    "Sid": "CreateReportStatement",
    "Effect": "Allow",
    "Action": [
        "trustedadvisor:GenerateReport"
    ],
    "Resource": "*"
},
{
    "Sid": "ManageOrganizationalViewStatement",
    "Effect": "Allow",
    "Action": [
        "organizations:EnableAWSServiceAccess",
        "organizations:DisableAWSServiceAccess",
        "trustedadvisor:SetOrganizationAccess"
    ],
    "Resource": "*"
},
{
    "Sid": "CreateServiceLinkedRoleStatement",
    "Effect": "Allow",
    "Action": "iam:CreateServiceLinkedRole",
    "Resource": "arn:aws:iam::*:role/aws-service-role/
reporting.trustedadvisor.amazonaws.com/AWSServiceRoleForTrustedAdvisorReporting"
}
]
}

```

Example: Read access to organizational view

The following policy allows read-only access to organizational view for Trusted Advisor. A user with these permissions can only view and download existing reports.

JSON

```

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

```

```

    "Statement": [
      {
        "Sid": "ReadStatement",
        "Effect": "Allow",
        "Action": [
          "organizations:ListAccountsForParent",
          "organizations:ListAccounts",
          "organizations:ListRoots",
          "organizations:DescribeOrganization",
          "organizations:ListOrganizationalUnitsForParent",
          "organizations:ListAWSServiceAccessForOrganization",
          "trustedadvisor:DescribeAccount",
          "trustedadvisor:DescribeChecks",
          "trustedadvisor:DescribeCheckSummaries",
          "trustedadvisor:DescribeAccountAccess",
          "trustedadvisor:DescribeOrganization",
          "trustedadvisor:DescribeReports",
          "trustedadvisor:ListAccountsForParent",
          "trustedadvisor:ListRoots",
          "trustedadvisor:ListOrganizationalUnitsForParent"
        ],
        "Resource": "*"
      }
    ]
  }
}

```

You can also create your own IAM policy. For more information, see [Creating IAM Policies](#) in the *IAM User Guide*.

Note

If you enabled AWS CloudTrail in your account, the following roles can appear in your log entries:

- `AWSServiceRoleForTrustedAdvisorReporting` – The service-linked role that Trusted Advisor uses to access accounts in your organization.
- `AWSServiceRoleForTrustedAdvisor` – The service-linked role that Trusted Advisor uses to access services in your organization.

For more information about service-linked roles, see [Using service-linked roles for Trusted Advisor](#).

Using other AWS services to view Trusted Advisor reports

Follow this tutorial to upload and view your data by using other AWS services. In this topic, you create an Amazon Simple Storage Service (Amazon S3) bucket to store your report and an CloudFormation template to create resources in your account. Then, you can use Amazon Athena to analyze or run queries for your report or Quick to visualize that data in a dashboard.

For information and examples for visualizing your report data, see the [View AWS Trusted Advisor recommendations at scale with AWS Organizations](#) in the *AWS Management & Governance Blog*.

Prerequisites

Before you start this tutorial, you must meet the following requirements:

- Sign in as an AWS Identity and Access Management (IAM) user with administrator permissions.
- Use the US East (N. Virginia) AWS Region to quickly set up your AWS services and resources.
- Create an Quick account. For more information, see [Getting Started with Data Analysis in Quick](#) in the *Amazon Quick User Guide*.

Upload the report to Amazon S3

After you download your `resources.json` report, upload the file to Amazon S3. You must use a bucket in the US East (N. Virginia) Region.

To upload the report to an Amazon S3 bucket

1. Sign in to the AWS Management Console at <https://console.aws.amazon.com/>.
2. Use the **Region selector** and choose the US East (N. Virginia) Region.
3. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
4. From the list of buckets, choose an S3 bucket, and then copy the name. You use the name in the next procedure.

5. On the *bucket-name* page, choose **Create folder**, enter the name **folder1**, and then choose **Save**.
6. Choose the **folder1**.
7. In **folder1**, choose **Upload** and choose the `resources.json` file.
8. Choose **Next**, keep the default options, and then choose **Upload**.

 Note

If you upload a new report to this bucket, rename the `.json` files each time you upload them so that you don't override the existing reports. For example, you can add the timestamp to each file, such as `resources-timestamp.json`, `resources-timestamp2.json`, and so on.

Create your resources using AWS CloudFormation

After you upload your report to Amazon S3, upload the following YAML template to CloudFormation. This template tells CloudFormation what resources to create for your account so that other services can use the report data in the S3 bucket. The template creates resources for IAM, AWS Lambda, and AWS Glue.

To create your resources with CloudFormation

1. Download the [trusted-advisor-reports-template.zip](#) file.
2. Unzip the file.
3. Open the template file in a text editor.
4. For the `BucketName` and `FolderName` parameters, replace the values for *your-bucket-name-here* and *folder1* with the bucket name and folder name in your account.
5. Save the file.
6. Open the CloudFormation console at <https://console.aws.amazon.com/cloudformation>.
7. If you haven't already, in the **Region selector**, choose the US East (N. Virginia) Region.
8. In the navigation pane, choose **Stacks**.
9. Choose **Create stack** and choose **With new resources (standard)**.
10. On the **Create stack** page, under **Specify template**, choose **Upload a template file**, and then choose **Choose file**.

11. Choose the YAML file and choose **Next**.
12. On the **Specify stack details** page, enter a stack name such as **Organizational-view-Trusted-Advisor-reports**, and choose **Next**.
13. On the **Configure stack options** page, keep the default options, and then choose **Next**.
14. On the **Review Organizational-view-Trusted-Advisor-reports** page, review your options. At the bottom of the page, select the check box for **I acknowledge that CloudFormation might create IAM resources**.
15. Choose **Create stack**.

The stack takes about 5 minutes to create.

16. After the stack creates successfully, the **Resources** tab appears like the following example.

Trusted-Advisor-reports

Delete

Update

Stack actions ▾

Stack info

Events

Resources

Outputs

Parameters

Template

Change sets

Resources (12)

Q Search resources

Logical ID	Physical ID	Type	Status
AWSPutS3TANotification	2020/05/27/[\$LATEST]5bfd3cb8b29a4b85bc0f8d861EXAMPLE1	Custom::AWSPutS3TANotification	✔ CREATE_COMPLETE
AWSS3TAEventLambdaPermission	Trusted-Advisor-reports-AWSS3TAEventLambdaPermission-10KT2EXAMPLE1	AWS::Lambda::Permission	✔ CREATE_COMPLETE
AWSS3TALambdaExecutor	Trusted-Advisor-reports-AWSS3TALambdaExecutor-1BJCOEXAMPLE1	AWS::IAM::Role	✔ CREATE_COMPLETE
AWSS3TANotification	Trusted-Advisor-reports-AWSS3TANotification-15J3KEXAMPLE1	AWS::Lambda::Function	✔ CREATE_COMPLETE
AWSSstartTACrawler	2020/05/27/[\$LATEST]66726149d3d64a1f9242cdccEXAMPLE1	Custom::AWSSstartTACrawler	✔ CREATE_COMPLETE
AWSTACrawler	AWSTACrawler	AWS::Glue::Crawler	✔ CREATE_COMPLETE

Query the data in Amazon Athena

After you have your resources, you can view the data in Athena. Use Athena to create queries and analyze the results of the report, such as looking up specific check results for accounts in the organization.


 **Notes**

- Use the US East (N. Virginia) Region.

- If you're new to Athena, you must specify a query result location before you can run a query for your report. We recommend that you specify a different S3 bucket for this location. For more information, see [Specifying a query result location](#) in the *Amazon Athena User Guide*.

To query the data in Athena

1. Open the Athena console at <https://console.aws.amazon.com/athena/>.
2. If you haven't already, in the **Region selector**, choose the US East (N. Virginia) Region.
3. Choose **Saved Queries** and in search field, enter **Show sample**.
4. Choose the query that appears, such as **Show sample entries of TA report**.



The query should look like the following.

```
SELECT * FROM "athenatacfn"."folder1" limit 10
```

5. Choose **Run query**. Your query results appear.

Example: Athena query

The following example shows 10 sample entries from the report.

The screenshot displays the Amazon Athena console interface. At the top, there's a query editor with a tab labeled 'New query 1'. A green checkmark icon and the text 'Show sample ent...' are visible next to the tab. The query entered is: `SELECT * FROM "athenatacfn"."folder1" limit 10`. Below the query editor, there are buttons for 'Run query', 'Save as', and 'Create'. To the right of these buttons, it says '(Run time: 0.83 seconds, Data scanned: 94.75 KB)'. Below the buttons, there's a hint: 'Use Ctrl + Enter to run query, Ctrl + Space to autocomplete'. To the right of the hint are buttons for 'Format query' and 'Clear'. Below the query editor, there's a section titled 'Results' with a download icon and a share icon. The results are displayed in a table with 7 columns: 'volume type', 'checkname', 'accountid', 'category', 'issuppressed', and 'snapshot'. The table contains 10 rows of data.

	volume type	checkname	accountid	category	issuppressed	snapshot
1	General purpose(SSD)	Underutilized Amazon EBS Volumes	123456789012	Cost Optimizing	false	snap-0d4
2	General purpose(SSD)	Underutilized Amazon EBS Volumes	123456789012	Cost Optimizing	false	snap-06b
3	General purpose(SSD)	Underutilized Amazon EBS Volumes	123456789012	Cost Optimizing	false	
4	General purpose(SSD)	Underutilized Amazon EBS Volumes	123456789012	Cost Optimizing	false	
5	General purpose(SSD)	Underutilized Amazon EBS Volumes	123456789012	Cost Optimizing	false	snap-0ef4
6	General purpose(SSD)	Underutilized Amazon EBS Volumes	123456789012	Cost Optimizing	false	snap-0a5
7	General purpose(SSD)	Underutilized Amazon EBS Volumes	123456789012	Cost Optimizing	false	snap-078
8	General purpose(SSD)	Underutilized Amazon EBS Volumes	123456789012	Cost Optimizing	false	
9	General purpose(SSD)	Underutilized Amazon EBS Volumes	123456789012	Cost Optimizing	false	snap-0ff6
10	General purpose(SSD)	Underutilized Amazon EBS Volumes	123456789012	Cost Optimizing	false	

For more information, see [Running SQL Queries Using Amazon Athena](#) in the *Amazon Athena User Guide*.

Create a dashboard in Quick


You can also set up Quick so that you can view your data in a dashboard and visualize your report information.

Note


You must use the US East (N. Virginia) Region.

To create a dashboard in Quick


1. Navigate to the Quick console and sign in to your [account](#).
2. Choose **New analysis**, **New dataset**, and then choose **Athena**.
3. In the **New Athena data source** dialog box, enter a data source name such as **AthenaTA**, and then choose **Create data source**.



4. In the **Choose your table** dialog box, choose the **athenatacfn** table, choose **folder1**, and then choose **Select**.



5. In the **Finish data set creation** dialog box, choose **Directly query your data**, and then choose **Visualize**.



You can now create a dashboard in Quick. For more information, see [Working with Dashboards](#) in the *Amazon Quick User Guide*.

Example: Quick dashboard

The following example dashboard shows information about the Trusted Advisor checks, such as the following:

- Affected account IDs
- Summary by AWS Regions
- Check categories
- Check statuses
- Number of entries in the report for each account



Note

If you have permission errors while creating your dashboard, make sure that Quick can use Athena. For more information, see [I Can't Connect to Amazon Athena](#) in the *Amazon Quick User Guide*.

For more information and examples for visualizing your report data, see the [View AWS Trusted Advisor recommendations at scale with AWS Organizations](#) in the *AWS Management & Governance Blog*.

Troubleshooting

If you have issues with this tutorial, see the following troubleshooting tips.

I'm not seeing the latest data in my report

When you create a report, the organizational view feature doesn't automatically refresh the Trusted Advisor checks in your organization. To get the latest check results, refresh the checks for the management account and each member account in the organization. For more information, see [Refresh Trusted Advisor checks](#).

I have duplicate columns in the report

The Athena console might show the following error in your table if your report has duplicate columns.

```
HIVE_INVALID_METADATA: Hive metadata for table folder1 is invalid: Table descriptor contains duplicate columns
```

For example, if you added a column in your report that already exists, this can cause issues when you try to view the report data in the Athena console. You can follow these steps to fix this issue.

Find duplicate columns

You can use the AWS Glue console to view the schema and quickly identify if you have duplicate columns in your report.

To find duplicate columns

1. Open the AWS Glue console at <https://console.aws.amazon.com/glue/>.

2. If you haven't already, in the **Region selector**, choose the US East (N. Virginia) Region.
3. In the navigation pane, choose **Tables**.
4. Choose your folder name, such as *folder1*, and then under **Schema**, view the values for **Column name**.

If you have a duplicate column, you must upload a new report to your Amazon S3 bucket. See the following [Upload a new report](#) section.

Upload a new report

After you identify the duplicate column, we recommend that you replace the existing report with a new one. This ensures that the resources created from this tutorial use the latest report data from your organization.

To upload a new report

1. If you haven't already, refresh your Trusted Advisor checks for the accounts in your organization. See [Refresh Trusted Advisor checks](#).
2. Create and download another JSON report in the Trusted Advisor console. See [Create organizational view reports](#). You must use a JSON file for this tutorial.
3. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
4. Choose your Amazon S3 bucket and choose the *folder1* folder.
5. Select the previous *resources.json* reports and choose **Delete**.
6. In the **Delete objects** page, under **Permanently delete objects?**, enter **permanently delete**, and then choose **Delete objects**.
7. In your S3 bucket, choose **Upload** and then specify the new report. This action automatically updates your Athena table and AWS Glue crawler resources with the latest report data. It can take a few minutes to refresh your resources.
8. Enter a new query in the Athena console. See [Query the data in Amazon Athena](#).

Note

If you still have issues with this tutorial, you can create a technical support case in the [AWS Support Center](#).

View AWS Trusted Advisor checks powered by AWS Config

Important

End of Support Notice: Developer Support will be discontinued January 1, 2027. Customers with Developer Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information, see [Business Support+ plan details](#)

End of Support Notice: Business Support will be discontinued January 1, 2027. Customers with Business Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information see, [Business Support+ plan details](#)

End of Support Notice: On January 1, 2027, AWS will discontinue Enterprise On-Ramp. Throughout 2026, Enterprise On-Ramp customers will be automatically upgraded to AWS Enterprise Support during contract renewal or in periodic batches. Customers will receive an email notification a month before their upgrade. No further action is required. Enterprise Support provides designated TAM assignment, 15-minute response times, and AWS Security Incident Response available at no additional cost, all at a lower \$5,000 minimum (reduced from \$15,000). For more information, see [AWS Enterprise Support plan details](#).

For more information, see [Developer, Business, and Enterprise On-Ramp end of support](#). Developer Support, Business Support, and Enterprise On-Ramp will remain available in the AWS GovCloud (US) Region.

AWS Config is a service that continually assesses, audits, and evaluates your resource configurations for your desired settings. AWS Config provides managed rules, which are predefined, customizable compliance checks that AWS Config uses to evaluate if your AWS resources comply with common best practices.

The AWS Config console guides you through the configuration and activation of managed rules. You can also use the AWS Command Line Interface (AWS CLI) or AWS Config API to pass the JSON code that defines your configuration of a managed rule. You can customize the behavior of a

managed rule to suit your needs. You can customize the rule's parameters to define attributes that your resources must have to comply with the rule. To learn more about enabling AWS Config, see the [AWS Config Developer Guide](#).

AWS Config managed rules power a set of Trusted Advisor checks across all categories. When you enable certain managed rules, the corresponding Trusted Advisor checks are automatically enabled. To see which Trusted Advisor checks are powered by specific AWS Config managed rules, see [AWS Trusted Advisor check reference](#).

The AWS Config powered checks are available to customers with an AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan. If you enable AWS Config and you have one of these AWS Support plans, then you automatically see recommendations powered by corresponding deployed AWS Config managed rules.

Note

Results for these checks are automatically refreshed based on change-triggered updates to AWS Config managed rules. Refresh requests are not allowed. Currently, you can't exclude resources from these checks.

Troubleshooting

If you have issues with this integration, see the following troubleshooting information.

Contents

- [I just enabled recording and managed rules for AWS Config, but I don't see corresponding Trusted Advisor checks.](#)
- [I deployed the same AWS Config managed rule twice, what will I see in Trusted Advisor?](#)
- [I turned off recording for AWS Config in an AWS Region. What will I see in Trusted Advisor?](#)

I just enabled recording and managed rules for AWS Config, but I don't see corresponding Trusted Advisor checks.

After the AWS Config rule generates evaluation results, you see the results in Trusted Advisor in near real-time. If you have issues with this feature, create a technical support case in the [AWS Support Center](#).

I deployed the same AWS Config managed rule twice, what will I see in Trusted Advisor?

You see separate entries in the Trusted Advisor check results for each managed rule that you install.

I turned off recording for AWS Config in an AWS Region. What will I see in Trusted Advisor?

If you turned off resource recording for AWS Config in an AWS Region, then Trusted Advisor no longer receives data for corresponding managed rules and checks in that Region. Existing managed rule results remain in AWS Config and in Trusted Advisor until AWS Config expires, based on the recorder retention policy. If you delete a managed rule, then the Trusted Advisor check data usually deletes in near real-time.

Viewing AWS Security Hub CSPM controls in AWS Trusted Advisor

After you enable AWS Security Hub CSPM for your AWS account, you can view your security controls and their findings in the Trusted Advisor console. You can use Security Hub CSPM controls to identify security vulnerabilities in your account in the same way that you can use Trusted Advisor checks. You can view the check's status, the list of affected resources, and then follow Security Hub CSPM recommendations to address your security issues. You can use this feature to find security recommendations from Trusted Advisor and Security Hub CSPM in one convenient location.

Notes

- From Trusted Advisor, you can view controls in the AWS Foundational Security Best Practices security standard *except* for controls that have the Category: Recover > Resilience. For a list of supported controls, see [AWS Foundational Security Best Practices controls](#) in the *AWS Security Hub CSPM User Guide*.

For more information about the Security Hub CSPM categories, see [Control categories](#).

- Trusted Advisor onboarded Security Hub CSPM controls up to July 8, 2026. Controls released after July 8, 2026 are not yet onboarded to Trusted Advisor. You can find controls released after that date in the [Security Hub CSPM log](#).

Topics

- [Prerequisites](#)
- [View your Security Hub CSPM findings](#)
- [Refresh your Security Hub CSPM findings](#)
- [Disable Security Hub CSPM from Trusted Advisor](#)
- [Troubleshooting](#)

Prerequisites

You must meet the following requirements to enable the Security Hub CSPM integration with Trusted Advisor:

- You must have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan for this feature. You can find your support plan from the [AWS Support Center](#) or from the [Support plans](#) page. For more information, see [Compare AWS Support plans](#).
- You must enable resource recording in AWS Config for the AWS Regions that you want for your Security Hub CSPM controls. For more information, see [Enabling and configuring AWS Config](#).
- You must enable Security Hub CSPM and select the **AWS Foundational Security Best Practices v1.0.0** security standard. If you haven't done so already, see [Setting up AWS Security Hub CSPM](#) in the *AWS Security Hub CSPM User Guide*.

Note

If you already completed these prerequisites, you can skip to [View your Security Hub CSPM findings](#).

About AWS Organizations accounts

If you already completed the prerequisites for a management account, this integration is enabled automatically for all member accounts in your organization. Individual member accounts don't need to contact Support to enable this feature. However, member accounts in your organization must enable Security Hub CSPM if they want to see their findings in Trusted Advisor.

If you want to disable this integration for a specific member account, see [Disable this feature for AWS Organizations accounts](#).

View your Security Hub CSPM findings

After you enable Security Hub CSPM for your account, it can take up to 24 hours for your Security Hub CSPM findings to appear in the **Security** page of the Trusted Advisor console.

To view your Security Hub CSPM findings in Trusted Advisor

1. Navigate to the [Trusted Advisor console](#), and then choose the **Security** category.
2. In the **Search by keyword** field, enter the control name or description in the field.

Tip

For **Source**, you can choose **AWS Security Hub CSPM** to filter for Security Hub CSPM controls.

3. Choose the Security Hub CSPM control name to view the following information:
 - **Description** – Describes how this control checks your account for security vulnerabilities.
 - **Source** – Whether the check comes from AWS Trusted Advisor or AWS Security Hub CSPM. For Security Hub CSPM controls, you can find the control ID.
 - **Alert Criteria** – The status of the control. For example, if Security Hub CSPM detects an important issue, the status might be **Red: Critical or High**.
 - **Recommended Action** – Use the Security Hub CSPM documentation link to find the recommended steps to fix the issue.
 - **Security Hub CSPM resources** – You can find the resources in your account where Security Hub CSPM has detected an issue.

Notes

- Results for these checks are automatically refreshed at least once daily, and refresh requests are not allowed. It might take a few hours for changes to appear. You can use the Trusted Advisor console to exclude resources from checks that automatically refresh. You can use the [BatchUpdateRecommendationResourceExclusion](#) API to exclude resources from any check except Trusted Advisor Priority recommendation resources.
- The organizational view feature supports this integration with Security Hub CSPM. You can view your findings for your Security Hub CSPM controls across your organization, and

then create and download reports. For more information, see [Organizational view for AWS Trusted Advisor](#).

Example : Security Hub CSPM control for IAM user access key should not exist

The following is an example finding for a Security Hub CSPM control in the Trusted Advisor console.

The screenshot displays the AWS Trusted Advisor console interface for a Security Hub CSPM control. At the top, the control is titled "IAM root user access key should not exist" with a red status icon and a dropdown arrow. To the right, it indicates "Last updated: an hour ago" with refresh and share icons. Below the title, a description states: "Checks if the root user access key is available." The "Source" section identifies the control as "AWS Security Hub" with ID "IAM.4". The "Alert Criteria" section notes: "Red: Critical or High. Security Hub control failed." The "Recommended Action" section advises: "Follow the [Security Hub documentation](#) to fix the issue."

Below this information is a summary box for the control, showing "IAM root user access key should not exist (1)" and "1 of 1 resources failed this Security Hub control." It includes buttons for "Exclude & Refresh" and "Included items" with a dropdown arrow. Navigation controls show "1" item with left and right arrows and a settings gear icon.

A table lists the findings with the following columns: Status, Region, Resource, and Last Updated Time.

Status	Region	Resource	Last Updated Time
☐	us-east-1	AWS::::Account:123456789012	2021-12-12T19:56:26.305Z

Refresh your Security Hub CSPM findings

After you enable a security standard, it can take up to two hours for Security Hub CSPM to have findings for your resources. It can then take up to 24 hours for that data to appear in the Trusted Advisor console. If you recently enabled the **AWS Foundational Security Best Practices v1.0.0** security standard, check the Trusted Advisor console again later.

Note

- The refresh schedule for each Security Hub CSPM control is *periodic* or *change triggered*. Currently, you can't use the Trusted Advisor console or the AWS Support API to refresh

your Security Hub CSPM controls. For more information, see [Schedule for running security checks](#).

- Results for these checks are automatically refreshed at least once daily, and refresh requests are not allowed. It might take a few hours for changes to appear. You can use the Trusted Advisor console to exclude resources from checks that automatically refresh. You can use the [BatchUpdateRecommendationResourceExclusion](#) API to exclude resources from any check except Trusted Advisor Priority recommendation resources.

Disable Security Hub CSPM from Trusted Advisor

Follow this procedure if you don't want your Security Hub CSPM information to appear in the Trusted Advisor console. This procedure only disables the Security Hub CSPM integration with Trusted Advisor. It won't affect your configurations with Security Hub CSPM. You can continue to use the Security Hub CSPM console to view your security controls, resources, and recommendations.

To disable the Security Hub CSPM integration

1. Contact [AWS Support](#) and request to disable the Security Hub CSPM integration with Trusted Advisor.

After AWS Support disables this feature, Security Hub CSPM no longer sends data to Trusted Advisor. Your Security Hub CSPM data will be removed from Trusted Advisor.

2. If you want to enable this integration again, contact [AWS Support](#).

Disable this feature for AWS Organizations accounts

If you already completed the previous procedure for a management account, Security Hub CSPM integration is automatically removed from all member accounts in your organization. Individual member accounts in your organization don't need to contact AWS Support separately.

If you're a member account in an organization, you can contact Support to remove this feature from only your account.

Troubleshooting

If you're having issues with this integration, see the following troubleshooting information.

Contents

- [I don't see Security Hub CSPM findings in the Trusted Advisor console](#)
- [I configured Security Hub CSPM and AWS Config correctly, but my findings are still missing](#)
- [I want to disable specific Security Hub CSPM controls](#)
- [I want to find my excluded Security Hub CSPM resources](#)
- [I want to enable or disable this feature for a member account that belongs to an AWS organization](#)
- [I see multiple AWS Regions for the same affected resource for a Security Hub CSPM check](#)
- [I turned off Security Hub CSPM or AWS Config in a Region](#)
- [My control is archived in Security Hub CSPM, but I still see the findings in Trusted Advisor](#)
- [I still can't view my Security Hub CSPM findings](#)

I don't see Security Hub CSPM findings in the Trusted Advisor console

Verify that you completed the following steps:

- You have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan.
- You enabled resource recording in AWS Config within the same Region as Security Hub CSPM.
- You enabled Security Hub CSPM and selected the **AWS Foundational Security Best Practices v1.0.0** security standard.
- New controls from Security Hub CSPM are added as checks in Trusted Advisor within two to four weeks. See the [note](#).

For more information, see the [Prerequisites](#).

I configured Security Hub CSPM and AWS Config correctly, but my findings are still missing

It can take up to two hours for Security Hub CSPM to have findings for your resources. It can then take up to 24 hours for that data to appear in the Trusted Advisor console. Check the Trusted Advisor console again later.

Notes

- Only your findings for controls in the AWS Foundational Security Best Practices security standard will appear in Trusted Advisor *except* for controls that have the **Category: Recover > Resilience**.
- If there's a service issue with Security Hub CSPM or Security Hub CSPM isn't available, it can take up to 24 hours for your findings to appear in Trusted Advisor. Check the Trusted Advisor console again later.

I want to disable specific Security Hub CSPM controls

Security Hub CSPM sends your data to Trusted Advisor automatically. If you disable a Security Hub CSPM control or no longer have resources for that control, your findings won't appear in Trusted Advisor.

You can sign in to the [Security Hub CSPM console](#) and verify if your control is enabled or disabled.

If you disable a Security Hub CSPM control or disable all controls for the AWS Foundational Security Best Practices security standard, your findings are archived within the next five days. This five-day period to archive is approximate and best effort only, and isn't guaranteed. When your findings are archived, they are removed from Trusted Advisor.

For more information, see the following topics:

- [Disabling and enabling individual controls](#)
- [Disabling or enabling a security standard](#)

I want to find my excluded Security Hub CSPM resources

From the Trusted Advisor console, you can choose your Security Hub CSPM control name, and then choose the **Excluded items** option. This option displays all resources that are suppressed in Security Hub CSPM.

If the workflow status for a resource is set to SUPPRESSED, then that resource is an excluded item in Trusted Advisor. You can't suppress Security Hub CSPM resources from the Trusted Advisor console. To do so, use the [Security Hub CSPM console](#). For more information, see [Setting the workflow status for findings](#).

I want to enable or disable this feature for a member account that belongs to an AWS organization

By default, member accounts inherit the feature from the management account for AWS Organizations. If the management account has enabled the feature, then all accounts in the organization will also have the feature. If you have a member account and want to make specific changes for your account, you must contact [AWS Support](#).

I see multiple AWS Regions for the same affected resource for a Security Hub CSPM check

Some AWS services are global and aren't specific to a Region, such as IAM and Amazon CloudFront. By default, global resources such as Amazon S3 buckets appear in the US East (N. Virginia) Region.

For Security Hub CSPM checks that evaluate resources for global services, you might see more than one item for affected resources. For example, if the Hardware MFA should be enabled for the `root` user check identifies that your account hasn't activated this feature, then you will see multiple Regions in the table for the same resource.

You can configure Security Hub CSPM and AWS Config so that multiple Regions won't appear for the same resource. For more information, see [AWS Foundational Best Practices controls that you might want to disable](#).

I turned off Security Hub CSPM or AWS Config in a Region

If you stop resource recording with AWS Config or disable Security Hub CSPM in an AWS Region, Trusted Advisor no longer receives data for any controls in that Region. Trusted Advisor removes your Security Hub CSPM findings within 7-9 days. This time frame is best effort and isn't guaranteed. For more information, see [Disabling Security Hub CSPM](#).

To disable this feature for your account, see [Disable Security Hub CSPM from Trusted Advisor](#).

My control is archived in Security Hub CSPM, but I still see the findings in Trusted Advisor

When the `RecordState` status changes to `ARCHIVED` for a finding, Trusted Advisor deletes the finding for that Security Hub CSPM control from your account. You might still see the finding in Trusted Advisor for up to 7-9 days before it's deleted. This time frame is best effort and isn't guaranteed.

I still can't view my Security Hub CSPM findings

If you still have issues with this feature, you can create a technical support case in the [AWS Support Center](#).

Opt in AWS Compute Optimizer for Trusted Advisor checks

Compute Optimizer is a service that analyzes the configuration and utilization metrics of your AWS resources. This service reports whether your resources are correctly configured for efficiency and reliability. It also suggests improvements you can implement to improve workload performance. With Compute Optimizer, you view the same recommendations in your Trusted Advisor checks.

You can opt in either your AWS account only, or all member accounts that are part of an organization in AWS Organizations. For more information, see [Getting started](#) in the *AWS Compute Optimizer User Guide*.

Once you opt in for Compute Optimizer, the following checks receive data from your Lambda functions and Amazon EBS volumes. It can take up to 12 hours to generate the findings and optimization recommendations. It can then take up to 48 hours to view your results in Trusted Advisor for the following checks:

[Cost optimization](#)

- Amazon EBS over-provisioned volumes
- AWS Lambda over-provisioned functions for memory size

[Performance](#)

- Amazon EBS under-provisioned volumes
- AWS Lambda under-provisioned functions for memory size

Notes

- Results for these check are automatically refreshed at least once daily, and refresh requests are not allowed. It might take a few hours for changes to appear. You can use the Trusted Advisor console to exclude resources from checks that automatically refresh. You can use the [BatchUpdateRecommendationResourceExclusion](#) API to exclude resources from any check except Trusted Advisor Priority recommendation resources.

- Trusted Advisor already has the Underutilized Amazon EBS Volumes and the Overutilized Amazon EBS Magnetic Volumes checks.

After you opt in with Compute Optimizer, we recommend that you use the new Amazon EBS over-provisioned volumes and Amazon EBS under-provisioned volumes checks instead.

Related information

For more information, see the following topics:

- [Viewing Amazon EBS volume recommendations](#) in the *AWS Compute Optimizer User Guide*
- [Viewing Lambda function recommendations](#) in the *AWS Compute Optimizer User Guide*
- [Configuring Lambda function memory](#) in the *AWS Lambda Developer Guide*
- [Request modifications to your Amazon EBS volumes](#) in the *Amazon EC2 User Guide*

Get started with AWS Trusted Advisor Priority

Trusted Advisor Priority helps you secure and optimize your AWS account to follow AWS best practices. With Trusted Advisor Priority, your AWS account team can proactively monitor your account and create prioritized recommendations when they identify opportunities for you.

For example, your account team can identify if your AWS account root user lacks multi-factor authentication (MFA). Your account team can create a recommendation so that you can take immediate action on a check, such as MFA on Root Account. The recommendation appears as an active **prioritized recommendation** on the Trusted Advisor Priority page of the Trusted Advisor console. You then follow the recommendations to resolve it.

Trusted Advisor Priority recommendations come from these two sources:

- AWS services – Services such as Trusted Advisor, AWS Security Hub CSPM, and AWS Well-Architected automatically create recommendations. Your account team shares these recommendations with you so that those recommendations appear in Trusted Advisor Priority.
- Your account team – Your account team can create manual recommendations.

Trusted Advisor Priority helps you focus on the most important recommendations. You and your account team can monitor the recommendation lifecycle, from the point when your account team shared the recommendation, up to the point when you acknowledge, resolve, or dismiss it. You can use Trusted Advisor Priority to find recommendations for all member accounts in your organization.

Topics

- [Prerequisites](#)
- [Enable Trusted Advisor Priority](#)
- [View prioritized recommendations](#)
- [Acknowledge a recommendation](#)
- [Dismiss a recommendation](#)
- [Resolve a recommendation](#)
- [Reopen a recommendation](#)
- [Download recommendation details](#)
- [Register delegated administrators](#)
- [Deregister delegated administrators](#)
- [Manage Trusted Advisor Priority notifications](#)
- [Disable Trusted Advisor Priority](#)

Prerequisites

You must meet the following requirements to use Trusted Advisor Priority:

- You must have an AWS Enterprise Support or AWS Unified Operations plan.
- Your account must be part of an organization that has enabled all features in AWS Organizations. For more information, see [Enabling all features in your organization](#) in the *AWS Organizations User Guide*.
- Your organization must have enabled trusted access to Trusted Advisor. To enable trusted access, log in as the management account. Open the [Your organization](#) page in the Trusted Advisor console.
- You must be signed in to your AWS account to view Trusted Advisor Priority recommendations for your account.

- You must be signed in to the organization's management account or a delegated administrator account to view aggregated recommendations across your organization. For instructions on how to register delegated administrator accounts, see [Register delegated administrators](#).
- You must have AWS Identity and Access Management (IAM) permissions to access Trusted Advisor Priority. For information on how to control access to Trusted Advisor Priority, see [Manage access to AWS Trusted Advisor](#) and [AWS managed policies for AWS Trusted Advisor](#).

Enable Trusted Advisor Priority

Ask your account team to enable this feature for you. You must have an AWS Enterprise Support or AWS Unified Operations plan and be the management account owner for your organization. If the Trusted Advisor Priority page in the console says that you need trusted access with AWS Organizations, then choose **Enable trusted access with AWS Organizations**. For more information, see the [Prerequisites](#) section.

View prioritized recommendations

After your account team enables Trusted Advisor Priority for you, you can view the latest recommendations for your AWS account.

To view your prioritized recommendations

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. On the **Trusted Advisor Priority** page, you can view the following items:

If you're using an AWS Organizations Management or Delegated Administrator account, then switch to the **My Account** tab.

- **Actions needed** – The number of recommendations that are pending a response or are in progress.
- **Overview** – The following information:
 - Dismissed recommendations in the last 90 days
 - Resolved recommendations in the last 90 days
 - Recommendations without an update in over 30 days
 - Average time to resolve recommendations

3. On the **Active** tab, the **Active prioritized recommendations** show recommendations that your account team prioritized for you. The **Closed** tab shows resolved or dismissed recommendations.
 - To filter your results, use the following options:
 - **Recommendation** – Enter keywords to search by name. This can be a check name, or a custom name that your account team created.
 - **Status** – Whether the recommendation is pending a response, in progress, dismissed, or resolved.
 - **Source** – The origin of a prioritized recommendation. The recommendation can come from AWS services, your AWS account team, or a planned service event.
 - **Category** – The recommendation category, such as security or cost optimization.
 - **Age** – When your account team shared the recommendation with you.
4. Choose a recommendation to learn more about its details, the affected resources, and the recommended actions. You can then [acknowledge](#) or [dismiss](#) the recommendation.

To view prioritized recommendations across all accounts in your AWS organization

Both the management account and the Trusted Advisor Priority delegated administrators can view recommendations aggregated across your organization.

Note

Member accounts don't have access to aggregated recommendations.

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. On the **Trusted Advisor Priority** page, make sure that you're on the **My Organization** tab.
3. To view recommendations for one account, select an account from the **Select an account from your organization** dropdown list. Or, you can view recommendations across all your accounts.

On the **My Organization** tab, you can view the following items:

- **Actions needed:** The number of recommendations across your organization that are pending a response or are in progress.

- **Overview:** Shows the following items:
 - Dismissed recommendations in the last 90 days.
 - Resolved recommendations in the last 90 days.
 - Recommendations without an update in over 30 days.
 - The average time taken to resolve recommendations.
- 4. Under the **Active** tab, the **Active prioritized recommendations** section shows recommendations that your account team prioritized for you. The **Closed** tab shows resolved or dismissed recommendations.

To filter your results, use the following options:

- **Recommendation** – Enter keywords to search by name. This can be either a check name, or a custom name that your account team created.
 - **Status** – Whether the recommendation is pending a response, in progress, dismissed, or resolved.
 - **Source** – The origin of a prioritized recommendation. The recommendation can come from AWS services, your AWS account team, or a planned service event.
 - **Category** – The recommendation category, such as security or cost optimization.
 - **Age** – When your account team shared the recommendation with you.
5. Choose a recommendation to see additional details, affected accounts and resources, and the recommended actions. You can then [acknowledge](#) or [dismiss](#) the recommendation.

Example: Trusted Advisor Priority recommendations

The following example shows 15 recommendations that are pending a response and 27 recommendations that are in progress under the **Action needed** section. The following image shows two of the recommendations that are pending response in the **Active prioritized recommendation** tab.

Trusted Advisor > Priority

Trusted Advisor Priority [info](#)

You can use this page to find critical recommendations, trends, and activities for your organization.

My organization My account

Select an account from your organization

All accounts

Action needed

⏸ Pending response
15

🔄 In progress
27

Overview

Dismissed in the last 90 days
🗑 5

Resolved in the last 90 days
🟢 22

No update in 30+ days
10

Average time to resolve
46 days

Active Closed

Active prioritized recommendations (42)
Your AWS account team has prioritized the following recommendations for your organization. Choose a recommendation to learn more.

Search

Recommendations	Status	Source	Category	Age (days)
☐ Low Utilization Amazon EC2 Instances test test	⏸ Pending response	AWS Trusted Advisor	Cost optimization	33 day(s) Shared on: Jun 20, 2023
☐ RDS DB instances should have deletion protection enabled	⏸ Pending response	AWS Security Hub	Security	20 day(s) Shared on: Jul 3, 2023

Acknowledge a recommendation

Under the **Active** tab, you can learn more about the recommendation and then decide if you want to acknowledge it.

To acknowledge a recommendation

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. If you're using an AWS Organizations Management or Delegated Administrator account, then switch to the **My Account** tab.
3. On the **Trusted Advisor Priority** page, under the **Active** tab, choose a recommendation name.
4. In the **Details** section, you can review the recommended actions to resolve the recommendation.
5. In the **Affected resources** section, you can review the affected resources and filter by *Status*.
6. Choose **Acknowledge**.
7. In the **Acknowledge recommendation** dialog box, choose **Acknowledge**.

The recommendation status changes to **In progress**. Recommendations in progress or pending a response appear in the **Active** tab on the Trusted Advisor Priority page.

8. Follow the recommended actions to resolve the recommendation. For more information, see [Resolve a recommendation](#).

Example: Manual recommendation from Trusted Advisor Priority

The following image shows the **Low Utilization EC2 Instances** recommendation that is pending a response.

Trusted Advisor > Priority > Low Utilization Amazon EC2 Instances - Production accounts

My organization My account

Low Utilization Amazon EC2 Instances - Production accounts

Copy recommendation link Download Acknowledge Dismiss

Details Affected resources

Overview

Source AWS Trusted Advisor	Category Cost optimization	Age 33 day(s) Shared on: Jun 20, 2023	Status Pending response
-------------------------------	-------------------------------	---	----------------------------

Shared by
person@amazon.com

Details

Description
Checks the Amazon Elastic Compute Cloud (Amazon EC2) instances that were running at any time during the last 14 days and alerts you if the daily CPU utilization was 10% or less and network I/O was 5 MB or less on 4 or more days. Running instances generate hourly usage charges. Although some scenarios can result in low utilization by design, you can often lower your costs by managing the number and size of your instances.

Estimated monthly savings are calculated by using the current usage rate for On-Demand Instances and the estimated number of days the instance might be underutilized. Actual savings will vary if you are using Reserved Instances or Spot Instances, or if the instance is not running for a full day. To get daily utilization data, download the report for this check.

Alert Criteria
Yellow: An instance had 10% or less daily average CPU utilization and 5 MB or less network I/O on at least 4 of the previous 14 days.

Recommended Action
Consider stopping or terminating instances that have low utilization, or scale the number of instances by using Auto Scaling. For more information, see [Stop and Start Your Instance](#), [Terminate Your Instance](#), and [What is Auto Scaling?](#)

Additional Resources
[Monitoring Amazon EC2 Instance Metadata and User Data](#)
[Amazon CloudWatch Developer Guide](#)
[Auto Scaling Developer Guide](#)

To acknowledge a recommendation for all accounts in your AWS organization

The management account or the Trusted Advisor delegated administrators can acknowledge a recommendation for all of the affected accounts.

Note

Member accounts don't have access to aggregated recommendations.

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. On the **Trusted Advisor Priority** page, make sure that you're on the **My organization** tab.
3. In the **Active** tab, select a recommendation name.
4. Choose **Acknowledge**.
5. In the **Acknowledge recommendation** dialog box, choose **Acknowledge**.

The recommendation status changes to **In progress**.

6. Follow the recommended actions to resolve the recommendation. For more information, see [Resolve a recommendation](#).
7. To view the recommendation details, choose the recommendation name.

In the **Details** section, you can review the following information about the recommendation:

- An **Overview** of the recommendation and a **Details** section covering the recommendation actions to complete.

A **Status summary** that shows recommendations across all affected accounts.

- In the **Affected accounts** section, you can review the affected resources across all your accounts. You can filter by **Account number** and **Status**.
- In the **Affected resources** section, you can review the affected resources across all your accounts. You can filter by **Account number** and **Status**.

Example: Manual recommendation from Trusted Advisor Priority

The following image shows the **Low Utilization Amazon EC2 Instances** recommendation that's pending a response. One affected account has acknowledged the recommendation. Another account is pending a response, making the recommendation status **Pending response**.

Trusted Advisor > Priority > Low Utilization Amazon EC2 Instances - Production accounts

My organization | My account

Low Utilization Amazon EC2 Instances - Production accounts

Copy recommendation link | Download | Acknowledge | Dismiss

Details | Affected accounts | Affected resources

Overview

Source AWS Trusted Advisor	Category Cost optimization	Age 0 day(s) Shared on: Jul 10, 2023	Status ⚠ Pending response
Shared by person@amazon.com			

Status Summary
This is a summary of the status of this recommendation across all your accounts

- ⚠ 1 account Pending response
- ⌚ 1 account In progress

Details

Description
Checks the Amazon Elastic Compute Cloud (Amazon EC2) instances that were running at any time during the last 14 days and alerts you if the daily CPU utilization was 10% or less and network I/O was 5 MB or less on 4 or more days. Running instances generate hourly usage charges. Although some scenarios can result in low utilization by design, you can often lower your costs by managing the number and size of your instances.

Estimated monthly savings are calculated by using the current usage rate for On-Demand Instances and the estimated number of days the instance might be underutilized. Actual savings will vary if you are using Reserved Instances or Spot Instances, or if the instance is not running for a full day. To get daily utilization data, download the report for this check.

Alert Criteria

Yellow: An instance had 10% or less daily average CPU utilization and 5 MB or less network I/O on at least 4 of the previous 14 days.

Recommended Action

Consider stopping or terminating instances that have low utilization, or scale the number of instances by using Auto Scaling. For more information, see [Stop and Start Your Instance](#), [Terminate Your Instance](#), and [What is Auto Scaling?](#)

Dismiss a recommendation

You can also dismiss a recommendation. This means that you acknowledge the recommendation, but you won't address it. You can dismiss a recommendation if it's not relevant to your account. For example, if you have a test AWS account that you plan to delete, you don't need to follow the recommended actions.

To dismiss a recommendation

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. If you're using an AWS Organizations Management or Delegated Administrator account, then switch to the **My Account** tab.
3. On the **Trusted Advisor Priority** page, under the **Active** tab, choose a recommendation name.
4. On the recommendation detail page, review the information about the affected resources.
5. If this recommendation doesn't apply for your account, choose **Dismiss**.
6. In the **Dismiss recommendation** dialog box, select a reason why you won't address the recommendation.
7. (Optional) Enter a note detailing why you're dismissing the recommendation. If you choose **Other**, you must enter a description in the **Note** section.
8. Choose **Dismiss**. The recommendation status changes to **Dismissed** and appears in the **Closed** tab on the Trusted Advisor Priority page.

To dismiss a recommendation for all the accounts in your AWS organization

The management account or the delegated administrator of Trusted Advisor Priority can dismiss a recommendation for all of their accounts.

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. On the Trusted Advisor Priority page, make sure that you're on the **My Organization** tab.
3. In the **Active** tab, select a recommendation name.
4. If this recommendation doesn't apply for your account, then choose **Dismiss**.
5. In the **Dismiss recommendation** dialog box, select a reason why you won't address the recommendation.

6. (Optional) Enter a note detailing why you're dismissing the recommendation. If you choose **Other**, then you must enter a description in the **Note** section.
7. Choose **Dismiss**. The recommendation status changes to **Dismissed**. The recommendation appears in the **Closed** tab on the Trusted Advisor Priority page.

Note

You can choose the recommendation name and choose **View note** to find the reason for dismissal. If your account team dismissed the recommendation for you, their email address appears next to the note.

Trusted Advisor Priority also notifies your account team that you dismissed the recommendation.

Example: Dismiss a recommendation from Trusted Advisor Priority

The following example shows how you can dismiss a recommendation.

Dismiss recommendation

 Please note: This action will apply to all accounts affected by this recommendation

Choose a reason for why you're dismissing this recommendation

The affected AWS account was temporarily created for an event 

Note - *optional*

These are test accounts that we will delete soon 

Cancel **Dismiss**

Resolve a recommendation

After you acknowledge the recommendation and complete the recommended actions, you can resolve the recommendation.

Tip

After you resolve a recommendation, you can't reopen it. If you want to revisit the recommendation again later, see [Dismiss a recommendation](#).

To resolve a recommendation

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. On the Trusted Advisor Priority page, make sure that you're on the **My Organization** tab.
3. On the **Trusted Advisor Priority** page, select the recommendation, and then choose **Resolve**.
4. In the **Resolve recommendation** dialog box, choose **Resolve**. Resolved recommendations appear under the **Closed** tab on the Trusted Advisor Priority page. Trusted Advisor Priority notifies your account team that you resolved the recommendation.

To resolve a recommendation for all accounts in your AWS organization

The management account or the Trusted Advisor Priority delegated administrators can resolve a recommendation for all their accounts.

Note

Member accounts don't have access to aggregated recommendations.

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. If you're using an AWS Organizations Management or Delegated Administrator account, switch to the **My Account** tab.
3. In the **Active** tab, select a recommendation name.

- If the recommendation doesn't apply for your account, choose **Resolve**.
- In the **Resolve recommendation** dialog box, choose **Resolve**. Resolved recommendations appear under the **Closed** tab on the Trusted Advisor Priority page. Trusted Advisor Priority notifies your account team that you resolved the recommendation.

Example: Manual recommendation from Trusted Advisor Priority

The following example shows a resolved **Low Utilization Amazon EC2 Instances** recommendation.

The screenshot shows the AWS Trusted Advisor console interface. At the top, there's a breadcrumb trail: "Trusted Advisor > Priority > Low Utilization Amazon EC2 Instances - Production accounts". Below this, there are tabs for "My organization" and "My account", with "My account" selected. The main heading is "Low Utilization Amazon EC2 Instances - Production accounts". To the right of the heading are buttons for "Copy recommendation link" and "Download". Below the heading are tabs for "Details", "Affected accounts", and "Affected resources", with "Details" selected. The "Details" tab shows an "Overview" section with a table containing the following information:

Source	Category	Age	Status
AWS Trusted Advisor	Cost optimization	0 day(s) Shared on: Jul 10, 2023	Resolved

Below the table, it says "Shared by person@amazon.com" and "Resolved on Jul 10, 2023". To the right of the table is a "Status Summary" section with the text "This is a summary of the status of this recommendation across all your accounts" and a green checkmark icon followed by "2 accounts Resolved".

Reopen a recommendation

After you dismiss a recommendation, you or your account team can reopen the recommendation.

To reopen a recommendation

- Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
- If you're using an AWS Organizations Management or Delegated Administrator account, then switch to the **My Account** tab.
- On the **Trusted Advisor Priority** page, choose the **Closed** tab.
- Under **Closed recommendations**, select a recommendation that was **Dismissed**, and then choose **Reopen**.
- In the **Reopen recommendation** dialog box, describe why you're reopening the recommendation.
- Choose **Reopen**. The recommendation status changes to **In progress** and appears under the **Active** tab.

Tip

You can choose the recommendation name and then choose **View note** to find the reason for reopening. If your account team reopened the recommendation for you, their name appears next to the note.

7. Follow the steps in the recommendation details.

To reopen a recommendation for all accounts in your AWS organization

The management account or the Trusted Advisor Priority delegated administrators can reopen a recommendation for all of their accounts.

Note

Member accounts don't have access to aggregated recommendations.

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. On the Trusted Advisor Priority page, make sure that you're on the **My Organization** tab.
3. Under **Closed** recommendations, select a recommendation that was **Dismissed**, and then choose **Reopen**.
4. In the **Reopen recommendation** dialog box, describe why you're reopening the recommendation.
5. Choose **Reopen**. The recommendation status changes to **In progress** and appears under the **Active** tab.


Tip

You can choose the recommendation name and choose **View note** to find the reason for reopening. If your account team reopened the recommendation for you, their name appears next to the note.

6. Follow the steps in the recommendation details.

Example: Reopen a recommendation from Trusted Advisor Priority

The following example shows a recommendation that you want to reopen.



The screenshot shows a modal dialog titled "Reopen recommendation" with a close button (X) in the top right corner. Inside the dialog, there is a warning icon (triangle with an exclamation mark) followed by the text: "Please note: This action will apply to all accounts affected by this recommendation". Below this, there is a label "Reason for reopening" and a text input field containing the text "I need to address this recommendation for my organization". At the bottom right of the dialog, there are two buttons: "Cancel" and "Reopen" (which is highlighted in orange).

Download recommendation details

You can also download the results of a prioritized recommendation from Trusted Advisor Priority.

Note

Currently, you can download only one recommendation at a time.

To download a recommendation

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home>.
2. On the **Trusted Advisor Priority** page, select the recommendation, and then choose **Download**.
3. Open the file to view the recommendation details.

Register delegated administrators

You can add member accounts that are part of your organization as delegated administrators. Delegated administrator accounts can review, acknowledge, resolve, dismiss, and reopen recommendations in Trusted Advisor Priority.

After you register an account, you must grant the delegated administrator the required AWS Identity and Access Management permissions to access Trusted Advisor Priority. For more information, see [Manage access to AWS Trusted Advisor](#) and [AWS managed policies for AWS Trusted Advisor](#).

You can register up to five member accounts. Only the management account can add delegated administrators for the organization. You must be signed in to the organization's management account to register or deregister a delegated administrator.

To register a delegated administrator

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home> as the management account.
2. In the navigation pane, under **Preferences**, choose **Your organization**.
3. Under **Delegated administrator**, choose **Register new account**.
4. In the dialog box, enter the member account ID, and then choose **Register**.
5. (Optional) To deregister an account, select an account and choose **Deregister**. In the dialog box, choose **Deregister** again.

Deregister delegated administrators

When you deregister a member account, that account no longer has the same access to Trusted Advisor Priority as the management account. Accounts that are no longer delegated administrators won't receive email notifications from Trusted Advisor Priority.

To deregister a delegated administrator

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home> as the management account.
2. In the navigation pane, under **Preferences**, choose **Your organization**.
3. Under **Delegated administrator**, select an account and then choose **Deregister**.

4. In the dialog box, choose **Deregister**.

Manage Trusted Advisor Priority notifications

Trusted Advisor Priority delivers notifications through email. This email notification includes a summary of the recommendations that your account team prioritized for you. You can specify the frequency that you receive updates from Trusted Advisor Priority.

If you registered member accounts as delegated administrators, they can also set up their accounts to receive Trusted Advisor Priority email notifications.

Trusted Advisor Priority email notifications don't include check results for individual accounts and are separate from the weekly notification for Trusted Advisor Recommendations. For more information, see [Set up notification preferences](#).

Note

Only the management account or delegated administrator can set up Trusted Advisor Priority email notifications.

To manage your Trusted Advisor Priority notifications

1. Sign in to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor/home> as a management or delegated administrator account.
2. In the navigation pane, under **Preferences**, choose **Notifications**.
3. Under **Priority**, you can select the following options.
 - a. **Daily** – Receive an email notification daily.
 - b. **Weekly** – Receive an email notification once a week.
 - c. Choose the notifications to receive:
 - Summary of prioritized recommendations
 - Resolution dates
4. For **Recipients**, select other contacts that you want to receive the email notifications. You can add and remove contacts from the [Account Settings](#) page in the AWS Billing and Cost Management console.

5. For **Language**, choose the language for the email notification.
6. Choose **Save your preferences**.

Note

Trusted Advisor Priority sends email notifications from the `noreply@notifications.trustedadvisor.us-west-2.amazonaws.com` address. You might need to verify that your email client doesn't identify these emails as spam.

Disable Trusted Advisor Priority

Contact your account team and ask that they disable this feature for you. After this feature is disabled, prioritized recommendations no longer appear in your Trusted Advisor console.

If you disable Trusted Advisor Priority and then enable it again later, you can still view the recommendations that your account team sent before you disabled Trusted Advisor Priority.

AWS Trusted Advisor check reference

Important

End of Support Notice: Developer Support will be discontinued January 1, 2027. Customers with Developer Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information, see [Business Support+ plan details](#)

End of Support Notice: Business Support will be discontinued January 1, 2027. Customers with Business Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information see, [Business Support+ plan details](#)

End of Support Notice: On January 1, 2027, AWS will discontinue Enterprise On-Ramp. Throughout 2026, Enterprise On-Ramp customers will be automatically upgraded to AWS Enterprise Support during contract renewal or in periodic batches. Customers will

receive an email notification a month before their upgrade. No further action is required. Enterprise Support provides designated TAM assignment, 15-minute response times, and AWS Security Incident Response available at no additional cost, all at a lower \$5,000 minimum (reduced from \$15,000). For more information, see [AWS Enterprise Support plan details](#).

For more information, see [Developer, Business, and Enterprise On-Ramp end of support](#). Developer Support, Business Support, and Enterprise On-Ramp will remain available in the AWS GovCloud (US) Region.

You can view all Trusted Advisor check names, descriptions, and IDs in the following reference. You can also sign in to the [Trusted Advisor](#) console to view more information about the checks, recommended actions, and their statuses.

If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, you can also use the [AWS Trusted Advisor API](#) and the AWS Command Line Interface (AWS CLI) to access your checks. For more information, see the following topics:

- [Get started with the Trusted Advisor API](#)
- [AWS Trusted Advisor API Reference](#)

Note

If you have a Basic Support or Developer Support plan, you can use the Trusted Advisor console to access all checks in the [Service limits](#) category and the following checks in the Security and Fault Tolerance categories:

- [Amazon EBS Public Snapshots](#)
- [Amazon RDS Public Snapshots](#)
- [Amazon S3 Bucket Permissions](#)
- [MFA on root account](#)
- [Security Groups – Specific Ports Unrestricted](#)
- [AWS STS global endpoint usage across AWS Regions](#)

Check categories

- [Cost optimization](#)
- [Performance](#)
- [Security](#)
- [Fault tolerance](#)
- [Service limits](#)
- [Operational Excellence](#)

Cost optimization

You can use the following checks for the cost optimization category.

Check names

- [AWS account not part of AWS Organizations](#)
- [Amazon Aurora cost optimization recommendations for DB cluster storage](#)
- [Amazon Comprehend underutilized endpoints](#)
- [Amazon DynamoDB reserved capacity purchase recommendations](#)
- [Amazon EBS cost optimization recommendations for volumes](#)
- [Amazon EBS over-provisioned volumes](#)
- [Amazon EC2 cost optimization recommendations for Amazon EC2 Auto Scaling groups](#)
- [Amazon EC2 cost optimization recommendations for instances](#)
- [Amazon EC2 instances consolidation for Microsoft SQL Server](#)
- [Amazon EC2 instances over-provisioned for Microsoft SQL Server](#)
- [Amazon EC2 instances stopped](#)
- [Amazon EC2 Reserved Instance lease expiration](#)
- [Amazon EC2 Reserved Instance optimization](#)
- [Amazon ECR Repository without lifecycle policy configured](#)
- [Amazon ElastiCache reserved node purchase recommendations](#)
- [AWS Fargate cost optimization recommendations for Amazon ECS](#)
- [Amazon MemoryDB reserved node purchase recommendations](#)
- [Amazon OpenSearch Service Reserved Instance purchase recommendations](#)
- [Amazon RDS cost optimization recommendations for DB instances](#)

- [Amazon RDS cost optimization recommendations for DB instance storage](#)
- [Amazon RDS idle DB instances](#)
- [Amazon RDS Reserved Instance purchase recommendations](#)
- [Amazon Redshift reserved node purchase recommendations](#)
- [Amazon Route 53 Latency Resource Record Sets](#)
- [Amazon S3 Bucket Lifecycle Policy Configured](#)
- [Amazon S3 Incomplete Multipart Upload Abort Configuration](#)
- [Amazon S3 version-enabled buckets without lifecycle policies configured](#)
- [AWS Lambda cost optimization recommendations for functions](#)
- [AWS Lambda functions with excessive timeouts](#)
- [AWS Lambda functions with high error rates](#)
- [AWS Lambda over-provisioned functions for memory size](#)
- [AWS Savings Plans purchase recommendations for compute](#)
- [AWS Savings Plans purchase recommendations for Amazon SageMaker AI](#)
- [AWS Well-Architected high risk issues for cost optimization](#)
- [Idle Load Balancers](#)
- [Idle NAT gateways](#)
- [Inactive AWS Network Firewall](#)
- [Inactive VPC interface endpoints](#)
- [Inactive Gateway Load Balancer endpoints](#)
- [Inactive NAT Gateways](#)
- [Low utilization Amazon EC2 instances](#)
- [Unassociated Elastic IP Addresses](#)
- [Underutilized Amazon EBS volumes](#)
- [Underutilized Amazon Redshift Clusters](#)

AWS account not part of AWS Organizations

Description

Checks if an AWS account is part of AWS Organizations under the appropriate management account.

AWS Organizations is an account management service for consolidating multiple AWS accounts into a centrally-managed organization. This enables you to centrally structure accounts for billing consolidation and implement ownership and security policies as your workloads scale on AWS.

You can specify the management account id using the **MasterAccountId** parameter of the AWS Config rules.

For more information, see [What is AWS Organizations?](#)

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz127

Source

AWS Config Managed Rule: account-part-of-organizations

Alert criteria

Yellow: This AWS account is not part of AWS Organizations.

Recommended action

Add this AWS account as part of AWS Organizations.

For more information, see [Tutorial: Creating and configuring an organization.](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters

- Last Updated Time

Amazon Aurora cost optimization recommendations for DB cluster storage

Description

Checks your Amazon Aurora DB cluster storage configurations and usage patterns to provide recommendations for potential cost savings.

For each resource, Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#) and [AWS Compute Optimizer](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr17n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: An Aurora DB cluster storage has a cost savings action identified by Cost Optimization Hub.

Recommended action

Consider implementing the recommendation. The recommendation is one of the recommendation types listed in [Understanding cost optimization strategies](#) in the *AWS Cost Management User Guide*. For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Use the detailed recommendations in AWS Compute Optimizer to understand the potential impact of these changes on cost and performance.

Additional resources

- [Viewing Aurora and RDS database recommendations](#)
- [Recommendations from Amazon Aurora](#)
- [Settings for Amazon Aurora](#)

Report columns

- Status
- Region
- Resource ID
- Recommended Action
- Current Resource Summary
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

Amazon Comprehend underutilized endpoints

Description

Checks the throughput configuration of your endpoints. This check alerts you when endpoints are not actively used for real-time inference requests. An endpoint that isn't used for more than 15 consecutive days is considered underutilized. All endpoints accrue charges based on both the throughput set, and the length of time that the endpoint is active.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

Cm24dfsM12

Alert criteria

Yellow: The endpoint is active, but hasn't been used for real-time inference requests in the past 15 days.

Recommended action

If the endpoint hasn't been used in the past 15 days, we recommend that you define a scaling policy for the resource by using [Application Autoscaling](#).

If the endpoint has a scaling policy defined and hasn't been used in the past 30 days, consider deleting the endpoint and using asynchronous inference. For more information, see [Deleting an endpoint with Amazon Comprehend](#).

Report columns

- Status
- Region
- Endpoint ARN
- Provisioned Inference Unit
- AutoScaling Status
- Reason
- Last Updated Time

Amazon DynamoDB reserved capacity purchase recommendations

Description

Checks your Amazon DynamoDB usage patterns and provides recommendations for potential cost savings through reserved capacity purchases.

This check generates recommendations at [payer scope for the payer account and at linked scope for a linked account](#).

Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr15n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: Account has a cost savings action identified by Cost Optimization Hub for DynamoDB.

Recommended action

Consider [implementing the recommendation](#). For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Additional resources

- [DynamoDB reserved capacity](#)
- [Accessing reservation recommendations](#)
- [Amazon DynamoDB Reserved Capacity](#)

Report columns

- Status
- Region
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

Amazon EBS cost optimization recommendations for volumes

Description

Checks your Amazon EBS volume configurations and usage patterns to provide recommendations for potential cost savings.

For each resource, Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#) and [AWS Compute Optimizer](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr02n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: EBS volume has a cost savings action identified by Cost Optimization Hub.

Recommended action

Consider implementing the recommendation. The recommendation is one of the recommendation types listed in [Understanding cost optimization strategies](#) in the *AWS Cost Management User Guide*. For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Use the detailed recommendations in AWS Compute Optimizer to understand the potential impact of these changes on cost and performance.

Additional resources

- [Viewing Amazon EBS volume recommendations](#)
- [EBS volume metrics](#)
- [Amazon EBS volume types](#)
- [Request Amazon EBS volume modifications](#)

Report columns

- Status
- Region
- Resource ID
- Recommended Action
- Current Resource Summary
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

Amazon EBS over-provisioned volumes

Description

Note

This is a legacy check. We recommend using the new check (Check ID: [c1z7kmr02n](#)) that offers additional customized recommendations.

Checks the Amazon Elastic Block Store (Amazon EBS) volumes that were running at any time during the lookback period. This check alerts you if any EBS volumes were over-provisioned for your workloads. When you have over-provisioned volumes, you're paying for unused resources. Although some scenarios can result in low optimization by design, you can often lower your costs by changing the configuration of your EBS volumes. Estimated monthly savings are calculated by using the current usage rate for EBS volumes. Actual savings vary if the volume isn't present for a full month.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

C016dfpM03

Alert Criteria

Yellow: An EBS volume that was over-provisioned during the lookback period. To determine if a volume is over-provisioned, we consider all default CloudWatch metrics (including IOPS and throughput). The algorithm used to identify over-provisioned EBS volumes follows AWS best practices. The algorithm is updated when a new pattern has been identified.

Recommended action

Consider downsizing volumes that have low utilization.

For more information, see [Opt in AWS Compute Optimizer for Trusted Advisor checks](#).

Report columns

- Status
- Region
- Volume ID
- Volume Type
- Volume Size (GB)
- Volume Baseline IOPS
- Volume Burst IOPS
- Volume Burst Throughput
- Recommended Volume Type
- Recommended Volume Size (GB)
- Recommended Volume Baseline IOPS

- Recommended Volume Burst IOPS
- Recommended Volume Baseline Throughput
- Recommended Volume Burst Throughput
- Lookback Period (days)
- Savings Opportunity (%)
- Estimated Monthly Savings
- Estimated Monthly Savings Currency
- Last Updated Time

Amazon EC2 cost optimization recommendations for Amazon EC2 Auto Scaling groups

Description

Checks your Amazon EC2 Auto Scaling group configurations and usage patterns to provide recommendations for potential cost savings.

For each resource, Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#) and [AWS Compute Optimizer](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr01n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: Amazon EC2 Auto Scaling group has a cost savings action identified by Cost Optimization Hub.

Recommended action

Consider implementing the recommendation. The recommendation is one of the recommendation types listed in [Understanding cost optimization strategies](#) in the *AWS Cost Management User Guide*. For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Use the detailed recommendations in AWS Compute Optimizer to understand the potential impact of these changes on cost and performance.

Additional resources

- [Viewing Amazon EC2 Auto Scaling volume recommendations](#)
- [Amazon EC2 Auto Scaling groups](#)
- [What is Amazon EC2 Auto Scaling](#)

Report columns

- Status
- Region
- Resource ID
- Recommended Action
- Current Resource Summary
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

Amazon EC2 cost optimization recommendations for instances

Description

Checks your Amazon EC2 instance configurations and usage patterns to provide recommendations for potential cost savings.

For each resource, Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#) and [AWS Compute Optimizer](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmi00n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: EC2 instance has a cost savings action identified by Cost Optimization Hub.

Recommended action

Consider implementing the recommendation. The recommendation is one of the recommendation types listed in [Understanding cost optimization strategies](#) in the *AWS Cost Management User Guide*. For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Use the detailed recommendations in AWS Compute Optimizer to understand the potential impact of these changes on cost and performance.

Additional resources

- [Viewing EC2 instance recommendations](#)
- [EC2 instance metrics](#)
- [Amazon EC2 instance type changes](#)

Report columns

- Status
- Region
- Resource ID
- Recommended Action
- Current Resource Summary
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

Amazon EC2 instances consolidation for Microsoft SQL Server

Description

Checks your Amazon Elastic Compute Cloud (Amazon EC2) instances that are running SQL Server in the past 24 hours. This check alerts you if your instance has less than the minimum number of SQL Server licenses. From the Microsoft SQL Server Licensing Guide, you are paying 4 vCPU licenses even if an instance has only 1 or 2 vCPUs. You can consolidate smaller SQL Server instances to help lower costs.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

Qsdfp3A4L2

Alert criteria

Yellow: An instance with SQL Server has less than 4 vCPUs.

Recommended Action

Consider consolidating smaller SQL Server workloads into instances with at least four vCPUs.

Additional resources

- [Microsoft SQL Server on AWS](#)
- [Microsoft Licensing on AWS](#)
- [Microsoft SQL Server Licensing Guide](#)

Report columns

- Status
- Region
- Instance ID
- Instance Type
- vCPU
- Minimum vCPU
- SQL Server Edition
- Last Updated Time

Amazon EC2 instances over-provisioned for Microsoft SQL Server

Description

Checks your Amazon Elastic Compute Cloud (Amazon EC2) instances that are running SQL Server in the past 24 hours. An SQL Server database has a compute capacity limit for each instance. An instance with SQL Server Standard edition can use up to 48 vCPUs. An instance with SQL Server Web can use up to 32 vCPUs. This check alerts you if an instance exceeds this vCPU limit.

If your instance is over-provisioned, you pay full price without realizing an improvement in performance. You can manage the number and size of your instances to help lower costs.

Estimated monthly savings are calculated by using the same instance family with the maximum number of vCPUs that an SQL Server instance can use and the On-Demand pricing. Actual savings will vary if you're using Reserved Instances (RI) or if the instance isn't running for a full day.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

Qsdfp3A4L1

Alert criteria

- Red: An instance with SQL Server Standard edition has more than 48 vCPUs.
- Red: An instance with SQL Server Web edition has more than 32 vCPUs.

Recommended action

For SQL Server Standard edition, consider changing to an instance in the same instance family with 48 vCPUs. For SQL Server Web edition, consider changing to an instance in the same instance family with 32 vCPUs. If it is memory intensive, consider changing to memory optimized R5 instances.

Additional resources

- [Microsoft SQL Server on AWS](#)
- You can use [Launch Wizard](#) to simplify your SQL Server deployment on EC2.

Report columns

- Status
- Region
- Instance ID
- Instance Type
- vCPU
- SQL Server Edition
- Maximum vCPU
- Recommended Instance Type

- Estimated Monthly Savings
- Last Updated Time

Amazon EC2 instances stopped

Description

Checks if there are Amazon EC2 instances that have been stopped for more than 30 days.

You can specify the allowed number of days value in the **AllowedDays** of AWS Config parameters.

For more information, see [Why am I being charged for Amazon EC2 when all my instances were terminated?](#)

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz150

Source

AWS Config Managed Rule: ec2-stopped-instance

Alert criteria

- Yellow: There are Amazon EC2 instances stopped for more than the allowed number of days.

Recommended action

Review the Amazon EC2 instances that have been stopped for 30 days or more. To avoid incurring unnecessary costs, terminate any instances that are no longer needed.

For more information, see [Terminate your instance](#).

Additional resources

- [Amazon EC2 On-Demand Pricing](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon EC2 Reserved Instance lease expiration

Description

Checks for Amazon EC2 Reserved Instances that are scheduled to expire within the next 30 days, or have expired in the preceding 30 days.

Reserved Instances don't renew automatically. You can continue using an Amazon EC2 instance covered by the reservation without interruption, but you will be charged On-Demand rates. New Reserved Instances can have the same parameters as the expired ones, or you can purchase Reserved Instances with different parameters.

The estimated monthly savings is the difference between the On-Demand and Reserved Instance rates for the same instance type.

Check ID

1e93e4c0b5

Alert criteria

- Yellow: The Reserved Instance lease expires in less than 30 days.
- Yellow: The Reserved Instance lease expired in the preceding 30 days.

Recommended action

Consider purchasing a new Reserved Instance to replace the one that is nearing the end of its term. For more information, see [How to Purchase Reserved Instances](#) and [Buying Reserved Instances](#).

Additional resources

- [Reserved Instances](#)
- [Instance Types](#)

Report columns

- Status
- Zone
- Instance Type
- Platform
- Instance Count
- Current Monthly Cost
- Estimated Monthly Savings
- Expiration Date
- Reserved Instance ID
- Reason

Amazon EC2 Reserved Instance optimization

Description

An important part of using AWS involves balancing your Reserved Instance (RI) purchase against your On-Demand Instance usage. This check provides recommendations on which RIs will help reduce the costs incurred from using On-Demand Instances.

We create these recommendations by analyzing your On-Demand usage for the past 30 days. We then categorizing the usage into eligible categories for reservations. We simulate every combination of reservations in the generated category of usage to identify the recommended number of each type of RI to purchase. This process of simulation and optimization allows us to maximize your cost savings. This check covers recommendations based on Standard Reserved Instances with the partial upfront payment option.

This check is not available to accounts linked in consolidated billing. The recommendations for this check are only available for the paying account.

Check ID

cX3c2R1chu

Alert criteria

Yellow: Optimizing the use of partial upfront RIs can help reduce costs.

Recommended action

See the [Cost Explorer](#) page for more detailed and customized recommendations. Additionally, refer to the [buying guide](#) to understand how to purchase RIs and the options available.

Additional resources

- Information on RIs and how they can save you money can be found [here](#).
- For more information on this recommendation, see [Reserved Instance Optimization Check Questions](#) in the Trusted Advisor FAQs.

Report columns

- Region
- Instance Type
- Platform
- Recommended Number of RIs to Purchase
- Expected Average RI Utilization
- Estimated Savings with Recommendations (Monthly)
- Upfront Cost of RIs
- Estimated costs of RIs (Monthly)
- Estimated On-Demand Cost Post Recommended RI Purchase (Monthly)
- Estimated Break Even (Months)
- Lookback Period (Days)
- Term (Years)

Amazon ECR Repository without lifecycle policy configured

Description

Checks if a private Amazon ECR repository has at least one lifecycle policy configured. Lifecycle policies allow you to define a set of rules to automatically clean up old or unused container images. This gives you control over the lifecycle management of the images, allows Amazon ECR repositories to be better organized, and helps to lower overall storage costs.

For more information, see [Lifecycle policies](#).

Check ID

c18d2gz128

Source

AWS Config Managed Rule: ecr-private-lifecycle-policy-configured

Alert criteria

Yellow: An Amazon ECR private repository doesn't have any lifecycle policies configured.

Recommended action

Consider creating at least one lifecycle policy for your private Amazon ECR repository.

For more information, see [Creating a lifecycle policy](#).

Additional resources

- [Lifecycle policies](#).
- [Creating a lifecycle policy](#).
- [Examples of lifecycle policies](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon ElastiCache reserved node purchase recommendations

Description

Checks your Amazon ElastiCache usage patterns to provide recommendations for potential cost savings through reserved node purchases.

This check generates recommendations at [payer scope for the payer account and at linked scope for a linked account](#).

Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#).

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr13n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: Account has a cost savings action identified by Cost Optimization Hub for Amazon ElastiCache.

Recommended action

Consider [implementing the recommendation](#) For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Additional resources

- [Reserved nodes](#)
- [Accessing reservation recommendations](#)
- [Amazon ElastiCache Reserved Nodes](#)

Report columns

- Status
- Region
- Recommended Resource Summary

- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

AWS Fargate cost optimization recommendations for Amazon ECS

Description

Checks your AWS Fargate for Amazon ECS configurations and usage patterns to provide recommendations for potential cost savings.

For each resource, Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#) and [AWS Compute Optimizer](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr06n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: Amazon ECS services on AWS Fargate have a cost savings action identified by Cost Optimization Hub.

Recommended action

Consider implementing the recommendation. The recommendation is one of the recommendation types listed in [Understanding cost optimization strategies](#) in the *AWS Cost*

Management User Guide. For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Use the detailed recommendations in AWS Compute Optimizer to understand the potential impact of these changes on cost and performance.

Additional resources

- [Viewing Amazon ECS services on Fargate recommendations](#)
- [Metrics for Amazon ECS services on Fargate](#)
- [AWS Fargate for Amazon ECS](#)
- [Updating an Amazon ECS service using the console](#)

Report columns

- Status
- Region
- Resource ID
- Recommended Action
- Current Resource Summary
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

Amazon MemoryDB reserved node purchase recommendations

Description

Checks your Amazon MemoryDB usage patterns to provide recommendations for potential cost savings through reserved node purchases.

This check generates recommendations at [payer scope for the payer account and at linked scope for a linked account](#).

For each resource, Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr16n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: Account has a cost savings action identified by Cost Optimization Hub for MemoryDB.

Recommended action

Consider [implementing the recommendation](#). For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Additional resources

- [MemoryDB reserved nodes](#)
- [Accessing reservation recommendations](#)
- [Working with reserved nodes](#)

Report columns

- Status
- Region
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

Amazon OpenSearch Service Reserved Instance purchase recommendations

Description

Checks your Amazon OpenSearch Service usage patterns to provide recommendations for potential cost savings through Reserved Instance (RI) purchases.

This check generates recommendations at [payer scope for the payer account and at linked scope for a linked account](#).

Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr14n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: Account has a cost savings action identified by Cost Optimization Hub for Amazon OpenSearch Service.

Recommended action

Consider [implementing the recommendation](#). For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Additional resources

- [Reserved Instances in Amazon OpenSearch Service](#)

- [Accessing reservation recommendations](#)
- [Purchasing Reserved Instances \(AWS CLI\)](#)

Report columns

- Status
- Region
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

Amazon RDS cost optimization recommendations for DB instances

Description

Checks your Amazon RDS DB instance configurations and usage patterns to provide recommendations for potential cost savings.

For each resource, Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#) and [AWS Compute Optimizer](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr03n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: RDS DB instance has a cost savings action identified by Cost Optimization Hub.

Recommended action

Consider implementing the recommendation. The recommendation is one of the recommendation types listed in [Understanding cost optimization strategies](#) in the *AWS Cost Management User Guide*. For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Use the detailed recommendations in AWS Compute Optimizer to understand the potential impact of these changes on cost and performance.

Additional resources

- [Viewing RDS database recommendations](#)
- [RDS database metrics](#)
- [DB instance classes](#)
- [Modifying an Amazon RDS DB instance](#)

Report columns

- Status
- Region
- Resource ID
- Recommended Action
- Current Resource Summary
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

Amazon RDS cost optimization recommendations for DB instance storage

Description

Checks your Amazon RDS DB instance storage configurations and usage patterns to provide recommendations for potential cost savings.

For each resource, Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#) and [AWS Compute Optimizer](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr04n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: RDS DB storage has a cost savings action identified by Cost Optimization Hub.

Recommended action

Consider implementing the recommendation. The recommendation is one of the recommendation types listed in [Understanding cost optimization strategies](#) in the *AWS Cost Management User Guide*. For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Use the detailed recommendations in AWS Compute Optimizer to understand the potential impact of these changes on cost and performance.

Additional resources

- [Viewing RDS database recommendations](#)
- [Amazon RDS DB instance storage](#)

Report columns

- Status
- Region

- Resource ID
- Recommended Action
- Current Resource Summary
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

Amazon RDS idle DB instances

Description

Note

This is a legacy check. We recommend using the new check (Check ID: [c1z7kmr03n](#)) that offers additional customized recommendations.

Checks the configuration of your Amazon Relational Database Service (Amazon RDS) for any database (DB) instances that appear to be idle.

If a DB instance has not had a connection for a prolonged period of time, you can delete the instance to reduce costs. A DB instance is considered idle if the instance hasn't had a connection in the past 7 days. If persistent storage is needed for data on the instance, you can use lower-cost options such as taking and retaining a DB snapshot. Manually created DB snapshots are retained until you delete them.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

Ti39halfu8

Alert Criteria

Yellow: An active DB instance has not had a connection in the last 7 days.

Recommended action

Consider taking a snapshot of the idle DB instance and then either stopping it or deleting it. Stopping the DB instance removes some of the costs for it, but does not remove storage costs. A stopped instance keeps all automated backups based upon the configured retention period. Stopping a DB instance usually incurs additional costs when compared to deleting the instance and then retaining only the final snapshot. See [Stopping an Amazon RDS instance temporarily](#) and [Deleting a DB Instance with a Final Snapshot](#).

Additional resources

[Back Up and Restore](#)

Report columns

- Region
- DB Instance Name
- Multi-AZ
- Instance Type
- Storage Provisioned (GB)
- Days Since Last Connection
- Estimated Monthly Savings (On Demand)

Amazon RDS Reserved Instance purchase recommendations

Description

Checks your Amazon RDS usage patterns to provide recommendations for potential cost savings through Reserved Instance (RI) purchases.

This check generates recommendations at [payer scope for the payer account and at linked scope for a linked account](#).

Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr11n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: EBS volume has a cost savings action identified by Cost Optimization Hub.

Recommended action

Consider [implementing the recommendation](#). For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Additional resources

- [Reserved DB instance for Amazon RDS](#)
- [Accessing reservation recommendations](#)
- [Purchasing reserved DB instances for Amazon RDS](#)

Report columns

- Status
- Region
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

Amazon Redshift reserved node purchase recommendations

Description

Checks your Amazon Redshift usage patterns to provide recommendations for potential cost savings through reserved node purchases.

This check generates recommendations at [payer scope for the payer account and at linked scope for a linked account](#).

Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr12n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: Account has a cost savings action identified by Cost Optimization Hub for Amazon Redshift.

Recommended action

Consider [implementing the recommendation](#). For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Additional resources

- [Reserved nodes](#)
- [Accessing reservation recommendations](#)
- [Purchasing a reserved node](#)

Report columns

- Status
- Region
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

Amazon Route 53 Latency Resource Record Sets

Description

Checks for Amazon Route 53 latency record sets that are configured inefficiently.

To allow Amazon Route 53 to route queries to the AWS Region with the lowest network latency, you should create latency resource record sets for a particular domain name (such as example.com) in different Regions. If you create only one latency resource record set for a domain name, all queries are routed to one Region, and you pay extra for latency-based routing without getting the benefits.

Hosted zones created by AWS services won't appear in your check results.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

51fC20e7I2

Alert Criteria

Yellow: Only one latency resource record set is configured for a particular domain name.

Recommended Action

If you have resources in multiple regions, be sure to define a latency resource record set for each region. See [Latency-Based Routing](#).

If you have resources in only one AWS Region, consider creating resources in more than one AWS Region and define latency resource record sets for each; see [Latency-Based Routing](#).

If you don't want to use multiple AWS Regions, you should use a simple resource record set. See [Working with Resource Record Sets](#).

Additional Resources

- [Amazon Route 53 Developer Guide](#)
- [Amazon Route 53 Pricing](#)

Report columns

- Hosted Zone Name
- Hosted Zone ID
- Resource Record Set Name
- Resource Record Set Type

Amazon S3 Bucket Lifecycle Policy Configured

Description

Checks if an Amazon S3 bucket has a lifecycle policy configured. An Amazon S3 lifecycle policy ensures that Amazon S3 objects inside the bucket are stored cost-effectively throughout their lifecycle. This is important for meeting regulatory requirements for data retention and storage. The policy configuration is a set of rules that define actions applied by the Amazon S3 service to a group of objects. A lifecycle policy allows you to automate transitioning objects to lower-cost storage classes or deleting them as they age. For example, you can transition an object to Amazon S3 Standard-IA storage 30 days after creation, or to Amazon Glacier after 1 year.

You can also define object expiration so that Amazon S3 deletes the object on your behalf after a certain period of time.

You can adjust the check configuration using the parameters in your AWS Config rules

For more information, see [Managing your storage lifecycle](#).

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz100

Source

AWS Config Managed Rule: s3-lifecycle-policy-check

Alert Criteria

Yellow: Amazon S3 bucket has no lifecycle policy configured.

Recommended Action

Make sure that you have a lifecycle policy configured in your Amazon S3 bucket.

If your organization does not have a retention policy in place, consider using Amazon S3 Intelligent-Tiering to optimize cost.

For information on how to define your Amazon S3 lifecycle policy, see [Setting lifecycle configuration on a bucket](#).

For information on Amazon S3 Intelligent-Tiering, see [Amazon S3 Intelligent-Tiering storage class](#)

Additional Resources

[Setting lifecycle configuration on a bucket](#)

[Examples of S3 Lifecycle configuration](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters

Amazon S3 Incomplete Multipart Upload Abort Configuration

Description

Checks that each Amazon S3 bucket is configured with a lifecycle rule to abort multipart uploads that remain incomplete after 7 days. Using a lifecycle rule to abort these incomplete uploads and delete the associated storage is recommended.

Note

Results for this check are automatically refreshed one or more times each day, and refresh requests are not allowed. It might take a few hours for changes to appear. It might take a few hours for changes to appear. For Business, Enterprise On-Ramp, or Enterprise Support customers, you can use the `BatchUpdateRecommendationResourceExclusion` API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1cj39rr6v

Alert criteria

Yellow: The lifecycle configuration bucket does not contain a lifecycle rule to abort all multipart uploads that remain incomplete after 7 days.

Recommended action

Review lifecycle configuration for buckets without a lifecycle rule that would cleanup all incomplete multipart uploads. Uploads that are not completed after 24 hours are unlikely to

be completed. Click [here](#) to follow instructions to create a lifecycle rule. It is recommended that this is applied to all objects in your bucket. If you have a need to apply other lifecycle actions to selected objects in your bucket, you can have multiple rules with different filters. Check the storage lens dashboard or call the ListMultipartUpload API for more information.

Additional Resources

[Creating a lifecycle configuration](#)

[Discovering and Deleting Incomplete Multipart Uploads to Lower Amazon S3 Costs](#)

[Uploading and copying objects using multipart upload](#)

[Lifecycle configuration elements](#)

[Elements to describe lifecycle actions](#)

[Lifecycle configuration to abort multipart uploads](#)

Report columns

- Status
- Region
- Bucket Name
- Bucket ARN
- Lifecycle rule for deleting incomplete MPU
- Days After Initiation
- Last Updated Time

Amazon S3 version-enabled buckets without lifecycle policies configured

Description

Checks if Amazon S3 version-enabled buckets have a lifecycle policy configured..

For more information, see [Managing your storage lifecycle](#).

You can specify the bucket names that you want to check using the **bucketNames** parameters in your AWS Config rules.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz171

Source

AWS Config Managed Rule: s3-version-lifecycle-policy-check

Alert criteria

Yellow: An Amazon S3 version-enabled bucket with doesn't have a lifecycle policy configured.

Recommended action

Configure lifecycle policies for your Amazon S3 buckets to manage your objects so that they are stored cost effectively throughout their lifecycle.

For more information, see [Setting lifecycle configuration on a bucket](#).

Additional resources

[Managing your storage lifecycle](#)

[Setting lifecycle configuration on a bucket](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters

- Last Updated Time

AWS Lambda cost optimization recommendations for functions

Description

Checks your AWS Lambda configurations and usage patterns to provide recommendations for potential cost savings.

For each resource, Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#) and [AWS Compute Optimizer](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr05n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: Lambda function has a cost savings action identified by Cost Optimization Hub.

Recommended action

Consider implementing the recommendation. The recommendation is one of the recommendation types listed in [Understanding cost optimization strategies](#) in the *AWS Cost Management User Guide*. For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Use the detailed recommendations in AWS Compute Optimizer to understand the potential impact of these changes on cost and performance.

Additional resources

- [Viewing Lambda function recommendations](#)
- [Lambda function metrics](#)
- [Configuring AWS Lambda functions](#)

Report columns

- Status
- Region
- Resource ID
- Recommended Action
- Current Resource Summary
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

AWS Lambda functions with excessive timeouts

Description

Checks for Lambda functions with high timeout rates that might result in high cost.

Lambda charges based on run time and number of requests for your function. Function timeouts result in errors that may cause retries. Retrying functions will incur additionally request and run time charges.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

L4dfs2Q3C3

Alert criteria

Yellow: Functions where > 10% of invocations end in an error due to a timeout on any given day within the last 7 days.

Recommended action

Inspect function logging and X-ray traces to determine the contributor to the high function duration. Implement logging in your code at relevant parts, such as before or after API calls or database connections. By default, AWS SDK clients timeouts may be longer than the configured function duration. Adjust API and SDK connection clients to retry or fail within the function timeout. If the expected duration is longer than the configured timeout, you can increase the timeout setting for the function. For more information, see [Monitoring and troubleshooting Lambda applications](#).

Additional resources

- [Monitoring and troubleshooting Lambda applications](#)
- [Lambda Function Retry Timeout SDK](#)
- [Using AWS Lambda with AWS X-Ray](#)
- [Accessing Amazon CloudWatch logs for AWS Lambda](#)
- [Error Processor Sample Application for AWS Lambda](#)

Report columns

- Status
- Region
- Function ARN
- Max Daily Timeout Rate
- Date of Max Daily Timeout Rate
- Average Daily Timeout Rate
- Function Timeout Settings (millisecond)
- Lost Daily Compute Cost
- Average Daily Invokes
- Current Day Invokes
- Current Day Timeout Rate

- Last Updated Time

AWS Lambda functions with high error rates

Description

Checks for Lambda functions with high error rates that might result in higher costs.

Lambda charges are based on the number of requests and aggregate run time for your function. Function errors may cause retries that incur additional charges.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

L4dfs2Q3C2

Alert criteria

Yellow: Functions where > 10% of invocations end in error on any given day within the last 7 days.

Recommended action

Consider the following guidelines to reduce errors. Function errors include errors returned by the function's code and errors returned by the function's runtime.

To help you troubleshoot Lambda errors, Lambda integrates with services like Amazon CloudWatch and AWS X-Ray. You can use a combination of logs, metrics, alarms, and X-Ray tracing to quickly detect and identify issues in your function code, API, or other resources that support your application. For more information, see [Monitoring and troubleshooting Lambda applications](#).

For more information on handling errors with specific runtimes, see [Error handling and automatic retries in AWS Lambda](#).

For additional troubleshooting, see [Troubleshooting issues in Lambda](#).

You can also choose from an ecosystem of monitoring and observability tools provided by AWS Lambda partners. For more information, see [AWS Lambda Partners](#).

Additional resources

- [Error Handling and Automatic Retries in AWS Lambda](#)
- [Monitoring and Troubleshooting Lambda applications](#)
- [Lambda Function Retry Timeout SDK](#)
- [Troubleshooting issues in Lambda](#)
- [API Invoke Errors](#)
- [Error Processor Sample Application for AWS Lambda](#)

Report columns

- Status
- Region
- Function ARN
- Max Daily Error Rate
- Date for Max Error Rate
- Average Daily Error Rate
- Lost Daily Compute Cost
- Current Day Invokes
- Current Day Error Rate
- *Average Daily Invokes
- Last Updated Time

AWS Lambda over-provisioned functions for memory size

Description

Note

This is a legacy check. We recommend using the new check (Check ID: [c1z7kmr05n](#)) that offers additional customized recommendations.

Checks the AWS Lambda functions that were invoked at least once during the lookback period. This check alerts you if any of your Lambda functions were over-provisioned for memory size. When you have Lambda functions that are over-provisioned for memory sizes, you're paying for unused resources. Although some scenarios can result in low utilization by design, you can often lower your costs by changing the memory configuration of your Lambda functions. Estimated monthly savings are calculated by using the current usage rate for Lambda functions.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

C0r6dfpM05

Alert criteria

Yellow: A Lambda function that was over-provisioned for memory size during the lookback period. To determine if a Lambda function is over-provisioned, we consider all default CloudWatch metrics for that function. The algorithm used to identify over-provisioned Lambda functions for memory size follows AWS best practices. The algorithm is updated when a new pattern has been identified.

Recommended action

Consider reducing the memory size of your Lambda functions.

For more information, see [Opt in AWS Compute Optimizer for Trusted Advisor checks](#).

Report columns

- Status
- Region
- Function Name
- Function Version
- Memory Size (MB)

- Recommended Memory Size (MB)
- Lookback Period (days)
- Savings Opportunity (%)
- Estimated Monthly Savings
- Estimated Monthly Savings Currency
- Last Updated Time

AWS Savings Plans purchase recommendations for compute

Description

Checks your AWS compute usage patterns across Amazon EC2, AWS Fargate, and AWS Lambda and provides Savings Plans purchase recommendations. With these recommendations, you can commit to a consistent usage amount measured in dollars per hour in exchange for discounted rates.

This check generates recommendations at [payer scope for the payer account and at linked scope for a linked account](#).

Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr09n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: Account has a cost savings action identified by Cost Optimization Hub for compute resources.

Recommended action

Consider [implementing the recommendation](#). For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Additional resources

- [What are Savings Plans?](#)
- [Savings Plans types](#)
- [Purchasing Savings Plans](#)

Report columns

- Status
- Region
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Time stamp

AWS Savings Plans purchase recommendations for Amazon SageMaker AI

Description

Checks your usage of Amazon SageMaker AI and provides Savings Plans purchase recommendations. With these recommendations, you can commit to a consistent usage amount measured in dollars per hour in exchange for discounted rates.

This check generates recommendations at [payer scope for the payer account and at linked scope for a linked account](#).

Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7kmr08n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: Account has a cost savings action identified by Cost Optimization Hub for Amazon SageMaker AI.

Recommended action

Consider [implementing the recommendation](#). For more information on implementing these recommendations, see the AWS Cloud Financial Management (CFM) [Service Cost Optimization Playbook](#).

Additional resources

- [What are Savings Plans?](#)
- [Savings Plans types](#)
- [Purchasing Savings Plans](#)

Report columns

- Status
- Region
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings

- Last Refresh Time stamp

AWS Well-Architected high risk issues for cost optimization

Description

Checks for high risk issues (HRIs) for your workloads in the cost optimization pillar. This check is based on your AWS-Well Architected reviews. Your check results depend on whether you completed the workload evaluation with AWS Well-Architected.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

Wxdfp4B1L1

Alert Criteria

- Red: At least one active high risk issue was identified in the cost optimization pillar for AWS Well-Architected.
- Green: No active high risk issues were detected in the cost optimization pillar for AWS Well-Architected.

Recommended Action

AWS Well-Architected detected high risk issues during your workload evaluation. These issues present opportunities to reduce risk and save money. Sign in to the [AWS Well-Architected](#) tool to review your answers and take action to resolve your active issues.

Report columns

- Status
- Region
- Workload ARN

- Workload Name
- Reviewer Name
- Workload Type
- Workload Started Date
- Workload Last Modified Date
- Number of identified HRIs for Cost Optimization
- Number of HRIs resolved for Cost Optimization
- Number of questions answered for Cost Optimization
- Total number of questions in Cost Optimization pillar
- Last Updated Time

Idle Load Balancers

Description

Checks your Elastic Load Balancing configuration for load balancers that are idle.

Any load balancer that is configured accrues charges. If a load balancer has no associated back-end instances, or if network traffic is severely limited, the load balancer is not being used effectively. This check currently only checks for Classic Load Balancer type within ELB service. It does not include other ELB types (Application Load Balancer, Network Load Balancer).

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

hjLMh88uM8

Alert Criteria

- Yellow: A load balancer has no active back-end instances.
- Yellow: A load balancer has no healthy back-end instances.

- Yellow: A load balancer has had less than 100 requests per day for the last 7 days.

Recommended Action

If your load balancer has no active back-end instances, consider registering instances or deleting your load balancer. See [Registering Your Amazon EC2 Instances with Your Load Balancer](#) or [Delete Your Load Balancer](#).

If your load balancer has no healthy back-end instances, see [Troubleshooting Elastic Load Balancing: Health Check Configuration](#).

If your load balancer has had a low request count, consider deleting your load balancer. See [Delete Your Load Balancer](#).

Additional Resources

- [Managing Load Balancers](#)
- [Troubleshoot Elastic Load Balancing](#)

Report columns

- Region
- Load Balancer Name
- Reason
- Estimated Monthly Savings

Idle NAT gateways

Description

Checks your NAT gateway configurations and usage patterns to identify idle or underutilized NAT gateways that might be candidates for cost optimization.

For each resource, Trusted Advisor shows only the top recommended action from AWS Cost Optimization Hub.

To use this check, you must opt in to [Cost Optimization Hub](#) and [AWS Compute Optimizer](#).

Check ID

c1z7kmr18n

Source

AWS Cost Optimization Hub

Alert criteria

Yellow: NAT gateway has a cost savings action identified by Cost Optimization Hub.

Recommended action

Consider implementing the recommendation to delete the idle NAT gateway. For more information on implementing this recommendation, see the [AWS Cloud Financial Management Service Cost Optimization Playbook](#).

Additional resources

- [Viewing idle resource recommendations](#)
- [Work with NAT gateways](#)
- [NAT gateways](#)

Report columns

- Status
- Region
- Resource ID
- Recommended Action
- Current Resource Summary
- Recommended Resource Summary
- Estimated Monthly Cost
- Estimated Monthly Savings
- Last Refresh Timestamp

Inactive AWS Network Firewall

Description

Checks your AWS Network Firewall endpoints and alerts you when the Network Firewall appears to be inactive.

A Network Firewall is considered to be inactive if all its endpoints have no data processed the last 30 days. Network Firewall endpoints incur hourly charges. This check alerts you to Network Firewall with no data processed in the last 30 days. It's a best practice to either remove unused Network Firewalls or update your architecture.

Check ID

c2v1fg0bfw

Alert Criteria

- Yellow: The Network Firewall processed 0 bytes in the last 30 days.
- Green: The Network Firewall processed more than 0 bytes in the last 30 days.

Recommended Action

If the Network Firewall wasn't used in the last 30 days, then consider deleting the Network Firewall.

If a Transit Gateway is used for inter-VPC communication, then consider deploying your Network Firewalls in a centralized network inspection architectures. This can reduce the hourly charges on inactive Network Firewalls.

Additional Resources

[AWS Network Firewall Pricing](#)

[Inspection Deployment Models with AWS Network Firewall](#)

Report columns

- Status
- Region
- Network Firewall Arn
- VPC Id
- Subnets
- TotalBytesProcessed
- Last Updated Time

Inactive VPC interface endpoints

Description

Checks your VPC interface endpoints and alerts you when the endpoints appear to be inactive. A VPC interface endpoint is considered to be inactive if it has no data processed in the last 30 days. VPC interface endpoints have hourly charges and data processing costs. This check alerts

you about VPC interface endpoints with 0 data processed in the last 30 days. It's a best practice to either remove unused VPC interface endpoints or update your architecture.

Check ID

c2v1fg0jp6

Alert Criteria

- Yellow: VPC interface endpoint has processed 0 bytes in the last 30 days.
- Green: VPC interface endpoint has processed more than 0 bytes in the last 30 days

Recommended Action

If the VPC interface endpoint had not been used in the last 30 days, consider deleting the VPC interface endpoint.

If Transit Gateway is used for inter-VPC communication, then consider deploying your VPC interface endpoints in a centralized architecture to reduce the hourly charges on inactive VPC interface endpoints.

Additional Resources

- [AWS PrivateLink Pricing](#)
- [Centralized access to VPC private endpoints](#)

Report columns

- Status
- Region
- VPC Endpoint Id
- VPC Id
- Subnet Ids
- Service Name
- TotalBytesProcessed
- Last Updated Time

Inactive Gateway Load Balancer endpoints

Description

Checks your Gateway Load Balancer endpoints and warns when they appear to be inactive. A Gateway Load Balancer endpoint is considered to be underutilized if it has no data processed

in the last 30 days. Gateway Load Balancer endpoints have hourly charges and data processed charges. This check alerts you to Gateway Load Balancer endpoints with 0 data processed in the last 30 days. We recommend that you either remove unused Gateway Load Balancer endpoints, or update your architecture.

Check ID

c2v1fg0k35

Alert Criteria

- Yellow: Gateway Load Balancer endpoint processed 0 bytes in the last 30 days
- Green: Gateway Load Balancer endpoint processed more than 0 bytes in the last 30 days

Recommended Action

If the Gateway Load Balancer endpoint has not been used in the last 30 days, consider deleting the VPC endpoint.

If Transit Gateway is used for inter-VPC communication, consider deploying your Gateway Load Balancer endpoints in a centralized network inspection architecture to reduce the hourly charges on inactive Gateway Load Balancer endpoints.

Additional Resources

[AWS PrivateLink Pricing](#)

[Centralized inspection architecture with AWS Gateway Load Balancer and AWS Transit Gateway](#)

Report columns

- Status
- Region
- VPC Endpoint Id
- VPC Id
- Subnet Id
- Service Name
- TotalBytesProcessed
- Last Updated Time

Inactive NAT Gateways

Description

Checks your NAT Gateways for inactive gateways. A NAT Gateway is considered to be inactive if no data (0 bytes) was processed in the last 30 days. NAT Gateways have hourly charges and data processed charges.

Check ID

c2v1fg022t

Alert Criteria

- Yellow: The NAT Gateway processed 0 bytes in the last 30 days
- Green: The NAT Gateway processed more than 0 bytes in the last 30 days

Recommended Action

Consider deleting any NAT Gateways that weren't used in the last 30 days and that aren't required for external network access outside the VPC.

If a Transit Gateway is used for inter-VPC communication, then consider deploying a centralized NAT Gateway for egress to internet architecture. This can reduce the hourly cost from inactive NAT Gateways.

Additional Resources

[NAT Gateway Pricing](#)

[Centralized egress to internet](#)

Report columns

- Status
- Region
- NAT Gateway Id
- Subnet Id
- VPC Id
- TotalBytesFromDest
- TotalBytesFromSrc
- TotalBytes
- Last Updated Time

Low utilization Amazon EC2 instances

Description

Note

This is a legacy check. We recommend using the new check (Check ID: [c1z7kmr00n](#)) that offers additional customized recommendations.

Checks the Amazon Elastic Compute Cloud (Amazon EC2) instances that were running at any time during the last 14 days. This check alerts you if the daily CPU utilization was 10% or less and network I/O was 5 MB or less for at least 4 days.

Running instances generate hourly usage charges. Although some scenarios can result in low utilization by design, you can often lower your costs by managing the number and size of your instances.

Estimated monthly savings are calculated by using the current usage rate for On-Demand Instances and the estimated number of days the instance might be underutilized. Actual savings will vary if you are using Reserved Instances or Spot Instances, or if the instance is not running for a full day. To get daily utilization data, download the report for this check.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

Qch7DwouX1

Alert Criteria

Yellow: An instance had 10% or less daily average CPU utilization and 5 MB or less network I/O on at least 4 of the previous 14 days.

Recommended Action

Consider stopping or terminating instances that have low utilization, or scale the number of instances by using Auto Scaling. For more information, see [Stop and Start Your Instance](#), [Terminate Your Instance](#), and [What is Auto Scaling?](#)

Additional Resources

- [Monitoring Amazon EC2](#)
- [Instance Metadata and User Data](#)
- [Amazon CloudWatch User Guide](#)
- [Auto Scaling Developer Guide](#)

Report columns

- Region/AZ
- Instance ID
- Instance Name
- Instance Type
- Estimated Monthly Savings
- CPU Utilization 14-day Average
- Network I/O 14-Day Average
- Number of Days Low Utilization

Unassociated Elastic IP Addresses

Description

Checks for Elastic IP addresses (EIPs) that are not associated with a running Amazon Elastic Compute Cloud (Amazon EC2) instance.

EIPs are static IP addresses designed for dynamic cloud computing. Unlike traditional static IP addresses, EIPs mask the failure of an instance or Availability Zone by remapping a public IP address to another instance in your account. A nominal charge is imposed for an EIP that is not associated with a running instance.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

Z4AUBRNSmz

Alert Criteria

Yellow: An allocated Elastic IP address (EIP) is not associated with a running Amazon EC2 instance.

Recommended Action

Associate the EIP with a running active instance, or release the unassociated EIP. For more information, see [Associating an Elastic IP Address with a Different Running Instance](#) and [Releasing an Elastic IP Address](#).

Additional Resources

[Elastic IP Addresses](#)

Report columns

- Region
- IP Address

Underutilized Amazon EBS volumes**Description**

Checks Amazon Elastic Block Store (Amazon EBS) volume configurations and warns when volumes appear to be underutilized.

Charges begin when a volume is created. If a volume remains unattached or has very low write activity (excluding boot volumes) for a period of time, the volume is underutilized. We recommend that you remove underutilized volumes to reduce costs.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

DAvU99Dc4C

Alert Criteria

Yellow: A volume is unattached or had less than 1 IOPS per day for the past 7 days.

Recommended Action

Consider creating a snapshot and deleting the volume to reduce costs. For more information, see [Creating an Amazon EBS Snapshot](#) and [Deleting an Amazon EBS Volume](#).

Additional Resources

- [Amazon Elastic Block Store \(Amazon EBS\)](#)
- [Monitoring the Status of Your Volumes](#)

Report columns

- Region
- Volume ID
- Volume Name
- Volume Type
- Volume Size
- Monthly Storage Cost
- Snapshot ID
- Snapshot Name
- Snapshot Age

Note

If you opted in your account for AWS Compute Optimizer, we recommend that you use the Amazon EBS over-provisioned volumes check instead. For more information, see [Opt in AWS Compute Optimizer for Trusted Advisor checks](#).

Underutilized Amazon Redshift Clusters

Description

Checks your Amazon Redshift configuration for clusters that appear to be underutilized.

If an Amazon Redshift cluster has not had a connection for a prolonged period of time, or is using a low amount of CPU, you can use lower-cost options such as downsizing the cluster, or shutting down the cluster and taking a final snapshot. Final snapshots are retained even after you delete your cluster.

Check ID

G31sQ1E9U

Alert Criteria

- Yellow: A running cluster has not had a connection in the last 7 days.
- Yellow: A running cluster had less than 5% cluster-wide average CPU utilization for 99% of the last 7 days.

Recommended Action

Consider shutting down the cluster and taking a final snapshot, or downsizing the cluster. See [Shutting Down and Deleting Clusters](#) and [Resizing a Cluster](#).

Additional Resources

[Amazon CloudWatch User Guide](#)

Report columns

- Status
- Region
- Cluster
- Instance Type

- Reason
- Estimated Monthly Savings

Performance

Improve the performance of your service by checking your service quotas (formerly referred to as limits), so that you can take advantage of provisioned throughput, monitor for overutilized instances, and detect any unused resources.

You can use the following checks for the performance category.

Check names

- [Amazon Aurora DB cluster under-provisioned for read workload](#)
- [Amazon DynamoDB Auto Scaling Not Enabled](#)
- [Amazon EBS Optimization Not Enabled](#)
- [Amazon EBS Provisioned IOPS \(SSD\) Volume Attachment Configuration](#)
- [Amazon EBS under-provisioned volumes](#)
- [Amazon EC2 Auto Scaling Group is not Associated with a Launch Template](#)
- [Amazon EC2 to EBS Throughput Optimization](#)
- [EC2 Virtualization Type is Paravirtual](#)
- [Amazon ECS Memory Hard Limit](#)
- [Amazon EFS Throughput Mode Optimization](#)
- [Amazon RDS autovacuum parameter is turned off](#)
- [Amazon RDS DB clusters support only up to 64 TiB volume](#)
- [Amazon RDS DB instances in the clusters with heterogeneous instance classes](#)
- [Amazon RDS DB instances in the clusters with heterogeneous instance sizes](#)
- [Amazon RDS DB memory parameters are diverging from default](#)
- [Amazon RDS enable_indexonlyscan parameter is turned off](#)
- [Amazon RDS enable_indexscan parameter is turned off](#)
- [Amazon RDS general_logging parameter is turned on](#)
- [Amazon RDS InnoDB_Change_Buffering parameter using less than optimum value](#)
- [Amazon RDS innodb_open_files parameter is low](#)

- [Amazon RDS innodb_stats_persistent parameter is turned off](#)
- [Amazon RDS instance under-provisioned for system capacity](#)
- [Amazon RDS magnetic volume is in use](#)
- [Amazon RDS parameter groups not using huge pages](#)
- [Amazon RDS query cache parameter is turned on](#)
- [Amazon RDS resources instance class update is required](#)
- [Amazon RDS resources major versions update is required](#)
- [Amazon RDS resources using end of support engine edition under license-included](#)
- [Amazon Route 53 Alias Resource Record Sets](#)
- [AWS Lambda under-provisioned functions for memory size](#)
- [AWS Lambda Functions without Concurrency Limit Configured](#)
- [AWS Well-Architected high risk issues for performance](#)
- [CloudFront Alternate Domain Names](#)
- [CloudFront Content Delivery Optimization](#)
- [CloudFront Header Forwarding and Cache Hit Ratio](#)
- [High CPU Utilization Amazon EC2 Instances](#)

Amazon Aurora DB cluster under-provisioned for read workload

Description

Checks whether Amazon Aurora DB cluster has the resources to support a read workload.

Check ID

c1qf5bt038

Alert Criteria

Yellow:

Increased database reads: The database load was high and the database was reading more rows than writing or updating the rows.

Recommended Action

We recommend that you tune your queries to decrease the database load or add a reader DB instance to your DB cluster with the same instance class and size as the writer DB instance in

the cluster. The current configuration has at least one DB instance with a continuously high database load caused mostly by read operations. Distribute these operations by adding another DB instance to the cluster and directing the read workload to the DB cluster read-only endpoint.

Additional Resources

An Aurora DB cluster has one reader endpoint for read-only connections. This endpoint uses load balancing to manage the queries contributing the most to database load in your DB cluster. The reader endpoint directs these statements to the Aurora Read Replicas and reduces the load on the primary instance. The reader endpoint also scales the capacity to handle concurrent SELECT queries with the number of Aurora Read Replicas in the cluster.

For more information, see [Adding Aurora Replicas to a DB Cluster](#) and [Managing performance and scaling for Aurora DB clusters](#).

Report columns

- Status
- Region
- Resource
- Increased database read (count)
- Last detection period
- Last Updated Time

Amazon DynamoDB Auto Scaling Not Enabled

Description

Checks if your Amazon DynamoDB tables and global secondary indexes have auto scaling or on-demand enabled.

Amazon DynamoDB auto scaling uses the Application Auto Scaling service to dynamically adjust provisioned throughput capacity on your behalf in response to actual traffic patterns. This enables a table or a global secondary index to increase its provisioned read and write capacity to handle sudden increases in traffic, without throttling. When the workload decreases, Application Auto Scaling decreases the throughput so that you don't pay for unused provisioned capacity.

You can adjust the check configuration using the parameters in your AWS Config rules.

For more information, see [Managing throughput capacity automatically with DynamoDB auto scaling](#).

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz136

Source

AWS Config Managed Rule: dynamodb-autoscaling-enabled

Alert Criteria

Yellow: Auto scaling is not enabled for your DynamoDB tables and/or global secondary indexes.

Recommended Action

Unless you already have a mechanism to automatically scale the provisioned throughput of your DynamoDB table and/or the global secondary indexes based on your workload requirement, consider turning on auto scaling for your Amazon DynamoDB tables.

For more information, see [Using the AWS Management Console with DynamoDB auto scaling](#).

Additional Resources

[Managing throughput capacity automatically with DynamoDB auto scaling](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon EBS Optimization Not Enabled

Description

Checks if Amazon EBS optimization is enabled for your Amazon EC2 instances.

An Amazon EBS–optimized instance uses an optimized configuration stack and provides additional, dedicated capacity for Amazon EBS I/O. This optimization provides the best performance for your Amazon EBS volumes by minimizing contention between Amazon EBS I/O and other traffic from your instance..

For more information, see [Amazon EBS–optimized instances](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz142

Source

AWS Config Managed Rule: ebs-optimized-instance

Alert Criteria

Yellow: Amazon EBS optimization is not enabled on supported Amazon EC2 instances.

Recommended Action

Turn on Amazon EBS optimization on supported instances.

For more information, see [Enable EBS optimization at launch](#).

Additional Resources

[Amazon EBS–optimized instances](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon EBS Provisioned IOPS (SSD) Volume Attachment Configuration

Description

Checks for Provisioned IOPS (SSD) volumes that are attached to an Amazon EBS optimizable Amazon Elastic Compute Cloud (Amazon EC2) instance that is not EBS-optimized.

Provisioned IOPS (SSD) volumes in the Amazon Elastic Block Store (Amazon EBS) are designed to deliver the expected performance only when they are attached to an EBS-optimized instance.

Check ID

PPkZrjsH2q

Alert Criteria

Yellow: An Amazon EC2 instance that can be EBS-optimized has an attached Provisioned IOPS (SSD) volume but the instance is not EBS-optimized.

Recommended Action

Create a new instance that is EBS-optimized, detach the volume, and reattach the volume to your new instance. For more information, see [Amazon EBS-Optimized Instances](#) and [Attaching an Amazon EBS Volume to an Instance](#).

Additional Resources

- [Amazon EBS Volume Types](#)
- [Amazon EBS Volume Performance](#)

Report columns

- Status
- Region/AZ

- Volume ID
- Volume Name
- Volume Attachment
- Instance ID
- Instance Type
- EBS Optimized

Amazon EBS under-provisioned volumes

Description

Checks the Amazon Elastic Block Store (Amazon EBS) volumes that were running at any time during the lookback period. This check alerts you if any EBS volumes were under-provisioned for your workloads. Consistent high utilization can indicate optimized, steady performance, but can also indicate that an application does not have enough resources.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

C0r6dfpM04

Alert Criteria

Yellow: An EBS Volume that was under-provisioned during the lookback period. To determine if a volume is under-provisioned, we consider all default CloudWatch metrics (including IOPS and throughput). The algorithm used to identify under-provisioned EBS volumes follows AWS best practices. The algorithm is updated when a new pattern has been identified.

Recommended Action

Consider upsizing volumes that have high utilization.

For more information, see [Opt in AWS Compute Optimizer for Trusted Advisor checks](#).

Report columns

- Status
- Region
- Volume ID
- Volume Type
- Volume Size (GB)
- Volume Baseline IOPS
- Volume Burst IOPS
- Volume Burst Throughput
- Recommended Volume Type
- Recommended Volume Size (GB)
- Recommended Volume Baseline IOPS
- Recommended Volume Burst IOPS
- Recommended Volume Baseline Throughput
- Recommended Volume Burst Throughput
- Lookback Period (days)
- Performance Risk
- Last Updated Time

Amazon EC2 Auto Scaling Group is not Associated with a Launch Template

Description

Checks if an Amazon EC2 Auto Scaling group is created from an Amazon EC2 launch template.

Use a launch template to create your Amazon EC2 Auto Scaling groups to ensure access to the latest Auto Scaling group features and improvements. For example, versioning and multiple instance types.

For more information, see [Launch templates](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz102

Source

AWS Config Managed Rule: autoscaling-launch-template

Alert Criteria

Yellow: The Amazon EC2 Auto Scaling group isn't associated with a valid launch template.

Recommended Action

Use an Amazon EC2 launch template to create your Amazon EC2 Auto Scaling groups.

For more information, see [Create a launch template for an Auto Scaling group](#).

Additional Resources

- [Launch templates](#)
- [Create a launch template](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon EC2 to EBS Throughput Optimization

Description

Checks for Amazon EBS volumes whose performance might be affected by the maximum throughput capability of the Amazon EC2 instance they are attached to.

To optimize performance, you should ensure that the maximum throughput of an Amazon EC2 instance is greater than the aggregate maximum throughput of the attached EBS volumes. This check computes the total EBS volume throughput for each five-minute period in the preceding day (based on Coordinated Universal Time (UTC)) for each EBS-optimized instance and alerts you if usage in more than half of those periods was greater than 95% of the maximum throughput of the EC2 instance.

Check ID

Bh2xRR2FGH

Alert Criteria

Yellow: In the preceding day (UTC), the aggregate throughput (megabytes/sec) of the EBS volumes attached to the EC2 instance exceeded 95% of the published throughput between the instance and the EBS volumes more than 50% of time.

Recommended Action

Compare the maximum throughput of your Amazon EBS volumes (see [Amazon EBS Volume Types](#)) with the maximum throughput of the Amazon EC2 instance they are attached to. See [Instance Types That Support EBS Optimization](#).

Consider attaching your volumes to an instance that supports higher throughput to Amazon EBS for optimal performance.

Additional Resources

- [Amazon EBS Volume Types](#)
- [Amazon EBS-Optimized Instances](#)
- [Monitoring the Status of Your Volumes](#)
- [Attaching an Amazon EBS Volume to an Instance](#)
- [Detaching an Amazon EBS Volume from an Instance](#)
- [Deleting an Amazon EBS Volume](#)

Report columns

- Status
- Region
- Instance ID
- Instance Type
- Time Near Maximum

EC2 Virtualization Type is Paravirtual

Description

Checks if the virtualization type of an Amazon EC2 instance is paravirtual.

It's a best practice that you use Hardware Virtual Machine (HVM) instances instead of paravirtual instances, when possible. This is because of enhancements in HVM virtualization and the availability of PV drivers for HVM AMIs, which have closed the performance gap that historically existed between PV and HVM guests. It's important to note that current generation instance types do not support PV AMIs. Therefore, choosing an HVM instance type provides the best performance and compatibility with modern hardware.

For more information, see [Linux AMI virtualization types](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz148

Source

AWS Config Managed Rule: ec2-paravirtual-instance-check

Alert Criteria

Yellow: The virtualization type of Amazon EC2 instances is paravirtual.

Recommended Action

Use HVM virtualization for your Amazon EC2 instances, and use a compatible instance type.

For information on choosing the appropriate virtualization type, see [Compatibility for changing the instance type](#).

Additional Resources

[Compatibility for changing the instance type](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon ECS Memory Hard Limit

Description

Checks if Amazon ECS task definitions have a set memory limit for its container definitions. The total amount of memory reserved for all containers within a task must be lower than the task memory value.

For more information, see [Container definitions](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz176

Source

AWS Config Managed Rule: ecs-task-definition-memory-hard-limit

Alert Criteria

Yellow: Amazon ECS memory hard limit is not set.

Recommended Action

Allocate memory for your Amazon ECS tasks to avoid running out of memory. If your container attempts to exceed the specified memory, then the container is terminated.

For more information, see [How can I allocate memory to tasks in Amazon ECS?](#).

Additional Resources

[Cluster reservation](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon EFS Throughput Mode Optimization

Description

Checks if the Amazon EFS file system is currently configured to use Bursting throughput mode. Amazon EFS Bursting throughput uses “credits” for performance above baseline throughput (50 KiB/s per GiB). After credits are depleted, performance throttles to baseline throughput, which potentially causes slowness, application failures, and timeouts for your users. To learn more about bursting mode, see [Throughput mode](#) in the *Amazon EFS User Guide*

Check ID

c1dfprch02

Alert Criteria

- Yellow: File system is using Bursting throughput mode.

Recommended Action

If your Bursting throughput credits are low or depleted, consider switching to Provisioned or Elastic throughput mode to provide your users and applications with their desired throughput. With Provisioned throughput, you can set your workload's required throughput, and you pay for the amount of throughput enabled for the file system. If you're unsure of your throughput requirements, you can use Elastic throughput mode where your throughput scales elastically with your workload and you only pay for what you use, measured by total data transferred. You can update your file system configuration to switch between the throughput modes at any time. To learn more about throughput pricing, see [Amazon EFS Pricing](#). You can also estimate your costs using the [AWS Pricing Calculator](#).

Additional Resources

- [Bursting throughput](#)
- [Amazon EFS Pricing](#)
- [AWS Pricing Calculator](#)

Report columns

- Status
- Region
- EFS File System ID
- Throughput mode
- Last Updated Time

Amazon RDS autovacuum parameter is turned off

Description

The autovacuum parameter is turned off for your DB instances. Turning autovacuum off increases the table and index bloat and impacts the performance.

We recommend that you turn on autovacuum in your DB parameter groups.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt025

Alert Criteria

Yellow: DB parameter groups have autovacuum turned off.

Recommended Action

Turn on the autovacuum parameter in your DB parameter groups.

Additional Resources

PostgreSQL database requires periodic maintenance which is known as vacuuming. Autovacuum in PostgreSQL automates running **VACUUM** and **ANALYZE** commands. This process gathers the table statistics and deletes the dead rows. When autovacuum is turned off, the increase of the table, index bloat, stale statistics will impact the database performance.

For more information, see [Understanding autovacuum in Amazon RDS for PostgreSQL environments](#).

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS DB clusters support only up to 64 TiB volume

Description

Your DB clusters support volumes up to 64 TiB. The latest engine versions support volumes up to 128 TiB. We recommend that you upgrade the engine version of your DB cluster to latest versions to support volumes up to 128 TiB.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt017

Alert Criteria

Yellow: DB clusters have support for volumes only up to 64 TiB.

Recommended Action

Upgrade the engine version of your DB clusters to support volumes up to 128 TiB.

Additional Resources

When you scale up your application on a single Amazon Aurora DB cluster, you may not reach the limit if the storage limit is 128 TiB. The increased storage limit helps to avoid deleting the data or splitting the database across multiple instances.

For more information, see [Amazon Aurora size limits](#).

Report columns

- Status
- Region
- Resource
- Engine Name
- Engine Version Current
- Recommended Value
- Last Updated Time

Amazon RDS DB instances in the clusters with heterogeneous instance classes

Description

We recommend that you use the same DB instance class and size for all the DB instances in your DB cluster.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

 Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt009

Alert Criteria

Red: DB clusters have the DB instances with heterogeneous instance classes.

Recommended Action

Use the same instance class and size for all the DB instances in your DB cluster.

Additional Resources

When the DB instances in your DB cluster use different DB instance classes or sizes, there can be an imbalance in the workload for the DB instances. During a failover, one of the reader DB instance changes to a writer DB instance. If the DB instances use the same DB instance class and size, the workload can be balanced for the DB instances in your DB cluster.

For more information, see [Aurora Replicas](#).

Report columns

- Status
- Region
- Resource

- Recommended Value
- Engine Name
- Last Updated Time

Amazon RDS DB instances in the clusters with heterogeneous instance sizes

Description

We recommend that you use the same DB instance class and size for all the DB instances in your DB cluster.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt008

Alert Criteria

Red: DB clusters have the DB instances with heterogeneous instance sizes.

Recommended Action

Use the same instance class and size for all the DB instances in your DB cluster.

Additional Resources

When the DB instances in your DB cluster use different DB instance classes or sizes, there can be an imbalance in the workload for the DB instances. During a failover, one of the reader DB instance changes to a writer DB instance. If the DB instances use the same DB instance class and size, the workload can be balanced for the DB instances in your DB cluster.

For more information, see [Aurora Replicas](#).

Report columns

- Status
- Region
- Resource
- Recommended Value
- Engine Name
- Last Updated Time

Amazon RDS DB memory parameters are diverging from default

Description

The memory parameters of the DB instances are significantly different from the default values. These settings can impact performance and cause errors.

We recommend that you reset the custom memory parameters for the DB instance to their default values in the DB parameter group.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt020

Alert Criteria

Yellow: DB parameter groups have memory parameters that diverge considerably from the default values.

Recommended Action

Reset the memory parameters to their default values.

Additional Resources

For more information, see [Best practices for configuring parameters for Amazon RDS for MySQL, part 1: Parameters related to performance](#).

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS enable_indexonlyscan parameter is turned off

Description

The query planner or optimizer can't use the index-only scan plan type when it is turned off.

We recommend that you set the **enable_indexonlyscan** parameter value to 1.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt028

Alert Criteria

Yellow: DB parameter groups have **enable_indexonlyscan** parameter turned off.

Recommended Action

Set the parameter **enable_indexonlyscan** to 1.

Additional Resources

When you turn off **enable_indexonlyscan** parameter, it prevents the query planner from selecting an optimal execution plan. The query planner uses a different plan type, such as index scan which can increase the query cost and execution time. The index only scan plan type retrieves the data without accessing the table data.

For more information, see [enable_indexonlyscan \(boolean\)](#) on the PostgreSQL documentation website.

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS enable_indexscan parameter is turned off

Description

The query planner or optimizer can't use the index scan plan type when it is turned off.

We recommend that you set the **enable_indexscan** parameter value to 1.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

 Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt029

Alert Criteria

Yellow: DB parameter groups have **enable_indexscan** parameter turned off.

Recommended Action

Set the parameter **enable_indexscan** to 1.

Additional Resources

When you turn off **enable_indexscan** parameter, it prevents the query planner from selecting an optimal execution plan. The query planner uses a different plan type, such as index scan which can increase the query cost and execution time.

For more information, see [enable_indexscan \(boolean\)](#) on the PostgreSQL documentation website.

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS general_logging parameter is turned on

Description

The general logging is turned on for your DB instance. This setting is useful while troubleshooting the database issues. However, turning on general logging increases the amount of I/O operations and allocated storage space, which might result in contention and performance degradation.

Check your requirements for general logging usage. We recommend that you set the **general_logging** parameter value to **0**.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt037

Alert Criteria

Yellow: DB parameter groups have **general_logging** turned on.

Recommended Action

Check your requirements for general logging usage. If it isn't mandatory, we recommend that you to set the **general_logging** parameter value to **0**.

Additional Resources

The general query log is turned on when the **general_logging** parameter value is 1. The general query log contains records of the database server operations. The server writes information to this log when clients connect or disconnect and the logs contain each SQL statement received from the clients. The general query log is useful when you suspect an error in a client and you want to find the information the client to sent to the database server.

For more information, see [Overview of RDS for MySQL database logs](#).

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS InnoDB_Change_Buffering parameter using less than optimum value

Description

Change buffering allows a MySQL DB instance to defer a few writes, which are required to maintain secondary indexes. This feature was useful in environments with slow disks. The change buffering configuration improved the DB performance slightly but caused a delay in crash recovery and long shutdown times during upgrade.

We recommend that you set the value of **innodb_change_buffering** parameter to **NONE**.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

 Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt021

Alert Criteria

Yellow: DB parameter groups have **innodb_change_buffering** parameter set to a low optimum value.

Recommended Action

Set **innodb_change_buffering** parameter value to **NONE** in your DB parameter groups.

Additional Resources

For more information, see [Best practices for configuring parameters for Amazon RDS for MySQL, part 1: Parameters related to performance](#).

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value

- Last Updated Time

Amazon RDS innodb_open_files parameter is low

Description

The `innodb_open_files` parameter controls the number of files InnoDB can open at one time. InnoDB opens all of the log and system tablespace files when `mysqld` is running.

Your DB instance has a low value for the maximum number of files InnoDB can open at one time. We recommend that you set the `innodb_open_files` parameter to a minimum value of 65.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt033

Alert Criteria

Yellow: DB parameter groups have the InnoDB open files setting misconfigured.

Recommended Action

Set the `innodb_open_files` parameter to a minimum value of 65.

Additional Resources

The `innodb_open_files` parameter controls the number of files InnoDB can open at one time. InnoDB keeps all the log files and the system tablespace files open when `mysqld` is running. InnoDB also needs to open a few `.ibd` files, if file-per-table storage model is used. When the `innodb_open_files` setting is low, it impacts the database performance and the server may fail to start.

For more information, see [InnoDB Startup Options and System Variables - innodb_open_files](#) on the MySQL documentation website.

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS `innodb_stats_persistent` parameter is turned off

Description

Your DB instance isn't configured to persist the InnoDB statistics to the disk. When the statistics aren't stored, they are recalculated each time the instance restarts and the table accessed. This leads to variations in the query execution plan. You can modify the value of this global parameter at the table level.

We recommend that you set the `innodb_stats_persistent` parameter value to **ON**.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

 Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt032

Alert Criteria

Yellow: DB parameter groups have optimizer statistics that aren't persisted to the disk.

Recommended Action

Set the **innodb_stats_persistent** parameter value to **ON**.

Additional Resources

If the **innodb_stats_persistent** parameter is set to **ON**, then the optimizer statistics are persisted when the instance restarts. This improves the execution plan stability and consistent query performance. You can modify global statistics persistence at the table level by using the clause **STATS_PERSISTENT** when you create or alter a table.

For more information, see [Best practices for configuring parameters for Amazon RDS for MySQL, part 1: Parameters related to performance](#).

Report columns

- Status
- Region

- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS instance under-provisioned for system capacity

Description

Checks whether Amazon RDS instance or Amazon Aurora DB instance has the required system capacity to operate.

Check ID

c1qf5bt039

Alert Criteria

Yellow:

Out-of-memory kills: When a process on the database host is stopped because of memory reduction at the OS level, the Out Of Memory (OOM) kills counter increases.

Excessive swapping: os.memory.swap.in and os.memory.swap.out metric values were high.

Recommended Action

We recommend that you tune your queries to use less memory or use a DB instance type with higher allocated memory. When the instance is running low on memory, this impacts the database performance.

Additional Resources

Out-of-memory kills were detected: Linux kernel invokes the Out of Memory (OOM) Killer when the processes running on the host require more than the memory physically available from the operating system. In this case, the OOM Killer reviews all the running processes, and stops one or more processes, in order to free up system memory and keep the system running.

Swapping is detected: When the memory isn't sufficient on the database host, the operating system sends a few minimum used pages to the disk in the swap space. This offloading process impacts the database performance.

For more information, see [Amazon RDS Instance Types](#) and [Scaling your Amazon RDS instance](#).

Report columns

- Status
- Region
- Resource
- Out-of-memory kills (count)
- Excessive swapping (count)
- Last detection period
- Last Updated Time

Amazon RDS magnetic volume is in use

Description

Your DB instances are using magnetic storage. Magnetic storage isn't recommended for most of the DB instances. Choose a different storage type: General Purpose (SSD) or Provisioned IOPS.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt000

Alert Criteria

Yellow: Amazon RDS resources are using magnetic storage.

Recommended Action

Choose a different storage type: General Purpose (SSD) or Provisioned IOPS.

Additional Resources

Magnetic storage is an earlier generation storage type. The General Purpose (SSD) or Provisioned IOPS is the recommended storage type for new storage requirements. These storage types provide higher and consistent performance, and improved storage size options.

For more information, see [Previous generation volumes](#).

Report columns

- Status
- Region
- Resource
- Recommended Value
- Engine Name
- Last Updated Time

Amazon RDS parameter groups not using huge pages

Description

Large pages can increase database scalability, but your DB instance isn't using large pages. We recommend that you set the **use_large_pages** parameter value to **ONLY** in the DB parameter group for your DB instance.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

 Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt024

Alert Criteria

Yellow: DB parameter groups don't use large pages.

Recommended Action

Set the **use_large_pages** parameter value to **ONLY** in your DB parameter groups.

Additional Resources

For more information, see [Turning on HugePages for an RDS for Oracle instance](#).

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS query cache parameter is turned on

Description

When changes require that your query cache is purged, your DB instance will appear to stall. Most workloads don't benefit from a query cache. The query cache was removed from MySQL version 8.0. We recommend that you set the `query_cache_type` parameter to 0.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt022

Alert Criteria

Yellow: DB parameter groups have query cache turned on.

Recommended Action

Set the `query_cache_type` parameter value to 0 in your DB parameter groups.

Additional Resources

For more information, see [Best practices for configuring parameters for Amazon RDS for MySQL, part 1: Parameters related to performance](#).

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS resources instance class update is required

Description

Your database is running a previous generation DB instance class. We have replaced DB instance classes from a previous generation with DB instance classes with better cost, performance, or both. We recommend that you run your DB instance with a DB instance class from a newer generation.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt015

Alert Criteria

Red: DB instances are using end of support DB instance class.

Recommended Action

Upgrade to latest DB instance class.

Additional Resources

For more information, see [Supported DB engines for DB instance classes](#).

Report columns

- Status
- Region
- Resource
- DB Instance Class
- Recommended Value
- Engine Name
- Last Updated Time

Amazon RDS resources major versions update is required

Description

Databases with the current major version for the DB engine won't be supported. We recommend that you upgrade to the latest major version which includes new functionality and enhancements.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt014

Alert Criteria

Red: RDS resources are using end of support major versions.

Recommended Action

Upgrade to the latest major version for the DB engine.

Additional Resources

Amazon RDS releases new versions for the supported database engines to maintain your databases with the latest version. The new released versions may include bug fixes, security enhancements, and other improvements for the database engine. You can minimize the downtime required for the DB instance upgrade by using a blue/green deployment.

For more information, see the following resources:

- [Upgrading a DB instance engine version](#)

- [Amazon Aurora updates](#)
- [Using Amazon RDS Blue/Green Deployments for database updates](#)

Report columns

- Status
- Region
- Resource
- Engine Name
- Engine Current Version
- Recommended Value
- Last Updated Time

Amazon RDS resources using end of support engine edition under license-included

Description

We recommend that you upgrade the major version to the latest engine version supported by Amazon RDS to continue with the current license support. The engine version of your database won't be supported with the current license.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt016

Alert Criteria

Red: Amazon RDS resources are using end of support engine edition under license-included model.

Recommended Action

We recommend that you upgrade your database to the latest supported version in Amazon RDS to continue using the licensed model.

Additional Resources

For more information, see [Oracle major version upgrades](#).

Report columns

- Status
- Region
- Resource
- Engine Name
- Engine Version Current
- Recommended Value
- Engine Name
- Last Updated Time

Amazon Route 53 Alias Resource Record Sets

Description

Checks for resource record sets that can be changed to alias resource record sets to improve performance and save money.

An alias resource record set routes DNS queries to an AWS resource (for example, an Elastic Load Balancing load balancer or an Amazon S3 bucket) or to another Route 53 resource record set. When you use alias resource record sets, Route 53 routes your DNS queries to AWS resources free of charge.

Hosted zones created by AWS services won't appear in your check results.

 Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

B913Ef6fb4

Alert Criteria

- Yellow: A resource record set is a CNAME to an Amazon S3 website.
- Yellow: A resource record set is a CNAME to an Amazon CloudFront distribution.
- Yellow: A resource record set is a CNAME to an Elastic Load Balancing load balancer.

Recommended Action

Replace the listed CNAME resource record sets with alias resource record sets; see [Choosing Between Alias and Non-Alias Resource Record Sets](#).

You also need to change the record type from CNAME to A or AAAA, depending on the AWS resource. See [Values that You Specify When You Create or Edit Amazon Route 53 Resource Record Sets](#).

Additional Resources

[Routing Queries to AWS Resources](#)

Report columns

- Status
- Hosted Zone Name
- Hosted Zone ID

- Resource Record Set Name
- Resource Record Set Type
- Resource Record Set Identifier
- Alias Target

AWS Lambda under-provisioned functions for memory size

Description

Checks the AWS Lambda functions that were invoked at least once during the lookback period. This check alerts you if any of your Lambda functions were under-provisioned for memory size. When you have Lambda functions that are under-provisioned for memory size, these functions take longer time to complete.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

C0r6dfpM06

Alert Criteria

Yellow: A Lambda function that was under-provisioned for memory size during the lookback period. To determine if a Lambda function is under-provisioned, we consider all default CloudWatch metrics for that function. The algorithm used to identify under-provisioned Lambda functions for memory size follows AWS best practices. The algorithm is updated when a new pattern has been identified.

Recommended Action

Consider increasing the memory size of your Lambda functions.

For more information, see [Opt in AWS Compute Optimizer for Trusted Advisor checks](#).

Report columns

- Status
- Region
- Function Name
- Function Version
- Memory Size (MB)
- Recommended Memory Size (MB)
- Lookback Period (days)
- Performance Risk
- Last Updated Time

AWS Lambda Functions without Concurrency Limit Configured

Description

Checks if AWS Lambda functions are configured with function-level concurrent execution limit.

Concurrency is the number of in-flight requests your AWS Lambda function is handling at the same time. For each concurrent request, Lambda provisions a separate instance of your execution environment.

You can specify the minimum and maximum reserved concurrency limit using the **concurrencyLimitLow** and **ConcurrencyLimitHigh** parameters in your AWS Config rules.

For more information, see [Lambda function scaling](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz181

Source

AWS Config Managed Rule: lambda-concurrency-check

Alert Criteria

Yellow: Lambda function has no concurrency limit configured.

Recommended Action

Make sure that your Lambda functions have concurrency configured. A concurrency limit for your Lambda functions helps make sure that your function processes requests reliably and predictably. A concurrency limit reduces the risk of your function being overwhelmed due to a sudden surge in traffic.

For more information, see [Configuring reserved concurrency](#).

Additional Resources

- [Lambda function scaling](#)
- [Configuring reserved concurrency](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

AWS Well-Architected high risk issues for performance

Description

Checks for high risk issues (HRIs) for your workloads in the performance pillar. This check is based on your AWS-Well Architected reviews. Your check results depend on whether you completed the workload evaluation with AWS Well-Architected.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

Wxdfp4B1L2

Alert Criteria

- Red: At least one active high risk issue was identified in the performance pillar for AWS Well-Architected.
- Green: No active high risk issues were detected in the performance pillar for AWS Well-Architected.

Recommended Action

AWS Well-Architected detected high risk issues during your workload evaluation. These issues present opportunities to reduce risk and save money. Sign in to the [AWS Well-Architected](#) tool to review your answers and take action to resolve your active issues.

Report columns

- Status
- Region
- Workload ARN
- Workload Name
- Reviewer Name
- Workload Type
- Workload Started Date
- Workload Last Modified Date
- Number of identified HRIs for Performance
- Number of HRIs resolved for Performance

- Number of questions answered for Performance
- Total number of questions in Performance pillar
- Last Updated Time

CloudFront Alternate Domain Names

Description

Note

This check applies to classic Amazon CloudFront distributions.

Checks that DNS is properly configured in classic Amazon CloudFront distributions that use alternate domain names (CNAMEs).

If a CloudFront distribution includes alternate domain names, the DNS configuration for the domains must route DNS queries to that distribution.

Note

This check assumes Amazon Route 53 DNS and Amazon CloudFront distribution are configured in the same AWS account. As such the alert list might include resources otherwise working as expected due to DNS setting outside of this AWS account.

Check ID

N420c450f2

Alert Criteria

- Yellow: A CloudFront distribution includes alternate domain names, but the DNS configuration is not correctly set up with a CNAME record or an Amazon Route 53 alias resource record.
- Yellow: A CloudFront distribution includes alternate domain names, but Trusted Advisor could not evaluate the DNS configuration because there were too many redirects.
- Yellow: A CloudFront distribution includes alternate domain names, but Trusted Advisor could not evaluate the DNS configuration for some other reason, most likely because of a timeout.

Recommended Action

Update the DNS configuration to route DNS queries to the CloudFront distribution; see [Using Alternate Domain Names \(CNAMEs\)](#).

If you're using Amazon Route 53 as your DNS service, see [Routing Traffic to an Amazon CloudFront Web Distribution by Using Your Domain Name](#). If the check timed out, try refreshing the check.

Additional Resources

[Amazon CloudFront Developer Guide](#)

Report columns

- Status
- Distribution ID
- Distribution Domain Name
- Alternate Domain Name
- Reason

CloudFront Content Delivery Optimization

Description

Checks for cases where data transfer from Amazon Simple Storage Service (Amazon S3) buckets could be accelerated by using Amazon CloudFront, the AWS global content delivery service.

When you configure CloudFront to deliver your content, requests for your content are automatically routed to the nearest edge location where content is cached. This routing allows content to be delivered to your users with the best possible performance. A high ratio of data transferred out compared to the data stored in the bucket indicates that you could benefit from using Amazon CloudFront to deliver the data.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

796d6f3D83

Alert Criteria

- Yellow: The amount of data transferred out of the bucket to your users by GET requests in the 30 days preceding the check is at least 25 times greater than the average amount of data stored in the bucket.
- Red: The amount of data transferred out of the bucket to your users by GET requests in the 30 days preceding the check is at least 10 TB and at least 25 times greater than the average amount of data stored in the bucket.

Recommended Action

Consider using CloudFront for better performance. See [Amazon CloudFront Product Details](#).

If the data transferred is 10 TB per month or more, see [Amazon CloudFront Pricing](#) to explore possible cost savings.

Additional Resources

- [Amazon CloudFront Developer Guide](#)
- [AWS Case Study: PBS](#)

Report columns

- Status
- Region
- Bucket Name
- S3 Storage (GB)
- Data Transfer Out (GB)
- Ratio of Transfer to Storage

CloudFront Header Forwarding and Cache Hit Ratio

Description

Note

This check applies to classic Amazon CloudFront distributions.

Checks the HTTP request headers that CloudFront currently receives from the client and forwards to your origin server.

Some headers, such as date, or user-agent, significantly reduce the cache hit ratio (the proportion of requests that are served from a CloudFront edge cache). This increases the load on your origin and reduces performance, because CloudFront must forward more requests to your origin.

Check ID

N415c450f2

Alert Criteria

Yellow: One or more request headers that CloudFront forwards to your origin might significantly reduce your cache hit ratio.

Recommended Action

Consider whether the request headers provide enough benefit to justify the negative effect on the cache hit ratio. If your origin returns the same object regardless of the value of a given header, we recommend that you don't configure CloudFront to forward that header to the origin. For more information, see [Configuring CloudFront to Cache Objects Based on Request Headers](#).

Additional Resources

- [Increasing the Proportion of Requests that Are Served from CloudFront Edge Caches](#)
- [CloudFront Cache Statistics Reports](#)
- [HTTP Request Headers and CloudFront Behavior](#)

Report columns

- Distribution ID
- Distribution Domain Name
- Cache Behavior Path Pattern
- Headers

High CPU Utilization Amazon EC2 Instances

Description

Checks the Amazon Elastic Compute Cloud (Amazon EC2) instances that were running at any time during the last 14 days. An alert is sent if daily CPU utilization was greater than 90% on four or more days.

Consistent high utilization can indicate optimized, steady performance. However, it can also indicate that an application does not have enough resources. To get daily CPU utilization data, download the report for this check.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

ZRxQ1Psb6c

Alert Criteria

Yellow: An instance had more than 90% daily average CPU utilization on at least 4 of the previous 14 days.

Recommended Action

Consider adding more instances. For information about scaling the number of instances based on demand, see [What is Auto Scaling?](#)

Additional Resources

- [Monitoring Amazon EC2](#)
- [Instance Metadata and User Data](#)
- [Amazon CloudWatch User Guide](#)
- [Amazon EC2 Auto Scaling User Guide](#)

Report columns

- Region/AZ

- Instance ID
- Instance Type
- Instance Name
- 14-Day Average CPU Utilization
- Number of Days over 90% CPU Utilization

Security

You can use the following checks for the security category.

Note

If you enabled Security Hub CSPM for your AWS account, you can view your findings in the Trusted Advisor console. For information, see [Viewing AWS Security Hub CSPM controls in AWS Trusted Advisor](#).

You can view all controls in the AWS Foundational Security Best Practices security standard *except* for controls that have the **Category: Recover > Resilience**. For a list of supported controls, see [AWS Foundational Security Best Practices controls](#) in the *AWS Security Hub CSPM User Guide*.

Check names

- [Application Load Balancer security group](#)
- [Amazon CloudWatch Log Group Retention Period](#)
- [Amazon EC2 instances with Microsoft SQL Server end of support](#)
- [Amazon EC2 instances with Microsoft Windows Server end of support](#)
- [Amazon EC2 instances with Ubuntu LTS end of standard support](#)
- [Amazon EFS clients not using data-in-transit encryption](#)
- [Amazon EBS Public Snapshots](#)
- [Amazon RDS Aurora storage encryption is turned off](#)
- [Amazon RDS engine minor version upgrade is required](#)
- [Amazon RDS Public Snapshots](#)
- [Amazon RDS Security Group Access Risk](#)

- [Amazon RDS storage encryption is turned off](#)
- [Amazon Route 53 mismatching CNAME records pointing directly to S3 buckets](#)
- [Amazon Route 53 MX Resource Record Sets and Sender Policy Framework](#)
- [Amazon S3 Bucket Permissions](#)
- [Amazon VPC Peering Connections with DNS Resolution Disabled](#)
- [Application Load Balancer Target Groups Encrypted Protocol](#)
- [AWS Backup Vault Without Resource-based Policy to Prevent Deletion of Recovery Points](#)
- [AWS CloudTrail Management Event Logging](#)
- [AWS Lambda Functions Using Deprecated Runtimes](#)
- [AWS Well-Architected high risk issues for security](#)
- [CloudFront Custom SSL Certificates in the IAM Certificate Store](#)
- [CloudFront SSL Certificate on the Origin Server](#)
- [ELB Listener Security](#)
- [Classic Load Balancer Security Groups](#)
- [Exposed Access Keys](#)
- [IAM Access Key Rotation](#)
- [IAM Access Analyzer External Access](#)
- [IAM Password Policy](#)
- [IAM SAML 2.0 Identity Provider](#)
- [MFA on root account](#)
- [Root User Access Key](#)
- [Security Groups – Specific Ports Unrestricted](#)
- [Security Groups – Unrestricted Access](#)

Application Load Balancer security group

Description

Checks the security groups attached to the Application Load Balancer and its Amazon EC2 targets. Application Load Balancer security groups should only allow inbound ports that are

configured in a listener. A target's security groups should not accept direct connections from the internet in the same port the target receives traffic from the load balancer.

If a security group allows access to ports that are not configured for the load balancer or allows direct access to targets, the risk of loss of data or malicious attacks increases.

This check excludes the following groups:

- Target Groups that are not associated with IP addresses or EC2 instances.
- Security group rules for IPv6 traffic.

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

8604e947f2

Alert Criteria

- Red: Target has a public IP and a security group that allows inbound connections on the target control port from everywhere (0.0.0.0/0).
- Red: Target has a public IP and a security group that allows inbound connections on the traffic port from everywhere (0.0.0.0/0).
- Red: Application Load Balancer has authentication enabled and target allows inbound connections on the traffic port from everywhere (0.0.0.0/0).
- Yellow: Target's security group allow inbound connections on the traffic port from everywhere (0.0.0.0/0).
- Yellow: Target's security group allow inbound connections on the target control port from everywhere (0.0.0.0/0).
- Yellow: Application Load Balancer security group allow inbound connections on ports that don't have a corresponding listener.

- Yellow: Target's security group allow inbound connections on the target control port from a security group that is not attached to Application Load Balancer.
- Green: Application Load Balancer security group only allows inbound connections on ports that match with a listener.

Recommended Action

For improved security, make sure that your security groups only allow the necessary traffic flows:

- The Application Load Balancer's security groups should allow inbound connections only for the same ports configured in its listeners.
- Use exclusive security groups for load balancers and targets.
- Target security groups should allow connections in the traffic port only from the load balancer(s) it's associated with.
- Target security groups should allow connections in the target control port only from the load balancer(s) it's associated with.

Additional Resources

- [Control traffic to your AWS resources using security groups](#)
- [Security groups for your Application Load Balancer](#)

Report columns

- Status
- Region
- Target Group
- ALB Name
- ALB SG ID
- Target SG ID
- Auth Enabled
- Last Updated Time

Amazon CloudWatch Log Group Retention Period

Description

Checks if Amazon CloudWatch log group retention period is set to 365 days or other specified number.

By default, logs are kept indefinitely and never expire. However, you can adjust the retention policy for each log group to comply with industry regulations or legal requirements for a specific period.

You can specify the minimum retention time and log group names using the **LogGroupNames** and **MinRetentionTime** parameters in your AWS Config rules.

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz186

Source

AWS Config Managed Rule: cw-loggroup-retention-period-check

Alert Criteria

Yellow: Retention period of an Amazon CloudWatch log group is less than the desired minimum number of days.

Recommended Action

Configure a retention period of more than 365 days for your log data stored in Amazon CloudWatch Logs to meet compliance requirements.

For more information, see [Change log data retention in CloudWatch Logs](#).

Additional Resources

[Altering CloudWatch log retention](#)

Report columns

- Status
- Region

- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon EC2 instances with Microsoft SQL Server end of support

Description

Checks the SQL Server versions for Amazon Elastic Compute Cloud (Amazon EC2) instances running in the past 24 hours. This check alerts you if the versions are near or have reached the end of support. Each SQL Server version offers 10 years of support, including 5 years of mainstream support and 5 years of extended support. After the end of support, the SQL Server version won't receive regular security updates. Running applications with unsupported SQL Server versions can bring security or compliance risks.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

Qsdfp3A4L3

Alert Criteria

- Red: An EC2 instance has an SQL Server version that reached the end of support.
- Yellow: An EC2 instance has an SQL Server version that will reach the end of support in 12 months.

Recommended Action

To modernize your SQL Server workloads, consider refactoring to AWS Cloud native databases like Amazon Aurora. For more information, see [Modernize Windows Workloads with AWS](#).

To move to a fully managed database, consider replatforming to Amazon Relational Database Service (Amazon RDS). For more information, see [Amazon RDS for SQL Server](#).

To upgrade your SQL Server on Amazon EC2, consider using the automation runbook to simplify your upgrade. For more information, see the [AWS Systems Manager documentation](#).

If you can't upgrade your SQL Server on Amazon EC2, consider the End-of-Support Migration Program (EMP) for Windows Server. For more information, see the [EMP Website](#).

Additional Resources

- [Get ready for SQL Server end of support with AWS](#)
- [Microsoft SQL Server on AWS](#)

Report columns

- Status
- Region
- Instance ID
- SQL Server Version
- Support Cycle
- End of Support
- Last Updated Time

Amazon EC2 instances with Microsoft Windows Server end of support

Description

This check alerts you if your Microsoft Windows Server versions are near or have reached the end of support. Each Windows Server version offers 10 years of support, including 5 years of mainstream support and 5 years of extended support. After the end of support, the Windows Server version won't receive regular security updates. Running applications with unsupported Windows Server versions can bring security or compliance risks.

Note

This check generates results based on the AMI used to launch the EC2 instance. It's possible for the current instance operating system to be different from its launch AMI.

For example, if you launched an instance from a Windows Server 2016 AMI and later upgrade to Windows Server 2019, the launch AMI doesn't change.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

Qsdfp3A4L4

Alert criteria

- Red: An EC2 instance runs on a Windows Server version that reached the end of support (Windows Server 2003, 2003 R2, 2008, and 2008 R2).
- Yellow: An EC2 instance runs on a Windows Server version that will reach the end of support in less than 18 months (Windows Server 2012 and 2012 R2).

Recommended action

To modernize your Windows Server workloads, consider the various options available on [Modernize Windows Workloads with AWS](#).

To upgrade your Windows Server workloads to run on more recent versions of Windows Server, you can use an automation runbook. For more information, see the [AWS Systems Manager documentation](#).

Complete the following steps:

- Upgrade the Windows Server version
- Hard stop and start upon upgrading
- If using EC2Config, migrate to EC2Launch

Report columns

- Status
- Region

- Instance ID
- Windows Server Version
- Support Cycle
- End of Support
- Last Updated Time

Amazon EC2 instances with Ubuntu LTS end of standard support

Description

This check alerts you if the versions are near or have reached the end of standard support. It is important to take action – either by migrating to the next LTS or upgrading to Ubuntu Pro. After the end of support, your 18.04 LTS machines will not receive any security updates. With an Ubuntu Pro subscription, your Ubuntu 18.04 LTS deployment can receive Expanded Security Maintenance (ESM) until 2028. Security vulnerabilities that remain unpatched open your systems to hackers and the potential of a major breach.

Note

Results for this check are automatically refreshed at least once daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1dfprch15

Alert Criteria

Red: An Amazon EC2 instance has an Ubuntu version that reached the end of standard support (Ubuntu 18.04 LTS, 18.04.1 LTS, 18.04.2 LTS, 18.04.3 LTS, 18.04.4 LTS, 18.04.5 LTS, and 18.04.6 LTS).

Yellow: An Amazon EC2 instance has an Ubuntu version that will reach the end of standard support in less than 6 months (Ubuntu 20.04 LTS, 20.04.1 LTS, 20.04.2 LTS, 20.04.3 LTS, 20.04.4 LTS, 20.04.5 LTS, and 20.04.6 LTS).

Green: All Amazon EC2 instances are compliant.

Recommended Action

To upgrade the Ubuntu 18.04 LTS instances to a supported LTS version, please follow the steps mentioned in [this article](#). To upgrade the Ubuntu 18.04 LTS instances to [Ubuntu Pro](#), visit AWS License Manager console and follow the steps mentioned in the [AWS License Manager user guide](#). You can also refer to the [Ubuntu blog](#) showing a step by step demo of upgrading Ubuntu instances to Ubuntu Pro.

Additional Resources

For information about pricing, reach out to [Support](#).

Report columns

- Status
- Region
- Ubuntu Lts Version
- Expected End Of Support Date
- Instance ID
- Support Cycle
- Last Updated Time

Amazon EFS clients not using data-in-transit encryption

Description

Checks if Amazon EFS file system is mounted using data-in-transit encryption. AWS recommends that customers use data-in-transit encryption for all data flows to protect data from accidental exposure or unauthorized access. Amazon EFS recommends clients use the '-o tls' mount setting using the Amazon EFS mount helper to encrypt data in transit using TLS v1.2.

Check ID

c1dfpnchv1

Alert Criteria

Yellow: One or more NFS clients for your Amazon EFS file system are not using the recommended mount settings that provide data-in-transit encryption.

Green: All NFS clients for your Amazon EFS file system are using the recommended mount settings that provide data-in-transit encryption.

Recommended Action

To take advantage of data-in-transit encryption feature on Amazon EFS, we recommend that you remount your file system using the Amazon EFS mount helper and the recommended mount settings.

Note

Some Linux distributions don't include a version of stunnel that supports TLS features by default. If you're using an unsupported Linux distribution (see [Supported distributions](#) in the *Amazon Elastic File System User Guide*), then it's a best practice that you upgrade it before remounting with the recommended mount setting.

Additional Resources

- [Encrypting data in transit](#)

Report columns

- Status
- Region
- EFS File System ID
- AZs with Unencrypted Connections
- Last Updated Time

Amazon EBS Public Snapshots

Description

Checks the permission settings for your Amazon Elastic Block Store (Amazon EBS) volume snapshots and alerts you if any snapshots are publicly accessible.

When you make a snapshot public, you give all AWS accounts and users access to all the data on the snapshot. To share a snapshot only with specific users or accounts, mark the snapshot as private. Then, specify the user or accounts that you want to share the snapshot data with.

Note that if you have Block Public Access enabled in 'block all sharing' mode, then your public snapshots aren't publicly accessible and don't appear in the results of this check.

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

Check ID

ePs02jT06w

Alert Criteria

Red: The EBS volume snapshot is publicly accessible.

Recommended Action

Unless you are certain that you want to share all the data in the snapshot with all AWS accounts and users, modify the permissions: mark the snapshot as private, and then specify the accounts that you want to give permissions to. For more information, see [Sharing an Amazon EBS Snapshot](#). Use Block Public Access for EBS Snapshots to control the settings that allow public access to your data. This check can't be excluded from view in the Trusted Advisor console.

To modify permissions for your snapshots directly, use a runbook in the AWS Systems Manager console. For more information, see [AWSSupport-ModifyEBSSnapshotPermission](#).

Additional Resources

[Amazon EBS Snapshots](#)

Report columns

- Status
- Region
- Volume ID
- Snapshot ID
- Description

Amazon RDS Aurora storage encryption is turned off

Description

Amazon RDS supports encryption at rest for all the database engines by using the keys that you manage in AWS Key Management Service. On an active DB instance with Amazon RDS encryption, the data stored at rest in the storage is encrypted, similar to automated backups, read replicas, and snapshots.

If encryption isn't turned on while creating an Aurora DB cluster, then you must restore a decrypted snapshot to an encrypted DB cluster.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Note

This check relies on Amazon RDS Recommendations. This check doesn't evaluate DB instances in AWS Regions where Amazon RDS Recommendations isn't available. For information about regional availability, see [Viewing and responding to Amazon RDS recommendations](#).

Check ID

c1qf5bt005

Alert Criteria

Red: Amazon RDS Aurora resources don't have encryption enabled.

Recommended Action

Turn on encryption of data at rest for your DB cluster.

Additional Resources

You can turn on encryption while creating a DB instance or use a workaround to turn on the encryption on an active DB instance. You can't modify a decrypted DB cluster to an encrypted DB cluster. However, you can restore a decrypted snapshot to an encrypted DB cluster. When you restore from the decrypted snapshot, you must specify a AWS KMS key.

For more information, see [Encrypting Amazon Aurora resources](#).

Report columns

- Status
- Region
- Resource
- Engine Name
- Last Updated Time

Amazon RDS engine minor version upgrade is required

Description

Your database resources aren't running the latest minor DB engine version. The latest minor version contains the latest security fixes and other improvements.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

 Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

 Note

This check relies on Amazon RDS Recommendations. This check doesn't evaluate DB instances in AWS Regions where Amazon RDS Recommendations isn't available. For information about regional availability, see [Viewing and responding to Amazon RDS recommendations](#).

Check ID

c1qf5bt003

Alert Criteria

Yellow: Amazon RDS resources aren't running the latest minor DB engine version.

Recommended Action

Upgrade to the latest engine version.

Additional Resources

We recommend that you maintain your database with the latest DB engine minor version as this version includes the latest security and functionality fixes. The DB engine minor version

upgrades contain only the changes which are backward-compatible with earlier minor versions of the same major version of the DB engine.

For more information, see [Upgrading a DB instance engine version](#).

Report columns

- Status
- Region
- Resource
- Engine Name
- Engine Version Current
- Recommended Value
- Last Updated Time

Amazon RDS Public Snapshots

Description

Checks the permission settings for your Amazon Relational Database Service (Amazon RDS) DB snapshots and alerts you if any snapshots are marked as public.

When you make a snapshot public, you give all AWS accounts and users access to all the data on the snapshot. If you want to share a snapshot only with specific users or accounts, mark the snapshot as private. Then, specify the user or accounts you want to share the snapshot data with.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

Check ID

rSs93HQwa1

Alert Criteria

Red: The Amazon RDS snapshot is marked as public.

Recommended Action

Unless you are certain you want to share all the data in the snapshot with all AWS accounts and users, modify the permissions: mark the snapshot as private, and then specify the accounts that you want to give permissions to. For more information, see [Sharing a DB Snapshot or DB Cluster Snapshot](#). This check can't be excluded from view in the Trusted Advisor console.

To modify permissions for your snapshots directly, you can use a runbook in the AWS Systems Manager console. For more information, see [AWSsupport-ModifyRDSSnapshotPermission](#).

Additional Resources

[Backing Up and Restoring Amazon RDS DB Instances](#)

Report columns

- Status
- Region
- DB Instance or Cluster ID
- Snapshot ID

Amazon RDS Security Group Access Risk

Description

Checks security group configurations for Amazon Relational Database Service (Amazon RDS) and warns when a security group rule grants overly permissive access to your database. The recommended configuration for a security group rule is to allow access only from specific Amazon Elastic Compute Cloud (Amazon EC2) security groups or from a specific IP address.

Note

This check evaluates only security groups that are attached to Amazon RDS instances running outside on an [Amazon VPC](#).

Check ID

nNauJisYIT

Alert Criteria

- Yellow: A DB security group rule references an Amazon EC2 security group that grants global access on one of these ports: 20, 21, 22, 1433, 1434, 3306, 3389, 4333, 5432, 5500.
- Red: A DB security group rule grants global access (the CIDR rule suffix is /0).
- Green: A DB security group doesn't include permissive rules.

Recommended Action

EC2-Classic was retired on August 15, 2022. It's recommend to move your Amazon RDS instances to a VPC and use Amazon EC2 security groups. For more information of moving your DB instance to a VPC see [Moving a DB instance not in a VPC into a VPC](#).

If you are unable to migrate your Amazon RDS instances to a VPC, then review your security group rules and restrict access to authorized IP addresses or IP ranges. To edit a security group, use the [AuthorizeDBSecurityGroupIngress](#) API or the AWS Management Console. For more information, see [Working with DB Security Groups](#).

Additional Resources

- [Amazon RDS Security Groups](#)
- [Classless Inter-Domain Routing](#)
- [List of TCP and UDP port numbers](#)

Report columns

- Status
- Region
- RDS Security Group Name
- Ingress Rule
- Reason

Amazon RDS storage encryption is turned off

Description

Amazon RDS supports encryption at rest for all the database engines by using the keys that you manage in AWS Key Management Service. On an active DB instance with Amazon RDS encryption, the data stored at rest in the storage is encrypted, similar to automated backups, read replicas, and snapshots.

If encryption isn't turned on while creating a DB instance, then you must restore an encrypted copy of the decrypted snapshot before you turn on the encryption.

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

 Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

 Note

This check relies on Amazon RDS Recommendations. This check doesn't evaluate DB instances in AWS Regions where Amazon RDS Recommendations isn't available. For information about regional availability, see [Viewing and responding to Amazon RDS recommendations](#).

Check ID

c1qf5bt006

Alert Criteria

Red: Amazon RDS resources don't have encryption enabled.

Recommended Action

Turn on encryption of data at rest for your DB instance.

Additional Resources

You can encrypt a DB instance only when you create the DB instance. To encrypt an existing active DB instance:

Create an encrypted copy of the original DB instance

1. Create a snapshot of your DB instance.
2. Create an encrypted copy of the snapshot created in step 1.
3. Restore a DB instance from the encrypted snapshot.

For more information, see the following resources:

- [Encrypting Amazon RDS resources](#)
- [Copying a DB snapshot](#)

Report columns

- Status
- Region
- Resource
- Engine Name
- Last Updated Time

Amazon Route 53 mismatching CNAME records pointing directly to S3 buckets

Description

Checks the Amazon Route 53 Hosted Zones with CNAME records pointing directly to Amazon S3 bucket hostnames and alerts if your CNAME does not match with your S3 bucket name.

Check ID

c1ng44jvbm

Alert Criteria

Red: Amazon Route 53 Hosted Zone has CNAME records pointing to mismatching S3 bucket hostnames.

Green: No mismatching CNAME records found in your Amazon Route 53 Hosted Zone.

Recommended Action

When pointing CNAME records to S3 bucket hostnames, you must make sure that a matching bucket exists for any CNAME or alias record you configure. By doing this, you avoid the risk of your CNAME records being spoofed. You also prevent any unauthorized AWS user from hosting faulty or malicious web content with your domain.

To avoid pointing CNAME records directly to S3 bucket hostnames, consider using origin access control (OAC) to access your S3 bucket web assets through Amazon CloudFront.

For more information about associating CNAME with an Amazon S3 bucket hostname, see [Customizing Amazon S3 URLs with CNAME records](#).

Additional Resources

- [How to associate a hostname with an Amazon S3 bucket](#)
- [Restricting access to an Amazon S3 origin with CloudFront](#)

Report columns

- Status
- Hosted Zone ID
- Hosted Zone ARN
- Matching CNAME Records
- Mismatching CNAME Records
- Last Updated Time

Amazon Route 53 MX Resource Record Sets and Sender Policy Framework

Description

For each MX record, checks for an associated TXT record that contains a valid SPF value. The TXT record value must start with "v=spf1". SPF record types are deprecated by the Internet Engineering Task Force (IETF). With Route 53, it's a best practice to use a TXT record instead of

an SPF record. Trusted Advisor reports this check as green when an MX record has at least one associated TXT record with a valid SPF value.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

c9D319e7sG

Alert Criteria

- Green: An MX resource record set has a TXT resource record that contains a valid SPF value.
- Yellow: An MX resource record set has a TXT or SPF resource record that contains a valid SPF value.
- Red: An MX resource record set doesn't have a TXT or SPF resource record that contains a valid SPF value.

Recommended Action

For each MX resource record set, create a TXT resource record set that contains a valid SPF value. For more information, see [Sender Policy Framework: SPF Record Syntax](#) and [Creating Resource Record Sets By Using the Amazon Route 53 Console](#).

Additional Resources

- [MX record type](#)
- [SPF record type](#)
- [re:Post Guidance](#)
- [RFC 7208](#)

Report columns

- Hosted Zone Name
- Hosted Zone ID
- Resource Record Set Name
- Status

Amazon S3 Bucket Permissions

Description

Checks buckets in Amazon Simple Storage Service (Amazon S3) that have open access permissions, or that allow access to any authenticated AWS user.

This check examines explicit bucket permissions, as well as bucket policies that might override those permissions. Granting list access permissions to all users for an Amazon S3 bucket is not recommended. These permissions can lead to unintended users listing objects in the bucket at high frequency, which can result in higher than expected charges. Permissions that grant upload and delete access to everyone can lead to security vulnerabilities in your bucket.

Check ID

Pfx0RwqBli

Alert criteria

- Red: The bucket ACL allows List access or Upload/Delete access for **Everyone** or **Any Authenticated AWS User** and **Block Public Access** settings are not enabled.
- Red: A bucket policy allows public access and **Block Public Access** settings are not enabled.
- Red: Trusted Advisor does not have permission to check the policy, or the policy could not be evaluated for other reasons.
- Yellow: A bucket policy allows public access, but the **Restrict Public Buckets** setting is turned on and restricts access to only authorized users of that account.
- Yellow: The bucket is compliant but does not have full **Block Public Access** protection enabled.
- Green: The bucket is compliant and has full **Block Public Access** protection enabled.

Note

Public ACL grants are not evaluated when Block Public Access **Ignore Public ACLs** is enabled.

Recommended action

If a bucket allows open access, determine if open access is truly needed. For example to host a static website, you can use Amazon CloudFront to serve the content hosted on Amazon S3. See

[Restricting access to an Amazon S3 origin](#) in the Amazon CloudFront Developer Guide. When possible,, update the bucket permissions to restrict access to the owner or specific users. Use Amazon S3 Block Public Access to control the settings that allow public access to your data. See [Setting Bucket and Object Access Permissions](#).

Additional resources

[Managing Access Permissions to Your Amazon S3 Resources](#)

[Configuring block public access settings for your Amazon S3 buckets](#)

Report columns

- Status
- Region Name
- Region API Parameter
- Bucket Name
- ACL Allows List
- ACL Allows Upload/Delete
- Policy Allows Access

Amazon VPC Peering Connections with DNS Resolution Disabled

Description

Checks if your VPC peering connections have DNS resolution turned on for both the acceptor and requester VPCs.

DNS resolution for a VPC peering connection allows the resolution of public DNS hostnames to private IPv4 addresses when queried from your VPC. This allows the use of DNS names for communication between resources in peered VPCs. DNS resolution in your VPC peering connections makes application development and management simpler and less error-prone, and it ensures that resources always communicate privately over the VPC peering connection.

You can specify the VPC IDs, using the **vpclids** parameters in your AWS Config rules.

For more information, see [Enable DNS resolution for a VPC peering connection](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz124

Source

AWS Config Managed Rule: vpc-peering-dns-resolution-check

Alert Criteria

Yellow: DNS resolution is not enabled for both the acceptor and the requestor VPCs in a VPC peering connection.

Recommended Action

Turn on DNS resolution for your VPC peering connections.

Additional Resources

- [Modify VPC peering connection options](#)
- [DNS attributes in your VPC](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Application Load Balancer Target Groups Encrypted Protocol

Description

Checks Application Load Balancer (ALB) target groups are using HTTPS protocol to encrypt communication in transit for back-end target types of instance or IP. HTTPS requests between ALB and back-end targets help to maintain data confidentiality for data in transit.

Check ID

c2v1fg0p1w

Alert Criteria

- Yellow: Application Load Balancer target group using HTTP.
- Green: Application Load Balancer target group using HTTPS.

Recommended Action

Configure back-end target types of instance or IP to support HTTPS access, and change target group to use HTTPS protocol to encrypt communication between ALB and back-end target types of instance or IP.

Additional Resources

[Enforce encryption in transit](#)

[Application Load Balancer Target Types](#)

[Application Load Balancer Routing Configuration](#)

[Data Protection in Elastic Load Balancing](#)

Report columns

- Status
- Region
- ALB Arn
- ALB Name
- ALB VPC Id
- Target Group Arn
- Target Group Name
- Target Group Protocol
- Last Updated Time

AWS Backup Vault Without Resource-based Policy to Prevent Deletion of Recovery Points

Description

Checks if AWS Backup vaults have an attached resource-based policy that prevents recovery point deletion.

The resource-based policy prevents unexpected deletion of recovery points, which allows you to enforce access control with least privileges against your backup data.

You can specify the AWS Identity and Access Management ARNs that you don't want the rule to check in the **principalArnList** parameter of your AWS Config rules.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz152

Source

AWS Config Managed Rule: backup-recovery-point-manual-deletion-disabled

Alert Criteria

Yellow: There are AWS Backup vaults that don't have a resource-based policy to prevent deletion of recovery points.

Recommended Action

Create resource-based policies for your AWS Backup vaults to prevent unexpected deletion of recovery points.

The policy must include a "Deny" statement with `backup:DeleteRecoveryPoint`, `backup:UpdateRecoveryPointLifecycle`, and `backup:PutBackupVaultAccessPolicy` permissions.

For more information, see [Set access policies on backup vaults](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

AWS CloudTrail Management Event Logging

Description

Checks your use of AWS CloudTrail. CloudTrail provides increased visibility into activity in your AWS account. It does this by recording information about AWS API calls that are made on the account. You can use these logs to determine, for example, what actions a particular user has taken during a specified time period, or which users have taken actions on a particular resource during a specified time period.

Because CloudTrail delivers log files to an Amazon Simple Storage Service (Amazon S3) bucket, CloudTrail must have write permissions for the bucket. If a trail applies to all AWS Regions (the default when creating a new trail), then the trail appears multiple times in the Trusted Advisor report.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c25hn9x03v

Alert Criteria

- Red: No trail is created for an AWS Region, or logging isn't enabled for any trail.
- Yellow: CloudTrail is enabled but all trails report log delivery errors.
- Green: CloudTrail is enabled and no log delivery errors are reported.

Recommended Action

To create a trail and start logging from the console, open the [AWS CloudTrail console](#).

To start logging, see [Stopping and Starting Logging for a Trail](#).

If you receive log delivery errors, then make sure that the bucket exists and that the necessary policy is attached to the bucket. See [Amazon S3 Bucket Policy](#).

Additional Resources

- [AWS CloudTrail User Guide](#)
- [Supported Regions](#)
- [Supported Services](#)
- [Creating a trail for an organization](#)

Report columns

- Status
- Region
- Logging Enabled
- Delivery Error Reported
- Last Updated Time

AWS Lambda Functions Using Deprecated Runtimes

Description

Checks for Lambda functions whose \$LATEST version is configured to use a runtime that is approaching deprecation, or is deprecated. Deprecated runtimes are not eligible for security updates or technical support

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Published Lambda function versions are immutable, which means they can be invoked but not updated. Only the \$LATEST version for a Lambda function can be updated. For more information, see [Lambda function versions](#).

Check ID

L4dfs2Q4C5

Alert Criteria

- Red: The function's \$LATEST version is configured to use a runtime that is already deprecated.
- Yellow: The function's \$LATEST version is running on a runtime that is approaching deprecation. Functions are included at least 180 days before the runtime deprecation date.

Recommended Action

If you have functions that are running on a runtime that is approaching deprecation, you should prepare for migration to a supported runtime. For more information, see [Runtime support policy](#).

We recommend that you delete earlier function versions that you're no longer using.

Additional Resources

[Lambda runtimes](#)

Report columns

- Status
- Region
- Function ARN
- Runtime
- Days to Deprecation
- Deprecation Date
- Average Daily Invokes
- Last Updated Time

AWS Well-Architected high risk issues for security

Description

Checks for high risk issues (HRIs) for your workloads in the security pillar. This check is based on your AWS-Well Architected reviews. Your check results depend on whether you completed the workload evaluation with AWS Well-Architected.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

Wxdfp4B1L3

Alert Criteria

- Red: At least one active high risk issue was identified in the security pillar for AWS Well-Architected.
- Green: No active high risk issues were detected in the security pillar for AWS Well-Architected.

Recommended Action

AWS Well-Architected detected high risk issues during your workload evaluation. These issues present opportunities to reduce risk and save money. Sign in to the [AWS Well-Architected](#) tool to review your answers and take action to resolve your active issues.

Report columns

- Status
- Region
- Workload ARN
- Workload Name
- Reviewer Name
- Workload Type

- Workload Started Date
- Workload Last Modified Date
- Number of identified HRIs for Security
- Number of HRIs resolved for Security
- Number of questions for Security
- Total number of questions in Security pillar
- Last Updated Time

CloudFront Custom SSL Certificates in the IAM Certificate Store

Description

Note

This check applies to classic Amazon CloudFront distributions.

Checks the SSL certificates for CloudFront alternate domain names in the IAM certificate store. This check alerts you if a certificate is expired, will expire soon, uses outdated encryption, or is not configured correctly for the distribution.

When a custom certificate for an alternate domain name expires, browsers that display your CloudFront content might show a warning message about the security of your website. Certificates that are encrypted by using the SHA-1 hashing algorithm are being deprecated by most web browsers such as Chrome and Firefox.

A certificate must contain a domain name that matches either the Origin Domain Name or the domain name in the host header of a viewer request. If it doesn't match, CloudFront returns an HTTP status code of 502 (bad gateway) to the user. For more information, see [Using Alternate Domain Names and HTTPS](#).

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

N425c450f2

Alert Criteria

- Red: A custom SSL certificate is expired.
- Yellow: A custom SSL certificate expires in the next seven days.
- Yellow: A custom SSL certificate was encrypted by using the SHA-1 hashing algorithm.
- Yellow: One or more of the alternate domain names in the distribution don't appear either in the Common Name field or the Subject Alternative Names field of the custom SSL certificate.

Recommended Action

We recommend using AWS Certificate Manager to provision, manage, and deploy your server certificates. With ACM, you can request a new certificate or deploy an existing ACM or external certificate to AWS resources. Certificates provided by ACM are free and can be automatically renewed. For more information about using ACM, see the [AWS Certificate Manager User Guide](#). To verify the AWS Regions ACM supports, see [AWS Certificate Manager endpoints](#) and quotas in the AWS General Reference.

Renew expired certificates or certificates that are about to expire. For more information on renewing a certificate see [Managing server certificates](#) in IAM.

Replace a certificate that was encrypted by using the SHA-1 hashing algorithm with a certificate that is encrypted by using the SHA-256 hashing algorithm.

Replace the certificate with a certificate that contains the applicable values in the Common Name or Subject Alternative Domain Names fields.

Additional Resources

[Using an HTTPS Connection to Access Your Objects](#)

[Importing Certificates](#)

[AWS Certificate Manager User Guide](#)

Report columns

- Status
- Distribution ID
- Distribution Domain Name
- Certificate Name

- Reason

CloudFront SSL Certificate on the Origin Server

Description

Checks your origin server for SSL certificates that are expired, about to expire, missing, or that use outdated encryption. If a certificate has one of these issues, CloudFront responds to requests for your content with HTTP status code 502, Bad Gateway.

Certificates that were encrypted by using the SHA-1 hashing algorithm are being deprecated by web browsers such as Chrome and Firefox. Depending on the number of SSL certificates that you have associated with your CloudFront distributions, this check might add a few cents per month to your bill with your web hosting provider, for example, AWS if you're using Amazon EC2 or Elastic Load Balancing as the origin for your CloudFront distribution. This check does not validate your origin certificate chain or certificate authorities. You can check these in your CloudFront configuration.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

N430c450f2

Alert Criteria

- Red: An SSL certificate on your origin has expired or is missing.
- Yellow: An SSL certificate on your origin expires in the next thirty days.
- Yellow: An SSL certificate on your origin was encrypted by using the SHA-1 hashing algorithm.
- Yellow: An SSL certificate on your origin can't be located. The connection might have failed due to timeout, or other HTTPS connection problems.

Recommended Action

Renew the certificate on your origin if it has expired or is about to expire.

Add a certificate if one does not exist.

Replace a certificate that was encrypted by using the SHA-1 hashing algorithm with a certificate that is encrypted by using the SHA-256 hashing algorithm.

Additional Resources

[Using Alternate Domain Names and HTTPS](#)

Report columns

- Status
- Distribution ID
- Distribution Domain Name
- Origin
- Reason

ELB Listener Security

Description

Checks for classic load balancers with listeners that don't use the recommended security configurations for encrypted communication. AWS recommends that you use a secure protocol (HTTPS or SSL), up-to-date security policies, and ciphers and protocols that are secure. When you use a secure protocol for a front-end connection (client to load balancer), the requests are encrypted between your clients and the load balancer. This creates a more secure environment. Elastic Load Balancing provides predefined security policies with ciphers and protocols that adhere to AWS security best practices. New versions of predefined policies are released as new configurations become available.

Check ID

a2sEc6ILx

Alert Criteria

- Red: A load balancer has no listeners configured with a secure protocol (HTTPS).
- Yellow: A load balancer HTTPS listener is configured with a Security Policy that contains a weak cipher.
- Yellow: A load balancer HTTPS listener is not configured with the recommended Security Policy.

- Green: A load balancer has at least one HTTPS listener AND all HTTPS listeners are configured with the recommended policy.

Recommended Action

If the traffic to your load balancer must be secure, use either the HTTPS or the SSL protocol for the front-end connection.

Upgrade your load balancer to the latest version of the predefined SSL security policy.

Use only the recommended ciphers and protocols.

For more information, see [Listener Configurations for Elastic Load Balancing](#).

Additional Resources

- [Listener Configurations Quick Reference](#)
- [Update SSL Negotiation Configuration of Your Load Balancer](#)
- [SSL Negotiation Configurations for Elastic Load Balancing](#)
- [SSL Security Policy Table](#)

Report columns

- Status
- Region
- Load Balancer Name
- Load Balancer Port
- Reason

Classic Load Balancer Security Groups

Description

Checks for load balancers configured with a security group that allows access to ports that are not configured for the load balancer.

If a security group allows access to ports that are not configured for the load balancer, the risk of loss of data or malicious attacks increases.

Check ID

xSqX82fQu

Alert Criteria

- Yellow: The inbound rules of an Amazon VPC security group associated with a load balancer allow access to ports that are not defined in the load balancer's listener configuration.
- Green: The inbound rules of an Amazon VPC security group associated with a load balancer do not allow access to ports that are not defined in the load balancers listener configuration.

Recommended Action

Configure the security group rules to restrict access to only those ports and protocols that are defined in the load balancer listener configuration, plus the ICMP protocol to support Path MTU Discovery. See [Listeners for Your Classic Load Balancer](#) and [Security Groups for Load Balancers in a VPC](#).

If a security group is missing, apply a new security group to the load balancer. Create security group rules that restrict access to only those ports and protocols that are defined in the load balancer listener configuration. See [Security Groups for Load Balancers in a VPC](#).

Additional Resources

- [Elastic Load Balancing User Guide](#)
- [Migrate your Classic Load Balancer](#)
- [Configure Your Classic Load Balancer](#)

Report columns

- Status
- Region
- Load Balancer Name
- Security Group IDs
- Reason

Exposed Access Keys

Description

Checks popular code repositories for access keys that have been exposed to the public and for irregular Amazon Elastic Compute Cloud (Amazon EC2) usage that could be the result of a compromised access key.

An access key consists of an access key ID and the corresponding secret access key. Exposed access keys pose a security risk to your account and other users, could lead to excessive charges from unauthorized activity or abuse, and violate the [AWS Customer Agreement](#).

If your access key is exposed, take immediate action to secure your account. To protect your account from excessive charges, AWS temporarily limits your ability to create some AWS resources. This does not make your account secure. It only partially limits the unauthorized usage for which you could be charged.

 Note

This check doesn't guarantee the identification of exposed access keys or compromised EC2 instances. You are ultimately responsible for the safety and security of your access keys and AWS resources.

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

If a deadline is shown for an access key, AWS may suspend your AWS account if the unauthorized usage is not stopped by that date. If you believe an alert is in error, [contact AWS Support](#).

The information displayed in Trusted Advisor might not reflect the most recent state of your account. No exposed access keys are marked as resolved until all exposed access keys on the account have been resolved. This data synchronization can take up to one week.

Check ID

12Fnkp18Y5

Alert Criteria

- Red: Potentially compromised – AWS has identified an access key ID and corresponding secret access key that have been exposed on the Internet and may have been compromised (used).

- **Red: Exposed** – AWS has identified an access key ID and corresponding secret access key that have been exposed on the Internet.
- **Red: Suspected** - Irregular Amazon EC2 usage indicates that an access key may have been compromised, but it has not been identified as exposed on the Internet.

Recommended Action

Delete the affected access key as soon as possible. If the key is associated with an IAM user, see [Managing Access Keys for IAM Users](#).

Check your account for unauthorized usage. Sign in to the [AWS Management Console](#) and check each service console for suspicious resources. Pay special attention to running Amazon EC2 instances, Spot Instance requests, access keys, and IAM users. You can also check overall usage on the [Billing and Cost Management console](#).

Additional Resources

- [Best Practices for Managing AWS Access Keys](#)
- [AWS Security Audit Guidelines](#)

Report columns

- Access Key ID
- User Name (IAM or Root)
- Fraud Type
- Case ID
- Time Updated
- Location
- Deadline
- Usage (USD per Day)

IAM Access Key Rotation

Description

Checks for active IAM access keys that have not been rotated in the last 90 days.

When you rotate your access keys regularly, you reduce the chance that a compromised key could be used without your knowledge to access resources. For the purposes of this check, the last rotation date and time is when the access key was created or most recently

activated. The access key number and date come from the `access_key_1_last_rotated` and `access_key_2_last_rotated` information in the most recent IAM credential report.

Because the regeneration frequency of a credential report is restricted, refreshing this check might not reflect recent changes. For more information, see [Getting Credential Reports for Your AWS account](#).

In order to create and rotate access keys, a user must have the appropriate permissions. For more information, see [Allow Users to Manage Their Own Passwords, Access Keys, and SSH Keys](#).

Check ID

DqdJqYeRm5

Alert Criteria

- Green: The access key is active and has been rotated in the last 90 days.
- Yellow: The access key is active and has been rotated in the last 2 years, but more than 90 days ago.
- Red: The access key is active and has not been rotated in the last 2 years.

Recommended Action

Rotate access keys on a regular basis. See [Rotating Access Keys](#) and [Managing Access Keys for IAM Users](#).

Additional Resources

- [IAM Best Practices](#)
- [How to rotate access keys for IAM users](#)

Report columns

- Status
- IAM user
- Access Key
- Key Last Rotated
- Reason

IAM Access Analyzer External Access

Description

Checks if the IAM Access Analyzer external access at the account level is present.

IAM Access Analyzer external access analyzers help identify resources in your accounts that are shared with an external entity. The analyzer then creates a centralized dashboard with the findings. After the new analyzer is activated in the IAM console, security teams can then prioritize which accounts to review based on excessive permissions. An external access analyzer creates public and cross-account access findings for resources, and is provided at no additional charge.

Check ID

07602fcad6

Alert Criteria

- Red: The analyzer external access isn't activated at the account level.
- Green: The analyzer external access is activated at the account level.

Recommended Action

The creation of an external access analyzer per account helps security teams to prioritize which accounts to review based on excessive permissions. For more information, see [Getting started with AWS Identity and Access Management Access Analyzer findings](#).

Additionally, it's a best practice to utilize the unused access analyzer, a paid feature that simplifies inspecting unused access to guide you toward least privilege. For more information, see [Identifying unused access granted to IAM users and roles](#).

Additional Resources

- [Using AWS Identity and Access Management Access Analyzer](#)
- [IAM Access Analyzer updates: Find unused access, check policies before deployment](#)

Report columns

- Status
- Region
- Account External Access Analyzer Arn
- Organization External Access Analyzer Arns
- Last Updated Time

IAM Password Policy

Description

Checks the password policy for your account and warns when a password policy is not enabled, or if password content requirements have not been enabled.

Password content requirements increase the overall security of your AWS environment by enforcing the creation of strong user passwords. When you create or change a password policy, the change is enforced immediately for new users but does not require existing users to change their passwords.

Check ID

Yw2K9puPz1

Alert Criteria

- Green: A password policy is enabled with recommended content requirement enabled.
- Yellow: A password policy is enabled, but at least one content requirement is not enabled.

Recommended Action

If some content requirements are not enabled, consider enabling them. If no password policy is enabled, create and configure one. See [Setting an Account Password Policy for IAM Users](#).

To access the AWS Management Console, IAM users need passwords. As a best practice, AWS highly recommends that instead of creating IAM users, you use federation. Federation allows users to use their existing corporate credentials to log into the AWS Management Console. Use IAM Identity Center to create or federate the user, and then assume an IAM role into an account.

To learn more about identity providers and federation, see [Identity providers and federation](#) in the IAM User Guide. To learn more about IAM Identity Center, see the [IAM Identity Center User Guide](#).

Additional Resources

[Managing Passwords](#)

Report columns

- Password Policy
- Uppercase

- Lowercase
- Number
- Non-alphanumeric

IAM SAML 2.0 Identity Provider

Description

Checks if the AWS account is configured for access via an identity provider (IdP) that supports SAML 2.0. Be sure to follow best practices when you centralize identities and configure users in an [external identity provider](#) or [AWS IAM Identity Center](#).

Check ID

c2v1fg0p86

Alert Criteria

- Yellow: This account isn't configured for access via an identity provider (IdP) that supports SAML 2.0.
- Green: This account is configured for access via an identity provider (IdP) that supports SAML 2.0.

Recommended Action

Activate IAM Identity Center for the AWS account. For more information, see [Enabling IAM Identity Center](#). After you turn on IAM Identity Center, you can then perform common tasks like creating a permission set and assigning access for Identity Center groups. For more information, see [Common tasks](#).

It's a best practice to manage human users in IAM Identity Center. But you can activate federated user access with IAM for human users in the short-term for small scale deployments. For more information see [SAML 2.0 federation](#).

Additional Resources

[What is IAM Identity Center?](#)

[What Is IAM?](#)

Report columns

- Status

- AWS account Id
- Last Updated Time

MFA on root account

Description

Checks the root user credentials of an account and warns if multi-factor authentication (MFA) is not enabled.

For increased security, we recommend that you protect your account by using MFA, which requires a user to enter a unique authentication code from their MFA hardware or virtual device when interacting with the AWS Management Console and associated websites.

Note

For your AWS Organizations management account, AWS requires multi-factor authentication (MFA) for the root user when accessing the AWS Management Console. For your AWS Organizations member accounts, we recommend that you centrally manage root credentials using AWS Identity and Access Management. Member account root user credentials can be deleted centrally, removing the need to manage MFA on root user credentials. For more information, see [Best practices for member accounts](#) in the *AWS Organizations User Guide*.

Check ID

7DAFEemoDos

Alert criteria

- Red: MFA is not enabled on the root account.
- Green: No root user credentials (root password) exist or MFA is enabled for the account.

Recommended action

If this is a member account in AWS Organizations: Log in to your management account, enable the root access management feature in IAM, and remove your root user credentials from this member account. See [Centralize root access for member accounts](#).

If this is a standalone or management account in AWS Organizations: Log in to your root account and activate an MFA device. For more information, see [Check MFA status](#) and [AWS Multi-factor authentication in IAM](#)

Additional resources

- [Centrally manage root access for member accounts](#)
- [AWS Multi-factor authentication in IAM](#)
- [Multi-factor authentication for AWS account root user](#)

Root User Access Key

Description

Checks if the root user access key is present. It's strongly recommended that you don't create access key pairs for your root user. Because [only a few tasks require the root user](#) and you typically perform those tasks infrequently, it's a best practice to log in to the AWS Management Console to perform the root user tasks. Before you create access keys, review the [alternatives to long-term access keys](#).

Check ID

c2v1fg0f4h

Alert Criteria

- Red: The root user access key is present
- Green: The root user access key isn't present

Recommended Action

Delete the access key(s) for the root user. See [Deleting access keys for the root user](#). This task must be performed by the root user. You can't perform these steps as an IAM user or role.

Additional Resources

- [Tasks that require root user credentials](#)
- [Resetting a lost or forgotten root user password](#)

Report columns

- Status
- Account ID

- Last Updated Time

Security Groups – Specific Ports Unrestricted

Description

Checks security groups for rules that allow unrestricted access (0.0.0.0/0) to specific ports.

Unrestricted access increases opportunities for malicious activity (hacking, denial-of-service attacks, loss of data). The ports with highest risk are flagged red, and those with less risk are flagged yellow. Ports flagged green are typically used by applications that require unrestricted access, such as HTTP and SMTP.

If you have intentionally configured your security groups in this manner, we recommend using additional security measures to secure your infrastructure (such as IP tables).

Note

This check only evaluates security groups that you create and their inbound rules for IPv4 addresses. Security groups created by AWS Directory Service are flagged as red or yellow, but they don't pose a security risk and can be excluded. For more information, see the [Trusted Advisor FAQ](#).

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

HCP4007jGY

Alert Criteria

- Green: Security Group provides unrestricted access on ports 80, 25, 443, or 465.
- Red: Security Group provides unrestricted access to port 20, 21, 22, 1433, 1434, 3306, 3389, 4333, 5432, or 5500.
- Yellow: Security Group provides unrestricted access to any other port.

Recommended Action

Restrict access to only those IP addresses that require it. To restrict access to a specific IP address, set the suffix to /32 (for example, 192.0.2.10/32). Be sure to delete overly permissive rules after creating rules that are more restrictive.

Review and delete unused security groups. You can use AWS Firewall Manager to centrally configure and manage security groups at scale across AWS accounts. For more information, see the [AWS Firewall Manager documentation](#).

Consider using Systems Manager Sessions Manager for SSH (Port 22) and RDP (Port 3389) access to EC2 instances. With sessions manager, you can access your EC2 instances without enabling port 22 and 3389 in the security group.

Additional Resources

- [Amazon EC2 Security Groups](#)
- [List of TCP and UDP port numbers](#)
- [Classless Inter-Domain Routing](#)
- [Working with Session Manager](#)
- [AWS Firewall Manager](#)

Report columns

- Status
- Region
- Security Group Name
- Security Group ID
- Protocol
- From Port
- To Port
- Association

Security Groups – Unrestricted Access

Description

Checks security groups for rules that allow unrestricted access to a resource.

Unrestricted access increases opportunities for malicious activity (hacking, denial-of-service attacks, loss of data).

 Note

This check evaluates only security groups that you create and their inbound rules for IPv4 addresses. Security groups created by AWS Directory Service are flagged as red or yellow, but they don't pose a security risk and can be excluded. For more information, see the [Trusted Advisor FAQ](#).

 Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

1iG5NDGVre

Alert Criteria

- Green: A security group rule has a source IP address with a /0 suffix for ports 25, 80, or 443.
- Yellow: A security group rule has a source IP address with a /0 suffix for ports other than 25, 80, or 443 and security group is attached to a resource.
- Red: A security group rule has a source IP address with a /0 suffix for ports other than 25, 80, or 443 and security group is not attached to a resource.

Recommended Action

Restrict access to only those IP addresses that require it. To restrict access to a specific IP address, set the suffix to /32 (for example, 192.0.2.10/32). Be sure to delete overly permissive rules after creating rules that are more restrictive.

Review and delete unused security groups. You can use AWS Firewall Manager to centrally configure and manage security groups at scale across AWS accounts. For more information, see the [AWS Firewall Manager documentation](#).

Consider using Systems Manager Sessions Manager for SSH (Port 22) and RDP (Port 3389) access to EC2 instances. With sessions manager, you can access your EC2 instances without enabling port 22 and 3389 in the security group.

Additional Resources

- [Amazon EC2 Security Groups](#)
- [Classless Inter-Domain Routing](#)
- [Working with Session Manager](#)
- [AWS Firewall Manager](#)

Report columns

- Status
- Region
- Security Group Name
- Security Group ID
- Protocol
- From Port
- To Port
- IP Range
- Association

Fault tolerance

You can use the following checks for the fault tolerance category.

Check names

- [ALB Multi-AZ](#)
- [Amazon Aurora MySQL cluster backtracking not enabled](#)
- [Amazon Aurora DB Instance Accessibility](#)
- [Amazon CloudFront Origin Failover](#)
- [Amazon Comprehend Endpoint Access Risk](#)
- [Amazon DocumentDB Single AZ Clusters](#)

- [Amazon DynamoDB Point-in-time Recovery](#)
- [Amazon DynamoDB Table Not Included in Backup Plan](#)
- [Amazon EBS Not Included in AWS Backup Plan](#)
- [Amazon EBS Snapshots](#)
- [Amazon EC2 Auto Scaling does not have ELB Health Check Enabled](#)
- [Amazon EC2 Auto Scaling Group has Capacity Rebalancing Enabled](#)
- [Amazon EC2 Auto Scaling is not deployed in multiple AZs or does not meet the minimum number of AZs](#)
- [Amazon EC2 Availability Zone Balance](#)
- [Amazon EC2 detailed monitoring not enabled](#)
- [Amazon ECS AWSLogs driver in blocking mode](#)
- [Amazon ECS service using a single AZ](#)
- [Amazon ECS Multi-AZ placement strategy](#)
- [Amazon EFS no mount target redundancy](#)
- [Amazon EFS not in AWS Backup plan](#)
- [Amazon ElastiCache Multi-AZ clusters](#)
- [ElastiCache \(Redis OSS\) Clusters Automatic Backup](#)
- [Amazon MemoryDB Multi-AZ clusters](#)
- [Amazon MSK brokers hosting too many partitions](#)
- [Amazon MSK Cluster Multi-AZ](#)
- [Amazon OpenSearch Service domains with less than three data nodes](#)
- [Amazon RDS Backups](#)
- [Amazon RDS Continuous Backup Not Enabled](#)
- [Amazon RDS DB clusters have one DB instance](#)
- [Amazon RDS DB clusters with all instances in the same Availability Zone](#)
- [Amazon RDS DB clusters with all reader instances in the same Availability Zone](#)
- [Amazon RDS DB Instance Enhanced Monitoring not enabled](#)
- [Amazon RDS DB instances have storage autoscaling turned off](#)
- [Amazon RDS DB instances not using Multi-AZ deployment](#)

- [Amazon RDS DiskQueueDepth](#)
- [Amazon RDS FreeStorageSpace](#)
- [Amazon RDS log_output parameter is set to table](#)
- [Amazon RDS innodb_default_row_format parameter setting is unsafe](#)
- [Amazon RDS innodb_flush_log_at_trx_commit parameter is not 1](#)
- [Amazon RDS max_user_connections parameter is low](#)
- [Amazon RDS Multi-AZ](#)
- [Amazon RDS Not In AWS Backup Plan](#)
- [Amazon RDS Read Replicas are open in writable mode](#)
- [Amazon RDS resource automated backups is turned off](#)
- [Amazon RDS sync_binlog parameter is turned off](#)
- [RDS DB Cluster has no Multi-AZ replication enabled](#)
- [RDS Multi-AZ Standby Instance Not Enabled](#)
- [Amazon RDS ReplicaLag](#)
- [Amazon RDS synchronous_commit parameter is turned off](#)
- [Amazon Redshift cluster automated snapshots](#)
- [Amazon Route 53 Deleted Health Checks](#)
- [Amazon Route 53 Failover Resource Record Sets](#)
- [Amazon Route 53 High TTL Resource Record Sets](#)
- [Amazon Route 53 Name Server Delegations](#)
- [Amazon Route 53 Resolver Endpoint Availability Zone Redundancy](#)
- [Amazon S3 Bucket Replication Not Enabled](#)
- [Amazon S3 Bucket Versioning](#)
- [Application, Network, and Gateway Load Balancers Not Spanning Multiple Availability Zones](#)
- [Auto Scaling available IPs in Subnets](#)
- [Auto Scaling Group Health Check](#)
- [Auto Scaling Group Resources](#)
- [AWS CloudHSM clusters running HSM instances in a single AZ](#)

- [Direct Connect Location Resiliency](#)
- [AWS Lambda functions without a dead-letter queue configured](#)
- [AWS Lambda On Failure Event Destinations](#)
- [AWS Lambda VPC-enabled Functions without Multi-AZ Redundancy](#)
- [AWS Outposts Single Rack deployment](#)
- [AWS Resilience Hub Application Component check](#)
- [AWS Resilience Hub policy breached](#)
- [AWS Resilience Hub resilience scores](#)
- [AWS Resilience Hub assessment age](#)
- [AWS Site-to-Site VPN has at least one tunnel in DOWN status](#)
- [AWS STS global endpoint usage across AWS Regions](#)
- [AWS Well-Architected high risk issues for reliability](#)
- [Classic Load Balancer has no multiple AZs configured](#)
- [CLB Connection Draining](#)
- [ELB Target Imbalance](#)
- [GWLB - endpoint AZ independence](#)
- [Load Balancer Optimization](#)
- [NAT Gateway AZ Independence](#)
- [Network Firewall endpoint AZ Independence](#)
- [Network Firewall Multi-AZ](#)
- [Network Load Balancers Cross Load Balancing](#)
- [NLB - Internet-facing resource in private subnet](#)
- [NLB Multi-AZ](#)
- [Number of AWS Regions in an Incident Manager replication set](#)
- [Single AZ Application Check](#)
- [VPC interface endpoint network interfaces in multiple AZs](#)
- [VPN Tunnel Redundancy](#)
- [ActiveMQ Availability Zone Redundancy](#)
- [RabbitMQ Availability Zone Redundancy](#)

ALB Multi-AZ

Description

Checks if your Application Load Balancers are configured to use more than one Availability Zone (AZ). An AZ is a distinct location that is insulated from failures in other zones. Configure your load balancer in multiple AZs in the same Region to help improve your workload availability.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1dfprch08

Alert Criteria

Yellow: ALB is in a single AZ.

Green: ALB has two or more AZs.

Recommended Action

Make sure that your load balancer is configured with at least two Availability Zones.

For more information, see [Availability Zones for your Application Load Balancer](#).

Additional Resources

For more information, see the following documentation:

- [How Elastic Load Balancing works](#)
- [Regions, Availability Zones, and Local Zones](#)

Report columns

- Status
- Region
- ALB Name

- ALB Rule
- ALB ARN
- Number of AZs
- Last Updated Time

Amazon Aurora MySQL cluster backtracking not enabled

Description

Checks if an Amazon Aurora MySQL cluster has backtracking enabled.

Amazon Aurora MySQL cluster backtracking is a feature that allows you to restore an Aurora DB cluster to a previous point in time without creating a new cluster. It enables you to roll back your database to a specific point in time within a retention period, without the need to restore from a snapshot.

You can adjust the backtracking time window (hours) in the **BacktrackWindowInHours** parameter of the AWS Config rules.

For more information, see [Backtracking an Aurora DB cluster](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz131

Source

AWS Config Managed Rule: aurora-mysql-backtracking-enabled

Alert Criteria

Yellow: Amazon Aurora MySQL clusters backtracking is not enabled.

Recommended Action

Turn on backtracking for your Amazon Aurora MySQL cluster.

For more information, see [Backtracking an Aurora DB cluster](#).

Additional Resources

[Backtracking an Aurora DB cluster](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon Aurora DB Instance Accessibility

Description

Checks for cases where an Amazon Aurora DB cluster has both private and public instances.

When your primary instance fails, a replica can be promoted to a primary instance. If that replica is private, users who have only public access would no longer be able to connect to the database after failover. We recommend that all the DB instances in a cluster have the same accessibility.

Check ID

xuy7H1avt1

Alert Criteria

Yellow: The instances in an Aurora DB cluster have different accessibility (a mix of public and private).

Recommended Action

Modify the `Publicly Accessible` setting of the instances in the DB cluster so that they are all either public or private. For details, see the instructions for MySQL instances at [Modifying a DB Instance Running the MySQL Database Engine](#).

Additional Resources

[Fault Tolerance for an Aurora DB Cluster](#)

Report columns

- Status
- Region
- Cluster
- Public DB Instances
- Private DB Instances
- Reason

Amazon CloudFront Origin Failover

Description

Checks that an origin group is configured for distributions that include two origins in Amazon CloudFront.

For more information, see [Optimizing high availability with CloudFront origin failover](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz112

Source

AWS Config Managed Rule: `cloudfront-origin-failover-enabled`

Alert Criteria

Yellow: Amazon CloudFront origin failover is not enabled.

Recommended Action

Make sure that you turn on the origin failover feature for your CloudFront distributions to help ensure high availability of your content delivery to end users. When you turn on this feature, traffic is automatically routed to the backup origin server if the primary origin server is unavailable. This minimizes potential downtime and ensures continuous availability of your content.

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon Comprehend Endpoint Access Risk

Description

Checks the AWS Key Management Service (AWS KMS) key permissions for an endpoint where the underlying model was encrypted by using customer managed keys. If the customer managed key is disabled, or the key policy was changed to alter the allowed permissions for Amazon Comprehend, the endpoint availability might be affected.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

Cm24dfsM13

Alert Criteria

Red: The customer managed key is disabled or the key policy was changed to alter the allowed permissions for Amazon Comprehend access.

Recommended Action

If the customer managed key was disabled, we recommend that you enable it. For more information, see [Enabling keys](#). If the key policy was altered and you want to keep using the endpoint, we recommend that you update the AWS KMS key policy. For more information, see [Changing a key policy](#).

Additional Resources

[AWS KMS Permissions](#)

Report columns

- Status
- Region
- Endpoint ARN
- Model ARN
- KMS KeyId
- Last Updated Time

Amazon DocumentDB Single AZ Clusters

Description

Checks if there are Amazon DocumentDB clusters configured as Single-AZ.

Running Amazon DocumentDB workloads in a Single-AZ architecture is not sufficient for highly critical workloads and it can take up to 10 minutes to recover from a component failure. Customers should deploy replica instances in additional availability zones to ensure availability during maintenance, instance failures, component failures, or availability zone failures.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c15vnddn2x

Alert Criteria

Yellow: Amazon DocumentDB cluster has instances in less than three availability zones.

Green: Amazon DocumentDB cluster has instances in three availability zones.

Recommended Action

If your application requires high availability, modify your DB instance to enable Multi-AZ using replica instances. See [Amazon DocumentDB High Availability and Replication](#)

Additional Resources

[Understanding Amazon DocumentDB Cluster Fault Tolerance](#)

[Regions and Availability Zones](#)

Report columns

- Status
- Region
- Availability Zone
- DB Cluster Identifier
- DB Cluster ARN
- Last Updated Time

Amazon DynamoDB Point-in-time Recovery

Description

Checks if point-in time-recovery is enabled for your Amazon DynamoDB tables.

Point-in time-recovery helps protect your DynamoDB tables from accidental write or delete operations. With point-in time-recovery, you don't have to worry about creating, maintaining,

or scheduling on-demand backups. Point-in time-recovery restores tables to any point in time during the last 35 days. DynamoDB maintains incremental backups of your table.

For more information, see [Point-in-time recovery for DynamoDB](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz138

Source

AWS Config Managed Rule: dynamodb-pitr-enabled

Alert Criteria

Yellow: Point-in-time recovery is not enabled for your DynamoDB tables.

Recommended Action

Turn on point-in-time recovery in Amazon DynamoDB to continuously back up your table data.

For more information, see [Point-in-time recovery: How it works](#).

Additional Resources

[Point-in-time recovery for DynamoDB](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon DynamoDB Table Not Included in Backup Plan

Description

Checks if Amazon DynamoDB tables are part of an AWS Backup plan.

AWS Backup provides incremental backups for DynamoDB tables that capture changes made since the last backup. Including DynamoDB tables in an AWS Backup plan helps protect your data from accidental data loss scenarios and automates the backup process. This provides a reliable and scalable backup solution for your DynamoDB tables, helping to ensure that your valuable data is protected and available for recovery as needed.

For more information, see [Creating backups of DynamoDB tables with AWS Backup](#)

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz107

Source

AWS Config Managed Rule: dynamodb-in-backup-plan

Alert Criteria

Yellow: Amazon DynamoDB table is not included in AWS Backup plan.

Recommended Action

Ensure that your Amazon DynamoDB tables are part of an AWS Backup plan.

Additional Resources

[Scheduled backups](#)

[What is AWS Backup?](#)

[Creating backup plans using the AWS Backup console](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon EBS Not Included in AWS Backup Plan

Description

Checks if Amazon EBS volumes are present in backup plans for AWS Backup.

Include Amazon EBS volumes in an AWS Backup plan to automate regular backups for the data stored on those volumes. This protects you against data loss, makes data management easier, and allows for data restoration when needed. A backup plan helps to ensure that your data is safe and that you're able to meet recovery time and point objectives (RTO/RPO) for your application and services.

For more information, see [Creating a backup plan](#)

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz106

Source

AWS Config Managed Rule: ebs-in-backup-plan

Alert Criteria

Yellow: Amazon EBS volume is not included in AWS Backup plan.

Recommended Action

Make sure that your Amazon EBS volumes are part of an AWS Backup plan.

Additional Resources

[Creating backup plans using the AWS Backup console](#)

[What is AWS Backup?](#)

[Getting started 3: Create a scheduled backup](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon EBS Snapshots

Description

Checks the age of the snapshots for your Amazon EBS volumes (either available or in-use). Failures can occur even if Amazon EBS volumes are replicated. Snapshots are persisted to Amazon S3 for durable storage and point-in-time recovery.

Check ID

H7IgTzjTYb

Alert Criteria

- Yellow: The most recent volume snapshot is between 7 and 30 days old.

- Red: The most recent volume snapshot is more than 30 days old.
- Red: The volume does not have a snapshot.

Recommended Action

Create weekly or monthly snapshots of your volumes. For more information, see [Creating an Amazon EBS Snapshot](#).

To automate the creation of EBS snapshots, you can consider using [AWS Backup](#) or [Amazon Data Lifecycle Manager](#).

Additional Resources

[Amazon Elastic Block Store \(Amazon EBS\)](#)

[Amazon EBS Snapshots](#)

[AWS Backup](#)

[Amazon Data Lifecycle Manager](#)

Report columns

- Status
- Region
- Volume ID
- Volume Name
- Snapshot ID
- Snapshot Name
- Snapshot Age
- Volume Attachment
- Reason

Amazon EC2 Auto Scaling does not have ELB Health Check Enabled

Description

Checks if your Amazon EC2 Auto Scaling groups that are associated with a Classic Load Balancer are using Elastic Load Balancing health checks. The default health checks for an Auto Scaling

group are Amazon EC2 status checks only. If an instance fails these status checks, it is marked unhealthy and is terminated. Amazon EC2 Auto Scaling launches a new replacement instance. The Elastic Load Balancing health check periodically monitors Amazon EC2 instances to detect and terminate unhealthy instances and then launch new instances.

For more information, see [Add Elastic Load Balancing health checks](#).

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz104

Source

AWS Config Managed Rule: autoscaling-group-elb-healthcheck-required

Alert Criteria

Yellow: Amazon EC2 Auto Scaling group attached to Classic Load Balancer has not enabled Elastic Load Balancing health checks.

Recommended Action

Ensure that your Auto Scaling groups that are associated with a Classic Load Balancer use Elastic Load Balancing health checks.

Elastic Load Balancing health checks report if the load balancer is healthy and available to handle requests. This ensures high availability for your application.

For more information, see [Add Elastic Load Balancing health checks to an Auto Scaling group](#)

Report columns

- Status
- Region

- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon EC2 Auto Scaling Group has Capacity Rebalancing Enabled

Description

Checks if Capacity Rebalancing is enabled for Amazon EC2 Auto Scaling groups that use multiple instance types.

Configuring Amazon EC2 Auto Scaling groups with capacity rebalancing helps ensure that Amazon EC2 instances are evenly distributed across Availability Zones, regardless of instance types and purchasing options. It uses a target tracking policy associated with the group, such as CPU utilization or network traffic.

For more information, see [Auto Scaling groups with multiple instance types and purchase options](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

AWS Config c18d2gz103

Source

AWS Config Managed Rule: autoscaling-capacity-rebalancing

Alert Criteria

Yellow: Amazon EC2 Auto Scaling group capacity rebalancing is not enabled.

Recommended Action

Ensure that capacity rebalancing is enabled for your Amazon EC2 Auto Scaling groups that use multiple instance types.

For more information, see [Enable Capacity Rebalancing \(console\)](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon EC2 Auto Scaling is not deployed in multiple AZs or does not meet the minimum number of AZs

Description

Checks if the Amazon EC2 Auto Scaling group is deployed in multiple Availability Zones, or the minimum number of Availability Zones specified. Deploy Amazon EC2 instances in multiple Availability Zones to ensure high availability.

You can adjust the minimum number of Availability Zones using the **minAvailabilityZones** parameter in your AWS Config rules.

For more information, see [Auto Scaling groups with multiple instance types and purchase options](#).

Check ID

c18d2gz101

Source

AWS Config Managed Rule: autoscaling-multiple-az

Alert Criteria

Red: The Amazon EC2 Auto Scaling group doesn't have multiple AZs configured, or doesn't meet the minimum number of AZs specified.

Recommended Action

Make sure that your Amazon EC2 Auto Scaling group is configured with multiple AZs. Deploy Amazon EC2 instances in multiple Availability Zones to ensure high availability.

Additional Resources

[Create an Auto Scaling group using a launch template](#)

[Create an Auto Scaling group using a launch configuration](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon EC2 Availability Zone Balance

Description

Checks the distribution of Amazon Elastic Compute Cloud (Amazon EC2) instances across Availability Zones in a Region.

Availability Zones are distinct locations that are insulated from failures in other Availability Zones. This allows inexpensive, low-latency network connectivity between Availability Zones in the same Region. By launching instances in multiple Availability Zones in the same Region, you can help protect your applications from a single point of failure.

Check ID

wuy7G1zxq1

Alert Criteria

- Yellow: The Region has instances in multiple zones, but the distribution is uneven (the difference between the highest and lowest instance counts in utilized Availability Zones is greater than 20%).
- Red: The Region has instances only in a single Availability Zone.

Recommended Action

Balance your Amazon EC2 instances evenly across multiple Availability Zones. You can do this by launching instances manually or by using Auto Scaling to do it automatically. For more information, see [Launch Your Instance](#) and [Load Balance Your Auto Scaling Group](#).

Additional Resources

- [Amazon EC2 Auto Scaling User Guide](#)
- [Placement groups for your Amazon EC2 instances](#)
- [Amazon EC2 instance types](#)

Report columns

- Status
- Region
- Zone a Instances
- Zone b Instances
- Zone c Instances
- Zone e Instances
- Zone f Instances
- Reason

Amazon EC2 detailed monitoring not enabled

Description

Checks if detailed monitoring is enabled for your Amazon EC2 instances.

Amazon EC2 detailed monitoring provides more frequent metrics, published at one-minute intervals, instead of the five-minute intervals used in Amazon EC2 basic monitoring. Enabling detailed monitoring for Amazon EC2 helps you better manage your Amazon EC2 resources, so that you can find trends and take action faster.

For more information, see [Basic monitoring and detailed monitoring](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

AWS Config c18d2gz144

Source

AWS Config Managed Rule: ec2-instance-detailed-monitoring-enabled

Alert criteria

Yellow: Detailed monitoring is not enabled for Amazon EC2 instances.

Recommended action

Turn on detailed monitoring for your Amazon EC2 instances to increase the frequency at which Amazon EC2 metric data is published to Amazon CloudWatch (from 5-minute to 1-minute intervals).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon ECS AWSLogs driver in blocking mode

Description

Checks for Amazon ECS task definitions that are actively in use by running tasks or are among the 5 most recent revisions per task definition family, configured with the AWSLogs logging driver in blocking mode. A driver configured in the blocking mode risks system availability.

This check does not consider account-level driver configuration settings.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1dvkm4z6b

Alert criteria

Yellow: The awslogs driver logging configuration parameter mode is set to blocking.

Green: Amazon ECS task definition is not using the awslogs driver or the awslogs driver is configured in non-blocking mode.

Recommended action

To mitigate the availability risk, consider changing the task definition AWSLogs driver configuration from blocking to non-blocking. With non-blocking mode, you will have to set a value for the max-buffer-size parameter. For more information and guidance on configuration parameters, see [Preventing log loss with non-blocking mode in the AWSLogs container log driver](#).

Additional resources

[Using the AWS logs log driver](#)

[Choosing container logging options to avoid backpressure](#)

[Preventing log loss with non-blocking mode in the AWSLogs container log driver](#)

Report columns

- Status
- Region
- Task Definition ARN
- Container Definition Names
- Last Updated Time

Amazon ECS service using a single AZ

Description

Checks that your service configuration uses a single Availability Zone (AZ).

An AZ is a distinct location that is insulated from failures in other zones. This supports inexpensive, low-latency network connectivity between AZs in the same AWS Region. By launching instances in multiple AZs in the same Region, you can help protect your applications from a single point of failure.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7dfpz01

Alert criteria

- Yellow: An Amazon ECS service is running all tasks in a single AZ.
- Green: An Amazon ECS service is running tasks in at least two different AZs.

Recommended action

Create at least one more task for the service in a different AZ.

Additional resources

[Amazon ECS capacity and availability](#)

Report columns

- Status
- Region
- ECS Cluster Name/ECS Service Name

- Number of Availability Zones
- Last Updated Time

Amazon ECS Multi-AZ placement strategy

Description

Checks that your Amazon ECS service uses the spread placement strategy based on Availability Zone (AZ). This strategy distributes tasks across Availability Zones in the same AWS Region and can help protect your applications from a single point of failure.

For tasks that run as part of an Amazon ECS service, spread is the default task placement strategy.

This check also verifies that spread is the first or only strategy in your list of enabled placement strategies.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1z7dfpz02

Alert criteria

- Yellow: Spread by availability zone is disabled or isn't the first strategy in your list of enabled placement strategies for your Amazon ECS service.
- Green: Spread by availability zone is the first strategy in your list of enabled placement strategies or the only placement strategy enabled for your Amazon ECS service.

Recommended action

Enable the spread task placement strategy to distribute tasks across multiple AZs. Verify that spread by availability zone is the first strategy for all enabled task placement strategies or the

only strategy used. If you choose to manage AZ placement, you can use a mirrored service in another AZ to mitigate these risks.

Additional resources

[Amazon ECS task placement strategies](#)

Report columns

- Status
- Region
- ECS Cluster Name/ECS Service Name
- Spread Task Placement Strategy Enabled and Applied Correctly
- Last Updated Time

Amazon EFS no mount target redundancy

Description

Checks if mount targets exist in multiple Availability Zones for an Amazon EFS file system.

An Availability Zone is a distinct location that is insulated from failures in other zones. By creating mount targets in multiple geographically separated Availability Zones within an AWS Region, you can achieve the highest levels of availability and durability for your Amazon EFS file systems.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1dfprch01

Alert criteria

- Yellow: File system has 1 mount target created in a single Availability Zone.

Green: File system has 2 or more mount targets created in multiple Availability Zones.

Recommended action

For EFS file systems using One Zone storage classes, we recommend you create new file systems that use Standard storage classes by restoring a backup to a new file system. Then create mount targets in multiple Availability Zones.

For EFS file systems using Standard storage classes, we recommend you create mount targets in multiple Availability Zones.

Additional resources

- [Managing mount targets using the Amazon EFS console](#)
- [Amazon EFS Quotas and Limits](#)

Report columns

- Status
- Region
- EFS File System ID
- Number of mount targets
- Number of AZs
- Last Updated Time

Amazon EFS not in AWS Backup plan

Description

Checks if Amazon EFS file systems are included in backup plans with AWS Backup.

AWS Backup is a unified backup service designed to simplify the creation, migration, restoration, and deletion of backups, while providing improved reporting and auditing.

For more information, see [Backing up your Amazon EFS file systems](#).

Check ID

c18d2gz117

Source

AWS Config Managed Rule: EFS_IN_BACKUP_PLAN

Alert criteria

Red: Amazon EFS file systems aren't included in AWS Backup plan.

Recommended action

Make sure that your Amazon EFS file systems are included in your AWS Backup plan to protect against accidental data loss or data corruption.

Additional resources

[Backing up your Amazon EFS file systems](#)

[Amazon EFS Backup and Restore using AWS Backup.](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon ElastiCache Multi-AZ clusters

Description

Checks for ElastiCache clusters that deploy in a single Availability Zone (AZ). This check alerts you if Multi-AZ is inactive in a cluster.

Deployments in multiple AZs enhance ElastiCache cluster availability by asynchronously replicating to read-only replicas in a different AZ. When planned cluster maintenance occurs, or a primary node is unavailable, ElastiCache automatically promotes a replica to primary. This failover allows cluster write operations to resume, and doesn't require an administrator to intervene.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

ECHdfsQ402

Alert criteria

- Green: Multi-AZ is active in the cluster.
- Yellow: Multi-AZ is inactive in the cluster.

Recommended action

Create at least one replica per shard, in an AZ that is different than the primary.

Additional resources

For more information, see [Minimizing downtime in ElastiCache \(Redis OSS\) with Multi-AZ](#).

Report columns

- Status
- Region
- Cluster Name
- Last Updated Time

ElastiCache (Redis OSS) Clusters Automatic Backup

Description

Checks if the Amazon ElastiCache (Redis OSS) clusters have automatic backup turned on and if the snapshot retention period is above the specified or 15 day default limit. When automatic backups are enabled, ElastiCache creates a backup of the cluster on a daily basis.

You can specify your desired snapshot retention limit using the **snapshotRetentionPeriod** parameters of your AWS Config rules.

For more information, see [Backup and restore for ElastiCache \(Redis OSS\)](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz178

Source

AWS Config Managed Rule: elasticache-redis-cluster-automatic-backup-check

Alert Criteria

Red: Amazon ElastiCache (Redis OSS) clusters do not have automatic backup turned on or the snapshot retention period is below the limit.

Recommended Action

Make sure that Amazon ElastiCache (Redis OSS) clusters have automatic backup turned on and the snapshot retention period is above the specified or 15 day default limit. Automatic backups can help guard against data loss. In the event of a failure, you can create a new cluster, restoring your data from the most recent backup.

For more information, see [Backup and restore for ElastiCache \(Redis OSS\)](#).

Additional Resources

For more information, see [Scheduling automatic backups](#).

Report columns

- Status
- Region
- Cluster Name
- Last Updated Time

Amazon MemoryDB Multi-AZ clusters

Description

Checks for MemoryDB clusters that deploy in a single Availability Zone (AZ). This check alerts you if Multi-AZ is inactive in a cluster.

Deployments in multiple AZs enhance MemoryDB cluster availability by asynchronously replicating to read-only replicas in a different AZ. When planned cluster maintenance occurs, or a primary node is unavailable, MemoryDB automatically promotes a replica to primary. This failover allows cluster write operations to resume, and doesn't require an administrator to intervene.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

MDBdfsQ401

Alert Criteria

- Green: Multi-AZ is active in the cluster.
- Yellow: Multi-AZ is inactive in the cluster.

Recommended Action

Create at least one replica per shard, in an AZ that is different than the primary.

Additional Resources

For more information, see [Minimizing downtime in MemoryDB with Multi-AZ](#).

Report columns

- Status
- Region

- Cluster Name
- Last Updated Time

Amazon MSK brokers hosting too many partitions

Description

Checks that the brokers of a Managed Streaming for Kafka (MSK) Cluster do not have more than the recommended number of partitions assigned.

Check ID

Cmsvnj8vf1

Alert Criteria

- Red: Your MSK broker has reached or exceeded 100% of the recommended maximum partition limit
- Yellow: Your MSK has reached 80% of the recommended maximum partition limit

Recommended Action

Follow the MSK [recommended best practices](#) to scale your MSK Cluster or delete any unused partitions.

Additional Resources

- [Right-sizing your Cluster](#)

Report columns

- Status
- Region
- Cluster ARN
- Broker ID
- Partition Count

Amazon MSK Cluster Multi-AZ

Description

Checks the number of Availability Zones (AZs) for your Amazon MSK provisioned cluster. The Amazon MSK cluster is formed of several brokers that work together and distribute the data

and load. Production might be interrupted during maintenance or broker issues in a 2-AZ cluster.

Check ID

90046ff5b5

Alert Criteria

- Yellow: The Amazon MSK cluster is provisioned with brokers in only two AZs
- Green: The Amazon MSK cluster is provisioned with brokers across three or more AZs

Recommended Action

To increase availability of the cluster, you can create another cluster in a 3 AZs setup. Then migrate the existing cluster to the new cluster that you created. You can use Amazon MSK replication for this migration.

Additional Resources

[Amazon MSK high availability](#)

[Amazon MSK migration](#)

Report columns

- Status
- Region
- MSK Cluster ARN
- Number of AZs
- Last Updated Time

Amazon OpenSearch Service domains with less than three data nodes

Description

Checks if Amazon OpenSearch Service domains are configured with at least three data nodes and `ZoneAwarenessEnabled` is true. With `ZoneAwarenessEnabled` enabled, Amazon OpenSearch Service ensures that each primary shard and its corresponding replica are allocated in different Availability Zones.

For more information, see [Configuring a multi-AZ domain in Amazon OpenSearch Service](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz183

Source

AWS Config Managed Rule: opensearch-data-node-fault-tolerance

Alert Criteria

Yellow: Amazon OpenSearch Service domains are configured with less than three data nodes.

Recommended Action

Make sure that Amazon OpenSearch Service domains are configured with a minimum of three data nodes. Configure a multi-AZ domain to enhance the availability of the Amazon OpenSearch Service cluster by allocating nodes and replicating data across three Availability Zones within the same Region. This prevents data loss and minimizes downtime in the event of node or data center (AZ) failure.

For more information, see [Increase availability for Amazon OpenSearch Service by deploying in three Availability Zones](#).

Additional Resources

- [Increase availability for Amazon OpenSearch Service by deploying in three Availability Zones](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule

- Input Parameters
- Last Updated Time

Amazon RDS Backups

Description

Checks for automated backups of Amazon RDS DB instances.

By default, backups are enabled with a retention period of one day. Backups reduce the risk of unexpected data loss and allow for point-in-time recovery.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

opQPADkZvH

Alert Criteria

Red: A DB instance has the backup retention period set to 0 days.

Recommended Action

Set the retention period for the automated DB instance backup to 1 to 35 days as appropriate to the requirements of your application. See [Working With Automated Backups](#).

Additional Resources

[Getting Started with Amazon RDS](#)

Report columns

- Status
- Region/AZ
- DB Instance

- VPC ID
- Backup Retention Period

Amazon RDS Continuous Backup Not Enabled

Description

Checks if an Amazon RDS instance is enabled with automated backups using Amazon RDS or with continuous backups of AWS Backup. Continuous backups reduce the risk of unexpected data loss and allow for point-in-time recovery.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

44fde09ab5

Alert Criteria

- Red: The instance doesn't have automated backup enabled in Amazon RDS or continuous backup in AWS Backup.
- Red: MySQL versions below 5.6 don't support automated or continuous backup. To provide resiliency, first upgrade the database version and then enable automated or continuous backup.
- Green: The instance has automated backup enabled in Amazon RDS.
- Green: The instance has continuous backup enabled in AWS Backup.

Recommended Action

Make sure that Amazon RDS instances have automated backup configured either by setting a retention period greater than 0 in Amazon RDS or by creating a continuous backup plan using AWS Backup.

Additional Resources

- [Getting Started with Amazon RDS](#)
- [Managing automated backups](#)
- [Introduction to backups](#)
- [Continuous backups and point-in-time restore \(PITR\)](#)
- [AWS Backup feature availability](#)

Report columns

- Status
- Region
- DB Instance Identifier
- DB Instance ARN
- Deployment Type
- Backup Type
- Reason
- Last Updated Time

Amazon RDS DB clusters have one DB instance

Description

Add at least another DB instance to the DB cluster to improve availability and performance.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the

recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt011

Alert Criteria

Yellow: DB clusters have only one DB instance.

Recommended Action

Add a reader DB instance to the DB cluster.

Additional Resources

In the current configuration, one DB instance is used for both read and write operations. You can add another DB instance to allow read redistribution and a failover option.

For more information, see [High availability for Amazon Aurora](#).

Report columns

- Status
- Region
- Resource
- Engine Name
- DB Instance Class
- Last Updated Time

Amazon RDS DB clusters with all instances in the same Availability Zone

Description

The DB clusters are currently in a single Availability Zone. Use multiple Availability Zones to improve the availability.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt007

Alert Criteria

Yellow: DB clusters have all the instances in the same Availability Zone.

Recommended Action

Add the DB instances to multiple Availability Zones in your DB cluster.

Additional Resources

We recommend that you add the DB instances to multiple Availability Zones in a DB cluster. Adding DB instances to multiple Availability Zones improves the availability of your DB cluster.

For more information, see [High availability for Amazon Aurora](#).

Report columns

- Status

- Region
- Resource
- Engine Name
- Last Updated Time

Amazon RDS DB clusters with all reader instances in the same Availability Zone

Description

Your DB cluster has all the reader instances in the same Availability Zone. We recommend that you distribute the Reader instances across multiple Availability Zones in your DB cluster.

Distribution increases the availability of the database, and improves the response time by reducing network latency between clients and the database.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt018

Alert Criteria

Red: DB clusters have the reader instances in the same Availability Zone.

Recommended Action

Distribute the reader instances across multiple Availability Zones.

Additional Resources

Availability Zones (AZs) are locations that are distinct from each other to provide isolation in case of outages within each AWS Region. We recommend that you distribute the primary instance and reader instances in your DB cluster across multiple AZs to improve the availability of your DB cluster. You can create a Multi-AZ cluster using the AWS Management Console, AWS CLI, or Amazon RDS API when you create the cluster. You can modify the existing Aurora cluster to a Multi-AZ cluster by adding a new reader instance and specifying a different AZ.

For more information, see [High availability for Amazon Aurora](#).

Report columns

- Status
- Region
- Resource
- Engine Name
- Last Updated Time

Amazon RDS DB Instance Enhanced Monitoring not enabled

Description

Checks if your Amazon RDS DB instances have Enhanced Monitoring enabled.

Amazon RDS Enhanced Monitoring provides metrics in real time for the operating system (OS) that your DB instance runs on. All system metrics and process information for your Amazon RDS DB instances can be view on the Amazon RDS console. And, you can customize the dashboard. With Enhanced Monitoring, you have visibility of your Amazon RDS instance operation status in near real time, allowing you to respond to operational issues faster.

You can specify your desired monitoring interval using the **monitoringInterval** parameter of your AWS Config rules.

For more information, see [Overview of Enhanced Monitoring](#) and [OS metrics in Enhanced Monitoring](#).

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz158

Source

AWS Config Managed Rule: rds-enhanced-monitoring-enabled

Alert Criteria

Yellow: Your Amazon RDS DB instances don't have Enhanced Monitoring enabled or are not configured with the desired interval.

Recommended Action

Enable Enhanced Monitoring for your Amazon RDS DB instances to improve the visibility of your Amazon RDS instance operation status.

For more information, see [Monitoring OS metrics with Enhanced Monitoring](#).

Additional Resources

[OS metrics in Enhanced Monitoring](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule

- Input Parameters
- Last Updated Time

Amazon RDS DB instances have storage autoscaling turned off

Description

Amazon RDS storage autoscaling isn't turned on for your DB instance. When there is an increase in the database workload, RDS Storage autoscaling automatically scales the storage capacity with zero downtime.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt013

Alert Criteria

Red: DB instances don't have storage autoscaling turned on.

Recommended Action

Turn on Amazon RDS storage autoscaling with a specified maximum storage threshold.

Additional Resources

Amazon RDS storage autoscaling automatically scales storage capacity with zero downtime when the database workload increases. Storage autoscaling monitors the storage usage and automatically scales up the capacity when the usage is close to the provisioned storage capacity. You can specify a maximum limit on the storage that Amazon RDS can allocate to the DB instance. There is no additional cost for storage autoscaling. You pay only for the Amazon RDS resources that are allocated to your DB instance. We recommend that you turn on Amazon RDS storage autoscaling.

For more information, see [Managing capacity automatically with Amazon RDS storage autoscaling](#).

Report columns

- Status
- Region
- Resource
- Recommended Value
- Engine Name
- Last Updated Time

Amazon RDS DB instances not using Multi-AZ deployment

Description

We recommend that you use Multi-AZ deployment. The Multi-AZ deployments enhance the availability and durability of the DB instance.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

 Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt019

Alert Criteria

Yellow: DB instances aren't using Multi-AZ deployment.

Recommended Action

Set up Multi-AZ for the impacted DB instances.

Additional Resources

In an Amazon RDS Multi-AZ deployment, Amazon RDS automatically creates a primary database instance and replicates the data to an instance in a different availability zone. When it detects a failure, Amazon RDS automatically fails over to a standby instance without manual intervention.

For more information, see [Pricing](#).

Report columns

- Status
- Region
- Resource

- Engine Name
- Last Updated Time

Amazon RDS DiskQueueDepth

Description

Checks to see if the CloudWatch metric DiskQueueDepth shows that number of queued writes to the RDS Instance database storage has grown to a level where an operational investigation should be suggested.

Check ID

Cmsvnj8db3

Alert Criteria

- Red: DiskQueueDepth CloudWatch metric has exceeded 10
- Yellow: DiskQueueDepth CloudWatch metric is greater than 5 but less than or equal to 10
- Green: DiskQueueDepth CloudWatch metric is less than or equal to 5

Recommended Action

Consider moving to instances and storage volumes that support the read/write characteristics.

Report columns

- Status
- Region
- DB Instance ARN
- DiskQueueDepth Metric

Amazon RDS FreeStorageSpace

Description

Checks to see if the FreeStorageSpace CloudWatch metric for an RDS database instance has decreased below an operationally reasonable threshold.

Check ID

Cmsvnj8db2

Alert Criteria

- Red: FreeStorageSpace has less than 10% of total capacity
- Yellow: FreeStorageSpace is between 10% and 20% of total capacity
- Green: FreeStorageSpace is more than 20% of total capacity

Recommended Action

Scale up the storage space for the RDS database instance that is running low on free storage using the Amazon RDS Management Console, Amazon RDS API or AWS Command Line Interface.

Report columns

- Status
- Region
- DB Instance ARN
- FreeStorageSpace Metric (MB)
- DB Instance Allocated Storage (MB)
- DB Instance Storage Used Percent

Amazon RDS log_output parameter is set to table

Description

When **log_output** is set to **TABLE**, more storage is used than when **log_output** is set to **FILE**. We recommend that you set the parameter to **FILE**, to avoid reaching the storage size limit.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

 Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt023

Alert Criteria

Yellow: DB parameter groups have **log_output** parameter set to **TABLE**.

Recommended Action

Set the **log_output** parameter value to **FILE** in your DB parameter groups.

Additional Resources

For more information, see [MySQL database log files](#).

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS `innodb_default_row_format` parameter setting is unsafe

Description

Your DB instance encounters a known issue: A table created in a MySQL version lower than 8.0.26 with the `row_format` set to **COMPACT** or **REDUNDANT** is inaccessible and unrecoverable when the index exceeds 767 bytes.

We recommend that you set the `innodb_default_row_format` parameter value to **DYNAMIC**.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt036

Alert Criteria

Red: DB parameter groups have an unsafe setting for the `innodb_default_row_format` parameter.

Recommended Action

Set the **innodb_default_row_format** parameter to **DYNAMIC**.

Additional Resources

When a table is created with MySQL version lower than 8.0.26 with **row_format** set to **COMPACT** or **REDUNDANT**, creating indexes with a key prefix shorter than 767 bytes isn't enforced. After the database restarts, these tables can't be accessed or recovered.

For more information, see [Changes in MySQL 8.0.26 \(2021-07-20, General Availability\)](#) on the MySQL documentation website.

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS innodb_flush_log_at_trx_commit parameter is not 1

Description

The value of the **innodb_flush_log_at_trx_commit** parameter of your DB instance isn't a safe value. This parameter controls the persistence of commit operations to disk.

We recommend that you set the **innodb_flush_log_at_trx_commit** parameter to 1.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt030

Alert Criteria

Yellow: DB parameter groups have **innodb_flush_log_at_trx_commit** set to other than 1.

Recommended Action

Set the **innodb_flush_log_at_trx_commit** parameter value to 1

Additional Resources

The database transaction is durable when the log buffer is saved to the durable storage. However, saving to the disk impacts performance. Depending on the value set for **innodb_flush_log_at_trx_commit** parameter, the behavior of how logs are written and saved to the disk can vary.

- When the parameter value is 1, the logs are written and saved to the disk after each committed transaction.
- When the parameter value is 0, the logs are written and saved to the disk once per second.
- When the parameter value is 2, the logs are written after each transaction is committed and saved to the disk once per second. The data moves from the InnoDB memory buffer to the operating system's cache which is also in the memory.

Note

When the parameter value is not 1, InnoDB doesn't assure ACID properties. The recent transactions for the last second may be lost when the database crashes.

For more information, see [Best practices for configuring parameters for Amazon RDS for MySQL, part 1: Parameters related to performance](#).

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS `max_user_connections` parameter is low

Description

Your DB instance has a low value for the maximum number of simultaneous connections for each database account.

We recommend setting the `max_user_connections` parameter to a number greater than 5.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt034

Alert Criteria

Yellow: DB parameter groups have **max_user_connections** misconfigured.

Recommended Action

Increase the value of the **max_user_connections** parameter to a number greater than 5.

Additional Resources

The **max_user_connections** setting controls the maximum number of simultaneous connections allowed for a MySQL user account. Reaching this connection limit cause failures in the Amazon RDS instance administration operations, such as backup, patching, and parameters changes.

For more information, see [Setting Account Resource Limits](#) on the MySQL documentation website.

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS Multi-AZ

Description

Checks for DB instances that are deployed in a single Availability Zone (AZ).

Multi-AZ deployments enhance database availability by synchronously replicating to a standby instance in a different Availability Zone. During planned database maintenance, or the failure

of a DB instance or Availability Zone, Amazon RDS automatically fails over to the standby. This failover allows database operations to resume quickly without administrative intervention. Because Amazon RDS does not support Multi-AZ deployment for Microsoft SQL Server, this check does not examine SQL Server instances.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

f2iK5R6Dep

Alert Criteria

Yellow: A DB instance is deployed in a single Availability Zone.

Recommended Action

If your application requires high availability, modify your DB instance to enable Multi-AZ deployment. See [High Availability \(Multi-AZ\)](#).

Additional Resources

[Regions and Availability Zones](#)

Report columns

- Status
- Region/AZ
- DB Instance
- VPC ID
- Multi-AZ

Amazon RDS Not In AWS Backup Plan**Description**

Checks if your Amazon RDS DB instances are included in a backup plan in AWS Backup.

AWS Backup is a fully managed backup service that makes it easy to centralize and automate backing up data across AWS services.

Including your Amazon RDS DB instance in a backup plan is important for regulatory compliance obligations, disaster recovery, business policies for data protection, and business continuity goals.

For more information, see [What is AWS Backup?](#).

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz159

Source

AWS Config Managed Rule: rds-in-backup-plan

Alert Criteria

Yellow: An Amazon RDS DB instance is not included in a backup plan with AWS Backup.

Recommended Action

Include your Amazon RDS DB instances in a backup plan with AWS Backup.

For more information, see [Amazon RDS Backup and Restore Using AWS Backup](#).

Additional Resources

[Assigning resources to a backup plan](#)

Report columns

- Status
- Region

- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon RDS Read Replicas are open in writable mode

Description

Your DB instance has a read replica in writable mode, which allows updates from clients.

We recommend that you set the `read_only` parameter to `TrueIfReplica` so that the read replicas isn't in writable mode.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt035

Alert Criteria

Yellow: DB parameter groups turn on writable mode for the read replicas.

Recommended Action

Set the **read_only** parameter value to **TruelfReplica**.

Additional Resources

The **read_only** parameter controls the write permission from the clients to a database instance. The default value for this parameter is **TruelfReplica**. For a replica instance, **TruelfReplica** sets the **read_only** value to ON (1) and disables any write activity from the clients. For a master/writer instance, **TruelfReplica** sets the value to OFF (0) and enables the write activity from the clients for the instance. When the read replica is opened in writable mode, the data stored in this instance may diverge from the primary instance which causes replication errors.

For more information, see [Best practices for configuring parameters for Amazon RDS for MySQL, part 2: Parameters related to replication](#) on the MySQL documentation website.

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon RDS resource automated backups is turned off

Description

Automated backups are disabled on your DB resources. Automated backups enable point-in-time recovery of your DB instance.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

 Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt001

Alert Criteria

Red: Amazon RDS resources don't have automated backups turned on

Recommended Action

Turn on automated backups with a retention period of up to 14 days.

Additional Resources

Automated backups enable point-in-time recovery of your DB instances. We recommend turning on automated backups. When you turn on automated backups for a DB instance, Amazon RDS automatically performs a full backup of your data daily during your preferred backup window. The backup captures transaction logs when there are updates to your DB instance. You get backup storage up to the storage size of your DB instance at no additional cost.

For more information, see the following resources:

- [Enabling automated backups](#)
- [Demystifying Amazon RDS backup storage costs](#)

Report columns

- Status
- Region
- Resource
- Recommended Value
- Engine Name
- Last Updated Time

Amazon RDS sync_binlog parameter is turned off

Description

The synchronization of the binary log to disk isn't enforced before the transaction commits are acknowledged in your DB instance.

We recommend that you set the **sync_binlog** parameter value to **1**.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt031

Alert Criteria

Yellow: DB parameter groups have synchronous binary logging turned off.

Recommended Action

Set the **sync_binlog** parameter to **1**.

Additional Resources

The **sync_binlog** parameter controls how MySQL pushes the binary log to disk. When the value of this parameter is set to **1**, it turns on binary log synchronization to disk before transactions are committed. When the value of this parameter is set to **0**, it turns off the binary log synchronization to the disk. Typically, MySQL server depends on the operating system to push the binary log to disk regularly similar to other files. The **sync_binlog** parameter value set to **0** can enhance the performance. However, during a power failure or an operating system crash, the server loses all the committed transactions that weren't synchronized to the binary logs.

For more information, see [Best practices for configuring parameters for Amazon RDS for MySQL, part 2: Parameters related to replication](#).

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

RDS DB Cluster has no Multi-AZ replication enabled

Description

Checks if your Amazon RDS DB clusters have Multi-AZ replication enabled.

A Multi-AZ DB cluster has a writer DB instance and two reader DB instances in three separate Availability Zones. Multi-AZ DB clusters provide high availability, increased capacity for read workloads, and lower latency when compared to Multi-AZ deployments.

For more information, see [Creating a Multi-AZ DB cluster](#).

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz161

Source

AWS Config Managed Rule: rds-cluster-multi-az-enabled

Alert Criteria

Yellow: Your Amazon RDS DB cluster does not have Multi-AZ replication configured

Recommended Action

Turn on Multi-AZ DB cluster deployment when you create an Amazon RDS DB cluster.

For more information, see [Creating a Multi-AZ DB cluster](#).

Additional Resources

[Multi-AZ DB cluster deployments](#)

Report columns

- Status
- Region
- Resource

- AWS Config Rule
- Input Parameters
- Last Updated Time

RDS Multi-AZ Standby Instance Not Enabled

Description

Checks if your Amazon RDS DB instances have a Multi-AZ standby replica configured.

Amazon RDS Multi-AZ provides high availability and durability for database instances by replicating data to a standby replica in a different Availability Zone. This provides automatic failover, improve performance, and enhances data durability. In a Multi-AZ DB instance deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to a standby replica to provide data redundancy and minimize latency spikes during system backups. Running a DB instance with high availability enhances availability during planned system maintenance. It can also help protect your databases against DB instance failure and Availability Zone disruption.

For more information, see [Multi-AZ DB instance deployments](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz156

Source

AWS Config Managed Rule: rds-multi-az-support

Alert Criteria

Yellow: An Amazon RDS DB instance does not have a Multi-AZ replica configured.

Recommended Action

Turn on Multi-AZ deployment when you create an Amazon RDS DB instance.

This check can't be excluded from view in the Trusted Advisor console.

Additional Resources

[Multi-AZ DB instance deployments](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon RDS ReplicaLag

Description

Checks to see if the ReplicaLag CloudWatch metric for an RDS database instance has increased above an operationally reasonable threshold over the past week.

ReplicaLag metric measures the number of seconds a read replica is behind the primary instance. Replication lag occurs when the asynchronous updates made to the read replica cannot keep up with the updates happening on the primary database instance. In the event of a failure to the primary instance, data could be missing from the read replica if the ReplicaLag is above an operationally reasonable threshold.

Check ID

Cmsvnj8db1

Alert Criteria

- Red: ReplicaLag metric exceeded 60 seconds at least once during the week.

- Yellow: ReplicaLag metric exceeded 10 seconds at least once during the week.
- Green: ReplicaLag is less than 10 seconds.

Recommended Action

There are several possible causes for ReplicaLag to increase beyond operationally safe levels. For example, it can be caused by recently replaced/launched replica instances from older backups and these replicas requiring substantial time to “catch-up” to the primary database instance and live transactions. This ReplicaLag may dwindle over time as catch-up occurs. Another example could be that the transaction velocity able to be achieved on the primary database instance is higher than the replication process or replica infrastructure is able to match. This ReplicaLag may grow over time as replication fails to keep pace with the primary database performance. Finally, the workload may be bursty throughout different periods of the day/month/etc. that result in occasional ReplicaLag to fall behind. Your team should investigate which possible root cause has contributed to high ReplicaLag for the database, and possibly change the database instance type or other characteristics of the workload to ensure data continuity on the replica matches your requirements.

Additional Resources

- [Working with read replicas for Amazon RDS for PostgreSQL](#)
- [Working with MySQL replication in Amazon RDS](#)
- [Working with MySQL read replicas](#)

Report columns

- Status
- Region
- DB Instance ARN
- ReplicaLag Metric

Amazon RDS `synchronous_commit` parameter is turned off

Description

When the **`synchronous_commit`** parameter is turned off, data can be lost in a database crash. The durability of the database is at risk.

We recommend that you turn on the **`synchronous_commit`** parameter.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt026

Alert Criteria

Red: DB parameter groups have **synchronous_commit** parameter turned off.

Recommended Action

Turn on **synchronous_commit** parameter in your DB parameter groups.

Additional Resources

The **synchronous_commit** parameter defines the Write-Ahead Logging (WAL) process completion before the database server sends a successful notification to the client. This commit is called as an asynchronous commit because the client acknowledges the commit before WAL saves the transaction in the disk. If the **synchronous_commit** parameter is turned off, then the transactions can be lost, DB instance durability might be compromised, and data might be lost when a database crashes.

For more information, see [MySQL database log files](#).

Report columns

- Status
- Region
- Resource
- Parameter Name
- Recommended Value
- Last Updated Time

Amazon Redshift cluster automated snapshots

Description

Checks if automated snapshots are enabled for your Amazon Redshift clusters.

Amazon Redshift automatically takes incremental snapshots that track changes to the cluster since the previous automated snapshot. Automated snapshots retain all of the data required to restore a cluster from a snapshot. To disable automated snapshots, set the retention period to zero. You can't disable automated snapshots for RA3 node types.

You can specify your desired minimum and maximum retention period using the **MinRetentionPeriod** and **MaxRetentionPeriod** parameter of your AWS Config rules.

[Amazon Redshift snapshots and backups](#)

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz135

Source

AWS Config Managed Rule: `redshift-backup-enabled`

Alert Criteria

Red: Amazon Redshift does not have automated snapshots configured within the desired retention period.

Recommended Action

Make sure that automated snapshots are enabled for your Amazon Redshift clusters.

For more information, see [Managing snapshots using the console](#).

Additional Resources

[Amazon Redshift snapshots and backups](#)

For more information, see [Working with backups](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon Route 53 Deleted Health Checks

Description

Checks for resource record sets that are associated with health checks that have been deleted.

Route 53 does not prevent you from deleting a health check that is associated with one or more resource record sets. If you delete a health check without updating the associated resource record sets, the routing of DNS queries for your DNS failover configuration will not work as intended.

Hosted zones created by AWS services won't appear in your check results.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

Cb877eB72b

Alert Criteria

Yellow: A resource record set is associated with a health check that has been deleted.

Recommended Action

Create a new health check and associate it with the resource record set. See [Creating, Updating, and Deleting Health Checks](#) and [Adding Health Checks to Resource Record Sets](#).

Additional Resources

- [Amazon Route 53 Health Checks and DNS Failover](#)
- [How Health Checks Work in Simple Amazon Route 53 Configurations](#)

Report columns

- Hosted Zone Name
- Hosted Zone ID
- Resource Record Set Name
- Resource Record Set Type
- Resource Record Set Identifier

Amazon Route 53 Failover Resource Record Sets**Description**

Checks for Amazon Route 53 failover resource record sets that have a misconfiguration.

When Amazon Route 53 health checks determine that the primary resource is unhealthy, Amazon Route 53 responds to queries with a secondary, backup resource record set. You must create correctly configured primary and secondary resource record sets for failover to work.

Hosted zones created by AWS services won't appear in your check results.

 Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

b73EEdD790

Alert Criteria

- Yellow: A primary failover resource record set does not have a corresponding secondary resource record set.
- Yellow: A secondary failover resource record set does not have a corresponding primary resource record set.
- Yellow: Primary and secondary resource record sets that have the same name are associated with the same health check.

Recommended Action

If a failover resource set is missing, create the corresponding resource record set. See [Creating Failover Resource Record Sets](#).

If your resource record sets are associated with the same health check, create separate health checks for each one. See [Creating, Updating, and Deleting Health Checks](#).

Additional Resources

[Amazon Route 53 Health Checks and DNS Failover](#)

Report columns

- Hosted Zone Name
- Hosted Zone ID
- Resource Record Set Name
- Resource Record Set Type
- Reason

Amazon Route 53 High TTL Resource Record Sets

Description

Checks for resource record sets that can benefit from having a lower time-to-live (TTL) value.

TTL is the number of seconds that a resource record set is cached by DNS resolvers. When you specify a long TTL, DNS resolvers take longer to request updated DNS records, which can cause unnecessary delay in rerouting traffic (for example, when DNS Failover detects and responds to a failure of one of your endpoints). This check looks only at records with a policy of Failover, or if there is an associated health check.

Hosted zones created by AWS services won't appear in your check results.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

C056F80cR3

Alert Criteria

- Yellow: A resource record set whose routing policy is Failover has a TTL greater than 60 seconds.
- Green: A resource record either has no failover policy, or has a failover policy with a TTL less than 60.

Recommended Action

Enter a TTL value of 60 seconds for the listed resource record sets. For more information, see [Working with Resource Record Sets](#).

Additional Resources

[Amazon Route 53 Health Checks and DNS Failover](#)

Report columns

- Status

- Hosted Zone Name
- Hosted Zone ID
- Resource Record Set Name
- Resource Record Set Type
- Resource Record Set ID
- TTL

Amazon Route 53 Name Server Delegations

Description

Checks for Amazon Route 53 hosted zones for which your domain registrar or DNS is not using the correct Route 53 name servers.

When you create a hosted zone, Route 53 assigns a delegation set of four name servers. The names of these servers are ns-###.awsdns-##.com, .net, .org, and .co.uk, where ### and ## typically represent different numbers. Before Route 53 can route DNS queries for your domain, you must update your registrar's name server configuration to remove the name servers that the registrar assigned. Then, you must add all four name servers in the Route 53 delegation set. For maximum availability, you must add all four Route 53 name servers.

Hosted zones created by AWS services won't appear in your check results.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

cF171Db240

Alert Criteria

Yellow: A hosted zone for which the registrar for your domain does not use all four of the Route 53 name servers in the delegation set.

Recommended Action

Add or update name server records with your registrar or with the current DNS service for your domain to include all four of the name servers in your Route 53 delegation set. To find these values, see [Getting the Name Servers for a Hosted Zone](#). For information about adding or updating name server records, see [Creating and Migrating Domains and Subdomains to Amazon Route 53](#).

Additional Resources

[Working with Hosted Zones](#)

Report columns

- Hosted Zone Name
- Hosted Zone ID
- Number of Name Server Delegations Used

Amazon Route 53 Resolver Endpoint Availability Zone Redundancy

Description

Checks to see if your service configuration has IP addresses specified in at least two Availability Zones (AZs) for redundancy. An AZ is a distinct location that is insulated from failures in other zones. By specifying IP addresses in multiple AZs in the same Region, you can help protect your applications from a single point of failure.

Check ID

Chrv231ch1

Alert Criteria

- Yellow: IP addresses are specified only in one AZ
- Green: IP addresses are specified in at least two AZs

Recommended Action

Specify IP addresses in at least two Availability Zones for redundancy.

Additional Resources

- If you require more than one elastic network interface endpoint to be available at all times, we recommend that you create at least one more network interface than you need, to make sure you have additional capacity available for handling possible traffic surges. The additional

network interface also ensures availability during service operations like maintenance or upgrades.

- [High availability for Resolver endpoints](#)

Report columns

- Status
- Region
- Resource ARN
- Number of AZs

Amazon S3 Bucket Replication Not Enabled

Description

Checks if your Amazon S3 buckets have replication rules enabled for Cross-Region Replication, Same-Region Replication, or both.

Replication is the automatic, asynchronous copying of objects across buckets in the same or different AWS Regions. Replication copies newly created objects and object updates from a source bucket to a destination bucket or buckets. Use Amazon S3 bucket replication to help improve the resilience and compliance of your applications and data storage.

For more information, see [Replicating objects](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz119

Source

AWS Config Managed Rule: s3-bucket-replication-enabled

Alert Criteria

Yellow: Amazon S3 bucket replication rules are not enabled for Cross-Region Replication, Same-Region Replication, or both.

Recommended Action

Turn on Amazon S3 bucket replication rules to improve the resiliency and compliance of your applications and data storage.

For more information, see [View your backup jobs and recovery points](#) and [Setting up replication](#).

Additional Resources

[Walkthroughs: Examples for configuring replication](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon S3 Bucket Versioning

Description

Checks for Amazon Simple Storage Service buckets that do not have versioning enabled, or that have versioning suspended.

When versioning is enabled, you can easily recover from both unintended user actions and application failures. Versioning allows you to preserve, retrieve, and restore any version of any object stored in a bucket. You can use lifecycle rules to manage all versions of your objects, as well as their associated costs, by automatically archiving objects to the Glacier storage class. Rules can also be configured to remove versions of your objects after a specified period of time. You can also require multi-factor authentication (MFA) for any object deletions or configuration changes to your buckets.

Versioning can't be deactivated after it has been enabled. However, it can be suspended, which prevents new versions of objects from being created. Using versioning can increase your costs for Amazon S3, because you pay for storage of multiple versions of an object.

Check ID

R365s2Qddf

Alert Criteria

- Green: Versioning is enabled for the bucket.
- Yellow: Versioning is not enabled for the bucket.
- Yellow: Versioning is suspended for the bucket.
- Yellow: Trusted Advisor doesn't have access to validate versioning.

Recommended Action

Enable bucket versioning on most buckets to prevent accidental deletion or overwriting. See [Using Versioning](#) and [Enabling Versioning Programmatically](#).

If bucket versioning is suspended, consider re-enabling versioning. For information on working with objects in a versioning-suspended bucket, see [Managing Objects in a Versioning-Suspended Bucket](#).

When versioning is enabled or suspended, you can define lifecycle configuration rules to mark certain object versions as expired or to permanently remove unneeded object versions. For more information, see [Object Lifecycle Management](#).

MFA Delete requires additional authentication when the versioning status of the bucket is changed or when versions of an object are deleted. It requires the user to enter credentials and a code from an approved authentication device. For more information, see [MFA Delete](#).

Additional Resources

[Working with Buckets](#)

Report columns

- Status
- Region
- Bucket Name
- Versioning
- MFA Delete Enabled

Application, Network, and Gateway Load Balancers Not Spanning Multiple Availability Zones

Description

Checks If your load balancers (Application, Network, and Gateway Load Balancer) are configured with subnets across multiple Availability Zones.

You can specify your desired minimum Availability Zones in the **minAvailabilityZones** parameters of your AWS Config rules.

For more information, see [Availability Zones for your Application Load Balancer](#), [Availability Zones - Network Load Balancers](#), and [Create a Gateway Load Balancer](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz169

Source

AWS Config Managed Rule: elbv2-multiple-az

Alert Criteria

Yellow: Application, Network, or Gateway Load Balancers configured with subnets in less than two Availability Zones.

Recommended Action

Configure your Application, Network, and Gateway Load Balancers with subnets across multiple Availability Zones.

Additional Resources

[Availability Zones for your Application Load Balancer](#)

[Availability Zones \(Elastic Load Balancing\)](#)[Create a Gateway Load Balancer](#)**Report columns**

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Auto Scaling available IPs in Subnets**Description**

Checks that sufficient available IPs remain among targeted Subnets. Having sufficient IPs available for use would help when Auto Scaling Group reaches its maximum size and needs to launch additional instances.

Check ID

Cjxm268ch1

Alert Criteria

- Red: The maximum number of instances and IP addresses that could be created by an ASG exceed the number of IP addresses remaining in the configured subnets.
- Green: There are sufficient IP addresses available for the remaining scale possible in the ASG.

Recommended Action

Increase the number of available IP addresses

Report columns

- Status
- Region
- Resource ARN
- Maximum instances that can be created
- Number of available instances

Auto Scaling Group Health Check

Description

Examines the health check configuration for Auto Scaling groups.

If Elastic Load Balancing is being used for an Auto Scaling group, the recommended configuration is to enable an Elastic Load Balancing health check. If an Elastic Load Balancing health check is not used, Auto Scaling can only act upon the health of the Amazon Elastic Compute Cloud (Amazon EC2) instance. Auto Scaling will not act on the application running on the instance.

Check ID

CLOG40CD08

Alert Criteria

- Yellow: An Auto Scaling group has an associated load balancer, but the Elastic Load Balancing health check is not enabled.
- Yellow: An Auto Scaling group does not have an associated load balancer, but the Elastic Load Balancing health check is enabled.

Recommended Action

If the Auto Scaling group has an associated load balancer, but the Elastic Load Balancing health check is not enabled, see [Add an Elastic Load Balancing Health Check to your Auto Scaling Group](#).

If the Elastic Load Balancing health check is enabled, but no load balancer is associated with the Auto Scaling group, see [Set Up an Auto-Scaled and Load-Balanced Application](#).

Additional Resources

[Amazon EC2 Auto Scaling User Guide](#)

Report columns

- Status
- Region
- Auto Scaling Group Name
- Load Balancer Associated
- Health Check

Auto Scaling Group Resources

Description

Checks the availability of resources associated with your launch configurations, launch templates, and your Auto Scaling groups.

Auto Scaling groups that point to unavailable resources cannot launch new Amazon Elastic Compute Cloud (Amazon EC2) instances. When properly configured, Auto Scaling causes the number of Amazon EC2 instances to increase seamlessly during demand spikes, and decrease automatically during demand lulls. Auto Scaling groups and launch configurations/launch templates that point to unavailable resources do not operate as intended.

Note

This check reports the resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.

Check ID

8CNsS11I5v

Alert Criteria

- Red: An Auto Scaling group is associated with a deleted load balancer.
- Red: A launch configuration is associated with a deleted Amazon Machine Image (AMI).
- Red: A launch template is associated with a deleted Amazon Machine Image (AMI).

Recommended Action

If the load balancer has been deleted, either create a new load balancer or target group and then associate it to the Auto Scaling group. or create a new Auto Scaling group without the load balancer. For information about creating a new Auto Scaling group with a new load balancer, see [Set Up an Auto-Scaled and Load-Balanced Application](#). For information about creating a new Auto Scaling group without a load balancer, see Create Auto Scaling Group in [Getting Started With Auto Scaling Using the Console](#).

If the AMI has been deleted, then create a new launch configuration or launch template version using a valid AMI and associate it with an Auto Scaling group. For information on how to create

a new launch configuration, see [Create a launch configuration](#) in the *Amazon EC2 Auto Scaling User Guide*. For information on how to create a launch template, see [Create a launch template for an Auto Scaling group](#) in the *Amazon EC2 Auto Scaling User Guide*.

Note

For security reasons, the check results don't include any resources referenced using AWS Systems Manager parameters in the launch template.

If your launch templates include an AWS Systems Manager parameter that include an Amazon Machine Image (AMI) ID , then review the launch template to make sure that the parameters reference a valid AMI ID, or make the appropriate changes in the AWS Systems Manager parameter store. For more information, see [Use AWS Systems Manager parameters instead of AMI IDs](#) in the *Amazon EC2 Auto Scaling User Guide*.

Additional Resources

- [Troubleshooting Auto Scaling: Amazon EC2 AMIs](#)
- [Troubleshooting Auto Scaling: Load Balancer Configuration](#)
- [Amazon EC2 Auto Scaling User Guide](#)
- [Use AWS Systems Manager parameters instead of AMI IDs](#)

Report columns

- Status
- Region
- Auto Scaling Group Name
- Launch Type
- Resource Type
- Resource Name

AWS CloudHSM clusters running HSM instances in a single AZ

Description

Checks your clusters that run HSM instances in a single Availability Zone (AZ). This check alerts you if your clusters are at risk of not having the most recent backup.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

hc0dfs7601

Alert Criteria

- Yellow: A CloudHSM cluster is running all HSM instances in a single Availability Zone for more than 1 hour.
- Green: A CloudHSM cluster is running all HSM instances in at least two different Availability Zones.

Recommended Action

Create at least one more instance for the cluster in a different Availability Zone.

Additional Resources

[Best practices for AWS CloudHSM](#)

Report columns

- Status
- Region
- Cluster ID
- Number of HSM Instances
- Last Updated Time

Direct Connect Location Resiliency**Description**

Checks the resilience of the Direct Connect used to connect your on-premises to each Direct Connect gateway or virtual private gateway.

This check alerts you if any Direct Connect gateway or virtual private gateway isn't configured with virtual interfaces across at least two distinct Direct Connect locations. Lack of location resiliency can result in unexpected downtime during maintenance, a fiber cut, a device failure, or a complete location failure.

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

 Note

Direct Connect is implemented with Transit Gateway using Direct Connect gateway.

Check ID

c1dfpnchv2

Alert Criteria

Red: The Direct Connect gateway or virtual private gateway is configured with one or more virtual interfaces on a single Direct Connect device.

Yellow: The Direct Connect gateway or virtual private gateway is configured with virtual interfaces across multiple Direct Connect devices in a single Direct Connect location.

Green: The Direct Connect gateway or virtual private gateway is configured with virtual interfaces across two or more distinct Direct Connect locations.

Recommended Action

To build Direct Connect location resiliency, you can configure the Direct Connect gateway or virtual private gateway to connect to at least two distinct Direct Connect locations. For more information, see [Direct Connect Resiliency Recommendation](#).

Additional Resources

[Direct Connect Resiliency Recommendations](#)

[Direct Connect Failover Test](#)

Report columns

- Status
- Region
- Last Updated Time
- Resiliency Status
- Location
- Connection ID
- Gateway ID

AWS Lambda functions without a dead-letter queue configured

Description

Checks if an AWS Lambda function is configured with a dead-letter queue.

A dead-letter queue is a feature of AWS Lambda that allows you to capture and analyze failed events, providing a way to handle those events accordingly. Your code might raise an exception, time out, or run out of memory, resulting in failed asynchronous executions of your Lambda function. A dead-letter queue stores messages from failed invocations, providing a way to handle the messages and troubleshoot the failures.

You can specify the dead-letter queue resource that you want to check using the **dlqArns** parameter in your AWS Config rules.

For more information, see [Dead-letter queues](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz182

Source

AWS Config Managed Rule: `lambda-dlq-check`

Alert Criteria

Yellow: AWS Lambda function has no dead-letter queue configured.

Recommended Action

Make sure that your AWS Lambda functions have a dead-letter queue configured to control message handling for all failed asynchronous invocations.

For more information, see [Dead-letter queues](#).

Additional Resources

- [Robust Serverless Application Design with AWS Lambda Dead Letter Queues](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

AWS Lambda On Failure Event Destinations

Description

Checks that Lambda functions in your account have On Failure event destination or Dead Letter Queue (DLQ) configured for asynchronous invocations, so that records from failed invocations can be routed to a destination for further investigation or processing.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1dfprch05

Alert Criteria

- Yellow: Function does not have any On Failure event destination or DLQ configured.

Recommended Action

Please set up On Failure event destination or DLQ for your Lambda functions to send failed invocations along with other details to one of the available destination AWS services for further debugging or processing.

Additional Resources

- [Asynchronous Invocation](#)
- [AWS Lambda On Failure Event Destinations](#)

Report columns

- Status
- Region
- The function with version which is flagged.
- Current day async requests dropped percentage
- Current day async requests
- Average daily async requests dropped percentage
- Average daily async requests
- Last Updated Time

AWS Lambda VPC-enabled Functions without Multi-AZ Redundancy

Description

Checks the \$LATEST version of VPC-enabled Lambda functions that are vulnerable to service interruption in a single Availability Zone. It's a best practice that VPC-enabled functions are connected to multiple Availability Zones for high availability.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

L4dfs2Q4C6

Alert Criteria

Yellow: The \$LATEST version of a VPC-enabled Lambda function is connected to subnets in a single Availability Zone.

Recommended Action

When configuring functions for access to your VPC, choose subnets in multiple Availability Zones to ensure high availability.

Additional Resources

- [Configuring a Lambda function to access resources in a VPC](#)
- [Resilience in AWS Lambda](#)

Report columns

- Status
- Region
- Function ARN
- VPC ID
- Average daily Invokes
- Last Updated Time

AWS Outposts Single Rack deployment

Description

Checks for Outposts Racks balance. This evaluates if a customers Outposts instances are deployed across multiple Outposts Racks or to a single Outpost Rack. A single Outposts rack creates a single point of failure for issues that involve a single Rack (for example, environmental failures). These scenarios can be mitigated by deploying outposts across multiple Racks.

Check ID

c243hjzrhn

Alert Criteria

- Yellow: Your Outpost is deployed on single Rack
- Green: Your Outpost is deployed across multiple Racks.

Recommended Action

If you are running production workloads on AWS Outposts, then its a best practice to use the following resilient architecture. A single AWS Outposts rack creates a single point of failure. Consider adding a second AWS Outposts rack to that location with enough capacity for a failover event, and then distribute workloads across racks.

Additional Resources

[Failure mode 4: Racks or data centers](#)

Report columns

- Status
- Resource ARN
- AZ
- Number of Racks
- Last Updated Time

AWS Resilience Hub Application Component check

Description

Checks if an Application Component (AppComponent) in your application is unrecoverable. If an AppComponent doesn't recover in the case of a disruption event, you might experience unknown data loss and system downtime.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

Check ID

RH23stmM04

Alert Criteria

Red: AppComponent is unrecoverable.

Recommended Action

To ensure that your AppComponent is recoverable, review and implement the resiliency recommendations, and then run a new assessment. For more information about reviewing the resiliency recommendations, see [Additional Resources](#).

Additional Resources

[Reviewing resiliency recommendations](#)

[AWS Resilience Hub concepts](#)

[AWS Resilience Hub User Guide](#)

Report columns

- Status
- Region
- Application Name
- AppComponent Name
- Last Updated Time

AWS Resilience Hub policy breached

Description

Checks Resilience Hub for applications that don't meet the recovery time objective (RTO) and recovery point objective (RPO) that the policy defines. The check alerts you if your application doesn't meet the RTO and RPO objectives you've set for an application in Resilience Hub.

Note

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

RH23stmM02

Alert Criteria

- Green: The application has a policy and meets the RTO and RPO objectives.
- Yellow: The application hasn't been assessed yet.
- Red: The application has a policy but doesn't meet the RTO and RPO objectives.

Recommended Action

Sign in to the Resilience Hub console and review the recommendations so that your application meets the RTO and RPO objectives.

Additional Resources

[Resilience Hub concepts](#)

Report columns

- Status

- Region
- Application Name
- Last Updated Time

AWS Resilience Hub resilience scores

Description

Checks if you have run an assessment for your applications in Resilience Hub. This check alerts you if your resilience scores are below a specific value.

Note

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

RH23stmM01

Alert Criteria

- Green: Your application has a resilience score of 70 or greater.
- Yellow: Your application has a resilience score of 40 through 69.
- Yellow: The application hasn't been assessed yet.
- Red: Your application has a resilience score of less than 40.

Recommended Action

Sign in to the Resilience Hub console and run an assessment for your application. Review the recommendations to improve the resilience score.

Additional Resources

[Resilience Hub concepts](#)

Report columns

- Status
- Region
- Application Name
- Application Resilience Score
- Last Updated Time

AWS Resilience Hub assessment age

Description

Checks how long since you last ran an application assessment. This check alerts you if you haven't run an application assessment for a specified number of days.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

RH23stmM03

Alert Criteria

- Green: Your application assessment ran in the last 30 days.
- Yellow: Your application assessment hasn't run in the last 30 days.

Recommended Action

Sign in to the Resilience Hub console and run an assessment for your application.

Additional Resources

[Resilience Hub concepts](#)

Report columns

- Status
- Region
- Application Name
- Days Since the Last Assessment Ran
- Last Assessment Run Time
- Last Updated Time

AWS Site-to-Site VPN has at least one tunnel in DOWN status

Description

Checks the number of tunnels that are active for each of your AWS Site-to-Site VPNs.

A VPN should have two tunnels configured at all times. This provides redundancy in case of outage or planned maintenance of the devices at the AWS endpoint. For some hardware, only one tunnel is active at a time. If a VPN has no active tunnels, charges for the VPN might still apply.

For more information, see [What is AWS Site-to-Site VPN?](#)

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz123

Source

AWS Config Managed Rule: `vpc-vpn-2-tunnels-up`

Alert Criteria

Yellow: A Site-to-Site VPN has at least one tunnel DOWN.

Recommended Action

Make sure that two tunnels are configured for VPN connections. And, if your hardware supports it, then make sure that both tunnels are active. If you no longer need a VPN connection, then delete it to avoid charges.

For more information, see [Your customer gateway device](#) and the content available on the [AWS Knowledge Center](#).

Additional Resources

- [AWS Site-to-Site VPN User Guide](#)
- [Adding a Virtual Private Gateway to Your VPC](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

AWS STS global endpoint usage across AWS Regions

Description

This check identifies a resiliency risk when AWS workloads make cross-region calls to the STS global endpoint in the US East (N. Virginia) Region. AWS workloads that make cross-region requests to the STS global endpoint in US East (N. Virginia) might be adversely impacted if connectivity to the US East (N. Virginia) STS endpoint is impacted.

In April 2025, AWS enhanced the global endpoint to automatically serve the STS global endpoint calls in the same Region as your workloads running in AWS, for all AWS Regions

enabled by default. However, some workloads, such as those running in opt-in Regions, or those not using Amazon DNS servers, do not benefit from the reduced latency and [fault isolation](#) offered as part of this enhancement. For these workloads, STS global endpoint requests are served in the US East (N. Virginia) Region, posing a resiliency risk to your application, and increasing latency of STS requests.

Note

Results for this check are automatically refreshed several times daily, each time covering the past 15 days. Refresh requests are not allowed. So, it might take up to 15 days for changes to appear in the Trusted Advisor check results.

For Business, Enterprise On-Ramp, or Enterprise Support customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c15m0mg1d3

Alert criteria

Red: Workloads are making cross-region calls to the AWS STS global endpoint in the US East (N. Virginia) Region.

Recommended action

To improve the resiliency and performance of your workloads, we recommend that you migrate from the STS global endpoint to the STS Regional endpoint. By using a Regional endpoint, you can use AWS STS in the same Region as your workloads.

The following is a list of AWS Identity and Access Management (IAM) Principals making cross-region calls to the AWS STS global endpoint in the US East (N. Virginia) Region. By following the steps in the blog post [How to use Regional AWS STS endpoints](#), you can reconfigure your workloads to use the Regional STS endpoint.

Status	Principal ARN	IAM Principal Type	Action	Last Updated Time	Originating Region
Red	arn:aws:iam::999999999999:role/example	Role	AssumeRole	2025-01-01T08:00.000Z	ca-west-1
Red	arn:aws:iam::123456789012:user/JohnDoe	User	GetCallerIdentity	2025-01-01T08:00.000Z	ca-west-1
Red	arn:aws:ssts::123456789012:federated-user/<Name>	Federated User	GetCallerIdentity	2025-01-01T08:00.000Z	eu-west-2
Red	arn:aws:iam::123456789012:root	Root	AssumeRole	2025-01-01T08:00.000Z	il-central-1
Red	arn:aws:ssts::0123456789012:assumed-role/aws-ec2-instance/i-0123456789example	EC2 Instance Identity Role	GetCallerIdentity	2025-01-01T08:00.000Z	us-west-1

Additional resources

- [Manage AWS STS in an AWS Region](#)
- [AWS STS Regional endpoints](#)

Report columns

- Status
- Principal Arn
- Action
- Principal Type
- Originating Region
- Last Updated Time

AWS Well-Architected high risk issues for reliability

Description

Checks for high risk issues (HRIs) for your workloads in the reliability pillar. This check is based on your AWS-Well Architected reviews. Your check results depend on whether you completed the workload evaluation with AWS Well-Architected.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

Wxdfp4B1L4

Alert Criteria

- Red: At least one active high risk issue was identified in the reliability pillar for AWS Well-Architected.

- Green: No active high risk issues were detected in the reliability pillar for AWS Well-Architected.

Recommended Action

AWS Well-Architected detected high risk issues during your workload evaluation. These issues present opportunities to reduce risk and save money. Sign in to the [AWS Well-Architected](#) tool to review your answers and take action to resolve your active issues.

Report columns

- Status
- Region
- Workload ARN
- Workload Name
- Reviewer Name
- Workload Type
- Workload Started Date
- Workload Last Modified Date
- Number of identified HRIs for Reliability
- Number of HRIs resolved for Reliability
- Number of questions answered for Reliability
- Total number of questions in Reliability pillar
- Last Updated Time

Classic Load Balancer has no multiple AZs configured

Description

Checks if Classic Load Balancer spans multiple Availability Zones (AZs).

A load balancer distributes incoming application traffic across multiple Amazon EC2 instances in multiple Availability Zones. By default, the load balancer distributes traffic evenly across the Availability Zones that you enable for your load balancer. If one Availability Zone experiences an outage, then load balancer nodes automatically forward requests to the healthy registered instances in one or more Availability Zones.

You can adjust the minimum number of Availability Zones using the **minAvailabilityZones** parameter in your AWS Config rules

For more information, see [What is a Classic Load Balancer?](#).

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz154

Source

AWS Config Managed Rule: clb-multiple-az

Alert Criteria

Yellow: Classic Load Balancer does not have Multi-AZ configured or does not meet the minimum number of AZs specified.

Recommended Action

Make sure that your Classic Load Balancers have multiple Availability Zones configured. Span your load balancer across multiple AZs to make sure that you have high availability of your application.

For more information, see [Tutorial: Create a Classic Load Balancer](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

CLB Connection Draining

Description

Checks for Classic load balancers that do not have connection draining enabled.

When connection draining is not enabled and you deregister an Amazon EC2 instance from a Classic load balancer, the Classic load balancer stops routing traffic to that instance and closes the connection. When connection draining is enabled, the Classic load balancer stops sending new requests to the deregistered instance but keeps the connection open to serve active requests.

Check ID

7qGXsKIUw

Alert Criteria

- Yellow: Connection draining is not enabled for a Classic load balancer.
- Green: Connection draining is enabled for Classic load balancer. .

Recommended Action

Enable connection draining for the Classic load balancer. For more information, see [Connection Draining](#) and [Enable or Disable Connection Draining for Your Load Balancer](#).

Additional Resources

[Elastic Load Balancing Concepts](#)

Report columns

- Status
- Region
- Load Balancer Name
- Reason

ELB Target Imbalance

Description

Checks the target groups' target distribution across Availability Zones (AZs) for Application Load Balancer (ALB), Network Load Balancer (NLB), and Gateway Load Balancer (GWLb).

This check excludes the following:

- Load balancers configured with a single Availability Zone (AZ).
- Load balancers where the difference in the number of targets between the most and least populated AZs is equal to or less than 1.
- Target groups with IP-based targets where the AvailabilityZone attribute is set to 'all'.

Check ID

b92b83d667

Alert Criteria

- Red: A single AZ represents more than 66% of the load balancer capacity.
- Yellow: A single AZ represents more than 50% of the load balancer capacity.
- Green: No AZs represents more than 50% of the load balancer capacity.

Recommended Action

For better resilience, make sure that your targets groups have same number of targets across AZs.

Additional Resources

[Target groups for your Application Load Balancers](#)

[Register targets with your Application Load Balancer target group](#)

Report columns

- Status
- Region
- Load Balancer Name
- Load Balancer Type
- Target Group ARN (arn)
- Difference in registered targets across AZs
- Last Updated Time

GWLB - endpoint AZ independence**Description**

Checks if your Gateway Load Balancer (GWLB) endpoints are configured as route destination from another Availability Zone (AZ).

Gateway Load Balancer endpoints forward network traffic to firewall appliances behind a Gateway Load Balancer for inspection. Each Gateway Load Balancer endpoint operates within a designated AZ and is built with redundancy in only that AZ. So any resources in a particular AZ must use a Gateway Load Balancer endpoint in the same AZ. This makes sure that any potential outage of a Gateway Load Balancer endpoint or its AZ doesn't impact your resources in another AZ.

Check ID

528d6f5ee7

Alert Criteria

- Yellow: Traffic from your subnet in one AZ is being routed through a Gateway Load Balancer endpoint in a different AZ.
- Green: Traffic from your subnet in one AZ is being routed through a Gateway Load Balancer endpoint in the same AZ.

Recommended Action

Check the AZ of your subnet and configure its route table to route traffic through a Gateway Load Balancer endpoint in the same AZ.

If there is no Gateway Load Balancer endpoint in the AZ, then create a new one and then route your subnet traffic through it.

If you have the same route table associated across subnets in different AZs, then keep this route table associated to the subnets that reside in the same AZ as the Gateway Load Balancer endpoint. For subnets in the other AZ, you can then associate a separate route table with a route to a Gateway Load Balancer endpoint in the this AZ.

It's a best practice to choose a maintenance window for architecture changes in your Amazon VPC.

Additional Resources

- [Availability Zone Independence](#)
- [Configure Routing for Gateway Load Balancer endpoint](#)
- [AWS Well-Architected Tool - Use bulkhead architectures to limit scope of impact](#)

Report columns

- Status
- Region

- Cross AZ Subnet Id List
- Gateway Load Balancer Endpoint Id
- Gateway Load Balancer Endpoint Subnet Id
- VPC Endpoint Subnet AZ
- Last Updated Time

Load Balancer Optimization

Description

Checks your load balancer configuration.

To help increase the level of fault tolerance in Amazon Elastic Compute Cloud (Amazon EC2) when using Elastic Load Balancing, we recommend running an equal number of instances across multiple Availability Zones in a Region. A load balancer that is configured accrues charges, so this is a cost-optimization check as well.

Check ID

iqdCTZKCUp

Alert Criteria

- Yellow: A load balancer is enabled for a single Availability Zone.
- Yellow: A load balancer is enabled for an Availability Zone that has no active instances.
- Yellow: The Amazon EC2 instances that are registered with a load balancer are unevenly distributed across Availability Zones. (The difference between the highest and lowest instance counts in utilized Availability Zones is more than 1, and the difference is more than 20% of the highest count.)

Recommended Action

Ensure that your load balancer points to active and healthy instances in at least two Availability Zones. For more information, see [Add Availability Zone](#).

If your load balancer is configured for an Availability Zone with no healthy instances, or if there is an imbalance of instances across the Availability Zones, determine if all the Availability Zones are necessary. Omit any unnecessary Availability Zones and ensure there is a balanced distribution of instances across the remaining Availability Zones. For more information, see [Remove Availability Zone](#).

Additional Resources

- [Availability Zones and Regions](#)
- [Managing Load Balancers](#)
- [Best Practices in Evaluating Elastic Load Balancing](#)

Report columns

- Status
- Region
- Load Balancer Name
- # of Zones
- Zone a Instances
- Zone b Instances
- Zone c Instances
- Zone d Instances
- Zone e Instances
- Zone f Instances
- Reason

NAT Gateway AZ Independence

Description

Checks if your NAT Gateways are configured with Availability Zone (AZ) independence.

A NAT Gateway enables resources in your private subnet to securely connect to services outside the subnet using the NAT Gateway's IP addresses and drops any unsolicited inbound traffic. Each NAT Gateway operates within a designated Availability Zone (AZ) and is built with redundancy in that AZ only. Therefore, your resources in a particular AZ should use a NAT Gateway in the same AZ so that any potential outage of a NAT Gateway or its AZ does not impact your resources in another AZ.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1dfptbg10

Alert Criteria

- Red: Traffic from your subnet in one AZ is being routed through a NATGW in a different AZ.
- Green: Traffic from your subnet in one AZ is being routed through a NATGW in the same AZ.

Recommended Action

Please check the AZ of your subnet and route traffic through a NAT Gateway in the same AZ.

If there is no NATGW in the AZ, please create one and then route your subnet traffic through it.

If you have the same route table associated across subnets in different AZs, keep this route table associated to the subnets that reside in the same AZ as the NAT Gateway and for subnets in the other AZ, please associate a separate route table with a route to a NAT Gateway in this other AZ.

We recommend choosing a maintenance window for architecture changes in your Amazon VPC.

Additional Resources

- [How to create a NAT Gateway](#)
- [How to configure routes for different NAT Gateway use cases](#)

Report columns

- Status
- Region
- NAT Availability Zone
- NAT ID
- Subnet Availability Zone
- Subnet ID
- Route Table ID

- NAT ARN
- Last Updated Time

Network Firewall endpoint AZ Independence

Description

Checks if your AWS Network Firewall endpoints are configured as a route destination from another Availability Zone (AZ).

Network Firewall endpoints forward network traffic to a Network Firewall for inspection. Each Network Firewall endpoint operates within a designated AZ and is built with redundancy only in that AZ. Your resources in a particular AZ should use a Network Firewall endpoint in the same AZ. This makes sure that any potential outage of a Network Firewall endpoint or its AZ doesn't impact your resources in another AZ. Network traffic that originates in a different AZ for traffic inspection incurs cross-AZ data transfer charges. It's a best practice to make sure that all resources in a specific AZ use a Network Firewall in the same AZ to avoid cross-AZ data charges.

Check ID

7040ea389a

Alert Criteria

- Yellow: Traffic from a subnet in one AZ is being routed through a Network Firewall endpoint in a different AZ.
- Green: Traffic from a subnet in one AZ is being routed through a Network Firewall endpoint in the same AZ.

Recommended Action

Check the AZ of your subnet and route traffic through a Network Firewall endpoint in the same AZ.

If there is no Network Firewall endpoint in the AZ, then create a new Network Firewall and route your subnet traffic through it.

If the same route table is associated across multiple subnets in different AZs, then keep this route table associated to the subnets that reside in the same AZ as the Network Firewall endpoint. For subnets in other AZs, associate a separate route table with a route to a Network Firewall endpoint in that AZ.

It's a best practice to choose a maintenance window for architecture changes in your Amazon VPC.

Additional Resources

[Data Transfer within the same AWS Region](#)

[Understanding data transfer charges](#)

[Availability Zone Independence](#)

[High Level steps for implementing a firewall](#)

[Creating a firewall](#)

[AWS Well-Architected Tool - Use bulkhead architectures to limit scope of impact](#)

Report columns

- Status
- Region
- Network Firewall Endpoint Id
- Network Firewall Arn
- Network Firewall Endpoint Subnet
- Network Firewall Endpoint AZ
- Cross AZ Subnets List
- Last Updated Time

Network Firewall Multi-AZ

Description

Checks if your Network Firewalls are configured to use more than one Availability Zone (AZ) for firewall endpoints.

An AZ is a distinct location that's insulated from failures in other zones. If the Network Firewall endpoint is deployed in only 1 AZ, then it can be a single point of failure and can impair workloads from other AZs using the Network Firewall for traffic inspection. It's a best practice to configure your Network Firewalls in multiple AZs in the same Region to improve your workload availability.

Check ID

c2v1fg0gqd

Alert Criteria

- Yellow: Network Firewall endpoint is deployed in 1 AZ.
- Green: Network Firewall endpoints is deployed in at least two AZs.

Recommended Action

Make sure that your Network Firewall is configured with at least two AZs for production workloads.

Additional Resources

[VPC subnet configuration for AWS Network Firewall](#)

[Creating a firewall](#)

[Availability Zone](#)

[AWS Well-Architected Tool - Deploy the workload to multiple locations](#)

[Appliance in a shared services VPC](#)

Report columns

- Status
- Region
- Network Firewall Arn
- VPC Id
- Network Firewall Subnets
- Network Firewall Subnets AZs
- Last Updated Time

Network Load Balancers Cross Load Balancing

Description

Checks if cross-zone load balancing is enabled on Network Load Balancers.

Cross-zone load balancing helps maintain even distribution of incoming traffic across instances in different Availability Zones. This prevents the load balancer from routing all traffic to

instances in the same Availability Zone, which can cause uneven traffic distribution and potential overloading. The feature also helps application reliability by automatically routing traffic to healthy instances in other Availability Zones in the event of a single Availability Zone failure.

For more information, see [Cross-zone load balancing](#).

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz105

Source

AWS Config Managed Rule: nlb-cross-zone-load-balancing-enabled

Alert Criteria

- Yellow: Network Load Balancer does not have cross-zone load balancing enabled.

Recommended Action

Ensure that cross-zone load balancing is enabled on Network Load Balancers.

Additional Resources

[Cross-zone load balancing \(Network Load Balancers\)](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

NLB - Internet-facing resource in private subnet

Description

Checks if an internet-facing Network Load Balancer (NLB) is configured with a private subnet. An internet-facing Network Load Balancer (NLB) must be configured in public subnets in order to receive traffic. A public subnet is defined as a subnet that has a direct route to an [internet gateway](#). If the subnet is configured as private, then its Availability Zone (AZ) doesn't receive traffic, which can cause availability issues.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1dfpnchv4

Alert Criteria

Red: NLB is configured with one or more private subnets

Green: No private subnet is configured for internet-facing NLB

Recommended Action

Confirm that the subnets configured in an internet-facing load balancer are public. A public subnet is defined as a subnet that has a direct route to an [internet gateway](#). Use one of following options:

- Create a new load balancer and select a different subnet with a direct route to an internet gateway.
- Change the subnet that's currently attached to the load balancer from private to public. To do this, change its route table and [associate an internet gateway](#).

Additional Resources

- [Configure a load balancer and a listener](#)

- [Subnets for your VPC](#)
- [Associate a gateway with a route table](#)

Report columns

- Status
- Region
- NLB Arn
- NLB Name
- Subnet ID
- NLB Scheme
- Subnet Type
- Last Updated Time

NLB Multi-AZ

Description

Checks if your Network Load Balancers are configured to use more than one Availability Zone (AZ). An AZ is a distinct location that is insulated from failures in other zones. Configure your load balancer in multiple AZs in the same Region to help improve your workload availability.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1dfprch09

Alert Criteria

Yellow: NLB is in a single AZ.

Green: NLB has two or more AZs.

Recommended Action

Make sure that your load balancer is configured with at least two Availability Zones.

Additional Resources

For more information, see the following documentation:

- [Availability Zones](#)
- [AWS Well-Architected - Deploy the workload to multiple locations](#)
- [Regions and Availability Zones](#)

Report columns

- Status
- Region
- Number of AZs
- NLB ARN
- NLB Name
- Last Updated Time

Number of AWS Regions in an Incident Manager replication set

Description

Checks that an Incident Manager replication set's configuration uses more than one AWS Region to support regional failover and response. For incidents created by CloudWatch alarms or EventBridge events, Incident Manager creates an incident in the same AWS Region as the alarm or event rule. If Incident Manager is temporarily unavailable in that Region, the system attempts to create an incident in another Region in the replication set. If the replication set includes only one Region, the system fails to create an incident record while Incident Manager is unavailable.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

cIdfp1js9r

Alert Criteria

- Green: The replication set contains more than one Region.
- Yellow: The replication set contains one Region.

Recommended Action

Add at least one more Region to the replication set.

Additional Resources

For more information, see [Cross-region Incident management](#).

Report columns

- Status
- Multi-region
- Replication Set
- Last Updated Time

Single AZ Application Check

Description

Checks through network patterns if your egress network traffic is routing through a single Availability Zone (AZ).

An AZ is a distinct location that is insulated from any impact in other zones. By spreading your service across multiple AZs, you limit the blast radius of an AZ failure.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1dfptbg11

Alert Criteria

- Yellow: Your application may be deployed in only one AZ based on observed egress network patterns. If this is true and your application expects high availability, we recommend that you provision your application resources and implement your network flows to utilize multiple Availability Zones.

Recommended Action

If your application requires high availability, consider implementing a multi-AZ architecture for higher availability.

Report columns

- Status
- Region
- VPC ID
- Last Updated Time

VPC interface endpoint network interfaces in multiple AZs

Description

Checks if your AWS PrivateLink VPC interface endpoints are configured to use more than one Availability Zone (AZ). An AZ is a distinct location that is insulated from failures in other zones. This supports inexpensive, low-latency network connectivity between AZs in the same AWS Region. Select subnets in multiple AZs when you create interface endpoints to help protect your applications from a single point of failure.

Note

This check currently includes only interface endpoints.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1dfprch10

Alert Criteria

Yellow: VPC endpoint is in a single AZ.

Green: VPC endpoint is in at least two AZs.

Recommended Action

Make sure that your VPC interface endpoint is configured with at least two Availability Zones.

Additional Resources

For more information, see the following documentation:

- [Access an AWS service using an interface VPC endpoint](#)
- [Private IP address of the endpoint network interface](#)
- [AWS PrivateLink concepts](#)
- [Regions and Availability Zones](#)

Report columns

- Status
- Region
- VPC Endpoint ID
- Is Multi AZ
- Last Updated Time

VPN Tunnel Redundancy

Description

Checks the number of tunnels that are active for each of your Site-to-Site VPNs.

A VPN should have two tunnels configured at all times. This provides redundancy in case of outage or planned maintenance of the devices at the AWS endpoint. For some hardware, only one tunnel is active at a time. If a VPN has no active tunnels, charges for the VPN might still apply. For more information, see [AWS Site-to-Site VPN User Guide](#).

Check ID

S45wrEXrLz

Alert Criteria

- Yellow: A VPN has one active tunnel (this is normal for some hardware).
- Yellow: A VPN has no active tunnels.

Recommended Action

Be sure that two tunnels are configured for your VPN connection, and that both are active if your hardware supports it. If you no longer need a VPN connection, you can delete it to avoid charges. For more information, see [Your customer gateway device](#) or [Delete a Site-to-Site VPN connection](#).

Additional Resources

- [AWS Site-to-Site VPN User Guide](#)
- [Create a target gateway](#)

Report columns

- Status
- Region
- VPN ID
- VPC
- Virtual Private Gateway
- Customer Gateway
- Active Tunnels

- Reason

ActiveMQ Availability Zone Redundancy

Description

Checks that Amazon MQ for ActiveMQ brokers are configured for high availability with an active/standby broker in multiple Availability Zones.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1t3k8mqv1

Alert Criteria

- Yellow: An Amazon MQ for ActiveMQ broker is configured in a single Availability Zone.

Green: An Amazon MQ for ActiveMQ broker is configured in at least two Availability Zones.

Recommended Action

Create a new broker with active/standby deployment mode.

Additional Resources

- [Creating an ActiveMQ broker](#)

Report columns

- Status
- Region
- ActiveMQ Broker ID
- Broker Engine Type

- Deployment Mode
- Last Updated Time

RabbitMQ Availability Zone Redundancy

Description

Checks that Amazon MQ for RabbitMQ brokers are configured for high availability with cluster instances in multiple Availability Zones.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1t3k8mqv2

Alert Criteria

- Yellow: An Amazon MQ for RabbitMQ broker is configured in a single Availability Zone.
Green: An Amazon MQ for RabbitMQ broker is configured in multiple Availability Zones.

Recommended Action

Create a new broker with the cluster deployment mode.

Additional Resources

- [Creating a RabbitMQ broker](#)

Report columns

- Status
- Region
- RabbitMQ Broker ID
- Broker Engine Type

- Deployment Mode
- Last Updated Time

Service limits

See the following checks for the service limits (also known as quotas) category.

All checks in this category have the following descriptions:

Alert Criteria

- Yellow: 80% of limit reached.
- Red: 100% of limit reached.
- Blue: Trusted Advisor was unable to retrieve utilization or limits in one or more AWS Regions.

Recommended Action

If you expect to exceed a service limit, request an increase directly from the [Service Quotas](#) console. If Service Quotas doesn't support your service yet, you can open a support case in [Support Center](#).

Report columns

- Status
- Service
- Region
- Limit Amount
- Current Usage

Note

- Values are based on a snapshot, so your current usage might differ. Quota and usage data can take up to 24 hours to reflect any changes. In cases where quotas have been recently increased, you might temporarily see utilization that exceeds the quota.

Check names

- [Auto Scaling Groups](#)

- [Auto Scaling Launch Configurations](#)
- [CloudFormation Stacks](#)
- [DynamoDB Read Capacity](#)
- [DynamoDB Write Capacity](#)
- [EBS Active Snapshots](#)
- [EBS Cold HDD \(sc1\) Volume Storage](#)
- [EBS General Purpose SSD \(gp2\) Volume Storage](#)
- [EBS General Purpose SSD \(gp3\) Volume Storage](#)
- [EBS Magnetic \(standard\) Volume Storage](#)
- [EBS Provisioned IOPS SSD \(io1\) Volume Aggregate IOPS](#)
- [EBS Provisioned IOPS SSD \(io1\) Volume Storage](#)
- [EBS Provisioned IOPS SSD \(io2\) Volume Storage](#)
- [EBS Throughput Optimized HDD \(st1\) Volume Storage](#)
- [EC2 On-Demand Instances](#)
- [EC2 Reserved Instance Leases](#)
- [EC2-Classic Elastic IP Addresses](#)
- [EC2-VPC Elastic IP Address](#)
- [ELB Application Load Balancers](#)
- [ELB Classic Load Balancers](#)
- [ELB Network Load Balancers](#)
- [IAM Group](#)
- [IAM Instance Profiles](#)
- [IAM Policies](#)
- [IAM Roles](#)
- [IAM Server Certificates](#)
- [IAM Users](#)
- [Kinesis Shards per Region](#)
- [Lambda Code Storage Usage](#)
- [RDS Cluster Parameter Groups](#)

- [RDS Cluster Roles](#)
- [RDS Clusters](#)
- [RDS DB Instances](#)
- [RDS DB Manual Snapshots](#)
- [RDS DB Parameter Groups](#)
- [RDS DB Security Groups](#)
- [RDS Event Subscriptions](#)
- [RDS Max Auths per Security Group](#)
- [RDS Option Groups](#)
- [RDS Read Replicas per Master](#)
- [RDS Reserved Instances](#)
- [RDS Subnet Groups](#)
- [RDS Subnets per Subnet Group](#)
- [RDS Total Storage Quota](#)
- [Route 53 Hosted Zones](#)
- [Route 53 Max Health Checks](#)
- [Route 53 Reusable Delegation Sets](#)
- [Route 53 Traffic Policies](#)
- [Route 53 Traffic Policy Instances](#)
- [SES Daily Sending Quota](#)
- [VPC](#)
- [VPC Internet Gateways](#)

Auto Scaling Groups

Description

Checks for usage that is more than 80% of the Auto Scaling Groups quota.

Check ID

fw7HH017J9

Additional Resources

[Auto Scaling quotas](#)

Auto Scaling Launch Configurations

Description

Checks for usage that is more than 80% of the Auto Scaling launch configurations quota.

Check ID

aW7HH017J9

Additional Resources

[Auto Scaling quotas](#)

CloudFormation Stacks

Description

Checks for usage that is more than 80% of the CloudFormation stacks quota.

Check ID

gW7HH017J9

Additional Resources

[CloudFormation quotas](#)

DynamoDB Read Capacity

Description

Checks for usage that is more than 80% of the DynamoDB provisioned throughput limit for reads per AWS account.

Check ID

6gtQddfEw6

Additional Resources

[DynamoDB quotas](#)

DynamoDB Write Capacity

Description

Checks for usage that is more than 80% of the DynamoDB provisioned throughput limit for writes per AWS account.

Check ID

c5ftjdfkMr

Additional Resources

[DynamoDB quotas](#)

EBS Active Snapshots

Description

Checks for usage that is more than 80% of the EBS active snapshots quota.

Check ID

eI7KK017J9

Additional Resources

[Amazon EBS limits](#)

EBS Cold HDD (sc1) Volume Storage

Description

Checks for usage that is more than 80% of the EBS Cold HDD (sc1) volume storage quota.

Check ID

gH5CC0e3J9

Additional Resources

[Amazon EBS limits](#)

EBS General Purpose SSD (gp2) Volume Storage

Description

Checks for usage that is more than 80% of the EBS General Purpose SSD (gp2) volume storage quota.

Check ID

dH7RR016J9

Additional Resources

[Amazon EBS limits](#)

EBS General Purpose SSD (gp3) Volume Storage

Description

Checks for usage that is more than 80% of the EBS General Purpose SSD (gp3) volume storage quota.

Check ID

dH7RR016J3

Additional Resources

[Amazon EBS limits](#)

EBS Magnetic (standard) Volume Storage

Description

Checks for usage that is more than 80% of the EBS Magnetic (standard) volume storage quota.

Check ID

cG7HH017J9

Additional Resources

[Amazon EBS limits](#)

EBS Provisioned IOPS SSD (io1) Volume Aggregate IOPS

Description

Checks for usage that is more than 80% of the EBS Provisioned IOPS SSD (io1) volume aggregate IOPS quota.

Check ID

tV7YY017J9

Additional Resources

[Amazon EBS limits](#)

EBS Provisioned IOPS SSD (io1) Volume Storage

Description

Checks for usage that is more than 80% of the EBS Provisioned IOPS SSD (io1) volume storage quota.

Check ID

gI7MM017J9

Additional Resources

[Amazon EBS limits](#)

EBS Provisioned IOPS SSD (io2) Volume Storage

Description

Checks for usage that is more than 80% of the EBS Provisioned IOPS SSD (io2) volume storage quota.

Check ID

gI7MM017J2

Additional Resources

[Amazon EBS limits](#)

EBS Throughput Optimized HDD (st1) Volume Storage

Description

Checks for usage that is more than 80% of the EBS Throughput Optimized HDD (st1) volume storage quota.

Check ID

wH7DD013J9

Additional Resources

[Amazon EBS limits](#)

EC2 On-Demand Instances

Description

Checks for usage that is more than 80% of the EC2 On-Demand Instances quota.

Check ID

0Xc6LMYG8P

Additional Resources

[Amazon EC2 quotas](#)

EC2 Reserved Instance Leases

Description

Checks for usage that is more than 80% of the EC2 Reserved Instance leases quota.

Check ID

iH7PP017J9

Additional Resources

[Amazon EC2 quotas](#)

EC2-Classic Elastic IP Addresses

Description

Checks for usage that is more than 80% of the EC2-Classic Elastic IP addresses quota.

Check ID

aW9HH018J6

Additional Resources

[Amazon EC2 quotas](#)

EC2-VPC Elastic IP Address

Description

Checks for usage that is more than 80% of the EC2-VPC Elastic IP address quota.

Check ID

1N7RR017J9

Additional Resources

[VPC Elastic IP quotas](#)

ELB Application Load Balancers

Description

Checks for usage that is more than 80% of the ELB Application Load Balancers quota.

Check ID

EM8b3yLRTx

Additional Resources

[Elastic Load Balancing quotas](#)

ELB Classic Load Balancers

Description

Checks for usage that is more than 80% of the ELB Classic Load Balancers quota.

Check ID

iK700017J9

Additional Resources

[Elastic Load Balancing quotas](#)

ELB Network Load Balancers

Description

Checks for usage that is more than 80% of the ELB Network Load Balancers quota.

Check ID

8wIqYSt25K

Additional Resources

[Elastic Load Balancing quotas](#)

IAM Group

Description

Checks for usage that is more than 80% of the IAM group quota.

Check ID

sU7XX017J9

Additional Resources

[IAM quotas](#)

IAM Instance Profiles

Description

Checks for usage that is more than 80% of the IAM instance profiles quota.

Check ID

n07SS017J9

Additional Resources

[IAM quotas](#)

IAM Policies

Description

Checks for usage that is more than 80% of the IAM policies quota.

Check ID

pR7UU017J9

Additional Resources

[IAM quotas](#)

IAM Roles

Description

Checks for usage that is more than 80% of the IAM roles quota.

Check ID

oQ7TT017J9

Additional Resources

[IAM quotas](#)

IAM Server Certificates

Description

Checks for usage that is more than 80% of the IAM server certificates quota.

Check ID

rT7WW017J9

Additional Resources

[IAM quotas](#)

IAM Users

Description

Checks for usage that is more than 80% of the IAM users quota.

Check ID

qS7VV017J9

Additional Resources

[IAM quotas](#)

Kinesis Shards per Region

Description

Checks for usage that is more than 80% of the Kinesis shards per Region quota.

Check ID

bW7HH017J9

Additional Resources

[Kinesis quotas](#)

Lambda Code Storage Usage

Description

Checks for code storage usage that is more than 80% of the account limit.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c1dfprch07

Alert Criteria

- Yellow: 80% of limit reached.

Recommended Action

Please identify unused lambda functions or versions and remove them to free up the code storage for your account in the region. If you need additional storage, please create a support case in Support Center. If you expect to exceed a service limit, request an increase directly from the Service Quotas console. If Service Quotas doesn't support your service yet, you can open a support case in Support Center.

Additional Resources

- [Lambda Code Storage Usage](#)

Report columns

- Status
- Region
- The qualified function ARN for this resource.
- The function code storage usage in MegaBytes with 2 decimals.
- The amount of versions in the function

- Last Updated Time

RDS Cluster Parameter Groups

Description

Checks for usage that is more than 80% of the RDS cluster parameter groups quota.

Check ID

jt1IM03qZM

Additional Resources

[Amazon RDS quotas](#)

RDS Cluster Roles

Description

Checks for usage that is more than 80% of the RDS cluster roles quota.

Check ID

7fuccf1Mx7

Additional Resources

[Amazon RDS quotas](#)

RDS Clusters

Description

Checks for usage that is more than 80% of the RDS clusters quota.

Check ID

gjqMBn6pjz

Additional Resources

[Amazon RDS quotas](#)

RDS DB Instances

Description

Checks for usage that is more than 80% of the RDS DB instances quota.

Check ID

XG0aXHpIEt

Additional Resources

[Amazon RDS quotas](#)

RDS DB Manual Snapshots

Description

Checks for usage that is more than 80% of the RDS DB manual snapshots quota.

Check ID

dV84wpqRUs

Additional Resources

[Amazon RDS quotas](#)

RDS DB Parameter Groups

Description

Checks for usage that is more than 80% of the RDS DB parameter groups quota.

Check ID

jEECYg2YVU

Additional Resources

[Amazon RDS quotas](#)

RDS DB Security Groups

Description

Checks for usage that is more than 80% of the RDS DB security groups quota.

Check ID

gfZAn3W7w1

Additional Resources

[Amazon RDS quotas](#)

RDS Event Subscriptions

Description

Checks for usage that is more than 80% of the RDS event subscriptions quota.

Check ID

keAhfbH5yb

Additional Resources

[Amazon RDS quotas](#)

RDS Max Auths per Security Group

Description

Checks for usage that is more than 80% of the RDS max auths per security group quota.

Check ID

dBkuNCvqn5

Additional Resources

[Amazon RDS quotas](#)

RDS Option Groups

Description

Checks for usage that is more than 80% of the RDS option groups quota.

Check ID

3Njm0DJQ09

Additional Resources

[Amazon RDS quotas](#)

RDS Read Replicas per Master

Description

Checks for usage that is more than 80% of the RDS read replicas per master quota.

Check ID

pYW8UkYz2w

Additional Resources

[Amazon RDS quotas](#)

RDS Reserved Instances

Description

Checks for usage that is more than 80% of the RDS Reserved Instances quota.

Check ID

UUDv0a5r34

Additional Resources

[Amazon RDS quotas](#)

RDS Subnet Groups

Description

Checks for usage that is more than 80% of the RDS subnet groups quota.

Check ID

dYWBaXaaMM

Additional Resources

[Amazon RDS quotas](#)

RDS Subnets per Subnet Group

Description

Checks for usage that is more than 80% of the RDS subnets per subnet group quota.

Check ID

jEhCtdJK0Y

Additional Resources

[Amazon RDS quotas](#)

RDS Total Storage Quota

Description

Checks for usage that is more than 80% of the RDS total storage quota.

Check ID

P1jhKWEMLa

Additional Resources

[Amazon RDS quotas](#)

Route 53 Hosted Zones

Description

Checks for usage that is more than 80% of the Route 53 hosted zones quota per account.

Check ID

dx3xfcdfMr

Additional Resources

[Route 53 quotas](#)

Route 53 Max Health Checks

Description

Checks for usage that is more than 80% of the Route 53 health checks quota per account.

Check ID

ru4xfcdfMr

Additional Resources

[Route 53 quotas](#)

Route 53 Reusable Delegation Sets

Description

Checks for usage that is more than 80% of the Route 53 reusable delegation sets quota per account.

Check ID

ty3xfcdfMr

Additional Resources

[Route 53 quotas](#)

Route 53 Traffic Policies

Description

Checks for usage that is more than 80% of the Route 53 traffic policies quota per account.

Check ID

dx3xfbjfMr

Additional Resources

[Route 53 quotas](#)

Route 53 Traffic Policy Instances

Description

Checks for usage that is more than 80% of the Route 53 traffic policy instances quota per account.

Check ID

dx8afcdfMr

Additional Resources

[Route 53 quotas](#)

SES Daily Sending Quota

Description

Checks for usage that is more than 80% of the Amazon SES daily sending quota.

Check ID

hJ7NN017J9

Additional Resources

[Amazon SES quotas](#)

VPC

Description

Checks for usage that is more than 80% of the VPC quota.

Check ID

jL7PP017J9

Additional Resources

[VPC quotas](#)

VPC Internet Gateways

Description

Checks for usage that is more than 80% of the VPC Internet gateways quota.

Check ID

kM7QQ017J9

Additional Resources

[VPC quotas](#)

Operational Excellence

You can use the following checks for the operational excellence category.

Check names

- [Amazon API Gateway Not Logging Execution Logs](#)
- [Amazon API Gateway REST APIs Without X-Ray Tracing Enabled](#)
- [Amazon CloudFront Access Log Configured](#)
- [Amazon CloudWatch Alarm Action is Disabled](#)
- [Amazon EC2 Instance Not Managed by AWS Systems Manager](#)
- [Amazon ECR Repository With Tag Immutability Disabled](#)
- [Amazon ECS clusters with Container Insights disabled](#)

- [Amazon ECS task logging not enabled](#)
- [Amazon OpenSearch Service logging CloudWatch not configured](#)
- [Amazon RDS DB instances in the clusters with heterogeneous parameter groups](#)
- [Amazon RDS Enhanced Monitoring is turned off](#)
- [Amazon RDS Performance Insights is turned off](#)
- [Amazon RDS track_counts parameter is turned off](#)
- [Amazon Redshift cluster audit logging](#)
- [Amazon S3 Access Logs Enabled](#)
- [Amazon S3 does not have Event Notifications enabled](#)
- [Amazon SNS Topics Not Logging Message Delivery Status](#)
- [Amazon VPC Without Flow Logs](#)
- [Application Load Balancers and Classic Load Balancers Without Access Logs Enabled](#)
- [CloudFormation Stack Notification](#)
- [AWS CloudTrail data events logging for objects in an S3 bucket](#)
- [AWS CodeBuild Project Logging](#)
- [AWS CodeDeploy Auto Rollback and Monitor Enabled](#)
- [AWS CodeDeploy Lambda is using all-at-once deployment configuration](#)
- [AWS Elastic Beanstalk Enhanced Health Reporting is not Configured](#)
- [AWS Elastic Beanstalk with Managed Platform Updates Disabled](#)
- [AWS Fargate platform version is not latest](#)
- [AWS Systems Manager State Manager Association in Non-compliant Status](#)
- [CloudTrail trails are not configured with Amazon CloudWatch Logs](#)
- [Elastic Load Balancing Deletion Protection Not Enabled for Load Balancers](#)
- [NitroV4 Amazon EC2 instances running outdated ENA driver](#)
- [RDS DB Cluster Deletion Protection Check](#)
- [RDS DB Instance Automatic Minor Version Upgrade Check](#)

Amazon API Gateway Not Logging Execution Logs

Description

Checks if Amazon API Gateway has CloudWatch Logs turned on at the desired logging level.

Turn on CloudWatch logging for REST API methods or WebSocket API routes in Amazon API Gateway to collect execution logs in CloudWatch Logs for requests received by your APIs. The information contained in the execution logs helps identify and troubleshoot issues related to your API.

You can specify the logging level (ERROR, INFO) ID in the **loggingLevel** parameter in the AWS Config rules.

Refer to the REST API or WebSocket API documentation for more information about CloudWatch logging in Amazon API Gateway.

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz125

Source

AWS Config Managed Rule: `api-gw-execution-logging-enabled`

Alert Criteria

Yellow: The CloudWatch logging setting for execution log collection isn't enabled at the desired logging level for an Amazon API Gateway.

Recommended Action

Turn on CloudWatch logging for execution logs for your Amazon API Gateway [REST APIs](#) or [WebSocket APIs](#) with the appropriate logging level (ERROR, INFO).

For more information, see [Create a flow log](#)

Additional Resources

- [Setting up CloudWatch logging for a REST API in API Gateway](#)

- [Configuring logging for a WebSocket API](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon API Gateway REST APIs Without X-Ray Tracing Enabled

Description

Checks if Amazon API Gateway REST APIs have AWS X-Ray tracing turned on.

Turn on X-Ray tracing for your REST APIs to allow API Gateway to sample API invocation requests with trace information. This allows you to take advantage of AWS X-Ray to trace and analyze requests as they travel through your API Gateway REST APIs to the downstream services.

For more information, see [Tracing user requests to REST APIs using X-Ray](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz126

Source

AWS Config Managed Rule: `api-gw-xray-enabled`

Alert Criteria

Yellow: X-Ray tracing is not turned on for an API Gateway REST API.

Recommended Action

Turn on X-Ray tracing for your API Gateway REST APIs.

For more information, see [Setting up AWS X-Ray with API Gateway REST APIs](#).

Additional Resources

- [Tracing user requests to REST APIs using X-Ray](#)
- [What is AWS X-Ray?](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon CloudFront Access Log Configured

Description

Checks if Amazon CloudFront distributions are configured to capture information from Amazon S3 server access logs. Amazon S3 server access logs contain detailed information about every user request that CloudFront receives.

You can adjust the the name of the Amazon S3 bucket for storing server access logs, using the **S3BucketName** parameter in your AWS Config rules.

For more information, see [Configuring and using standard logs \(access logs\)](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz110

Source

AWS Config Managed Rule: cloudfront-accesslogs-enabled

Alert Criteria

Yellow: Amazon CloudFront access logging is not enabled

Recommended Action

Make sure that you turn on CloudFront access logging to capture detailed information about every user request that CloudFront receives.

You can turn on standard logs when you create or update a distribution.

For more information, see [Values that you specify when you create or update a distribution](#).

Additional Resources

- [Values that you specify when you create or update a distribution](#)
- [Configuring and using standard logs \(access logs\)](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon CloudWatch Alarm Action is Disabled

Description

Checks if your Amazon CloudWatch alarm action is in a disabled state.

You can use the AWS CLI to enable or disable the action feature in your alarm. Or, you can programmatically disable or enable the action feature using the AWS SDK. When the alarm action feature is turned off, CloudWatch doesn't perform any defined action in any state (OK, INSUFFICIENT_DATA, ALARM).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz109

Source

AWS Config Managed Rule: cloudwatch-alarm-action-enabled-check

Alert Criteria

Yellow: Amazon CloudWatch alarm action is not enabled. No action is performed in any alarm state.

Recommended Action

Enable actions in your CloudWatch alarms unless you have a valid reason to disable them, such as for testing purposes.

If the CloudWatch alarm is no longer needed, delete it to avoid incurring unnecessary costs.

For more information, see [enable-alarm-actions](#) in the AWS CLI Command Reference and [func \(*CloudWatch\) EnableAlarmActions](#) in the AWS SDK for Go API Reference.

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon EC2 Instance Not Managed by AWS Systems Manager

Description

Checks if the Amazon EC2 instances in your account are managed by AWS Systems Manager.

Systems Manager helps you understand and control the current state of your Amazon EC2 instance and OS configurations. With Systems Manager, you can collect software configuration and inventory information about your fleet of instances, including the software installed on them. This allows you to track detailed system configuration, OS patch levels, application configurations, and other details about your deployment.

For more information, see [Setting up Systems Manager for EC2 instances](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz145

Source

AWS Config Managed Rule: `ec2-instance-managed-by-systems-manager`

Alert Criteria

Yellow: The Amazon EC2 instances are not managed by Systems Manager.

Recommended Action

Configure your Amazon EC2 instance to be managed by Systems Manager.

This check can't be excluded from view in the Trusted Advisor console.

For more information, see [Why is my EC2 instance not displaying as a managed node or showing a "Connection lost" status in Systems Manager?](#).

Additional Resources

[Setting up Systems Manager for EC2 instances](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon ECR Repository With Tag Immutability Disabled

Description

Checks if a private Amazon ECR repository has image tag immutability turned on.

Turn on image tag immutability for a private Amazon ECR repository to prevent image tags from being overwritten. This allows you to rely on descriptive tags as a reliable mechanism to track and uniquely identify images. For example, if image tag immutability is turned on, then users can reliably use an image tag to correlate a deployed image version with the build that produced such image.

For more information, see [Image tag mutability](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz129

Source

AWS Config Managed Rule: ecr-private-tag-immutability-enabled

Alert Criteria

Yellow: An Amazon ECR private repository doesn't have tag immutability turned on.

Recommended Action

Turn on image tag immutability for your Amazon ECR private repositories.

For more information, see [Image tag mutability](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon ECS clusters with Container Insights disabled**Description**

Checks if Amazon CloudWatch Container Insights is turned on for your Amazon ECS clusters.

CloudWatch Container Insights collects, aggregates, and summarizes metrics and logs from your containerized applications and microservices. The metrics include utilization for resources such as CPU, memory, disk, and network.

For more information, see [Amazon ECS CloudWatch Container Insights](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz173

Source

AWS Config Managed Rule: ecs-container-insights-enabled

Alert Criteria

Yellow: Amazon ECS cluster does not have container insights enabled.

Recommended Action

Turn on CloudWatch Container Insights on your Amazon ECS clusters.

For more information, see [Using Container Insights](#).

Additional Resources

[Amazon ECS CloudWatch Container Insights](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters

- Last Updated Time

Amazon ECS task logging not enabled

Description

Checks if log configuration is set on active Amazon ECS task definitions.

Checking the log configuration in your Amazon ECS task definitions makes sure that logs generated by containers are properly configured and stored. This helps identify and troubleshoot issues more quickly, optimize performance, and meet compliance requirements.

By default, the logs that are captured show the command output that you typically see in an interactive terminal if you ran the container locally. The awslogs driver passes these logs from Docker to Amazon CloudWatch Logs.

For more information, see [Using the awslogs log driver](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz175

Source

AWS Config Managed Rule: ecs-task-definition-log-configuration

Alert Criteria

Yellow: Amazon ECS task definition does not have a logging configuration.

Recommended Action

Consider specifying the log driver configuration in container definition to send log information to CloudWatch Logs or a different logging driver.

For more information, see [LogConfiguration](#).

Additional Resources

Consider specifying the log driver configuration in container definition to send log information to CloudWatch Logs or a different logging driver.

For more information, see [Example task definitions](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon OpenSearch Service logging CloudWatch not configured

Description

Checks if Amazon OpenSearch Service domains are configured to send logs to Amazon CloudWatch Logs.

Monitoring logs is crucial for maintaining the reliability, availability, and performance of OpenSearch Service.

Search slow logs, indexing slow logs, and error logs are useful for troubleshooting performance and stability issues your workload. These logs need to be enabled to capture data.

You can specify which log types that you want to filter (error, search, index) using the **logTypes** parameter in your AWS Config rules.

For more information, see [Monitoring Amazon OpenSearch Service domains](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz184

Source

AWS Config Managed Rule: opensearch-logs-to-cloudwatch

Alert Criteria

Yellow: Amazon OpenSearch Service does not have a logging configuration with Amazon CloudWatch Logs

Recommended Action

Configure OpenSearch Service domains to publish logs to CloudWatch Logs.

For more information, see [Enabling log publishing \(console\)](#).

Additional Resources

- [Monitoring OpenSearch Service cluster metrics with Amazon CloudWatch](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon RDS DB instances in the clusters with heterogeneous parameter groups

Description

We recommend that all of the DB instances in the DB cluster use the same DB parameter group.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt010

Alert Criteria

Yellow: DB clusters have the DB instances with heterogeneous parameter groups.

Recommended Action

Associate the DB instance with the DB parameter group associated with the writer instance in your DB cluster.

Additional Resources

When the DB instances in your DB cluster use different DB parameter groups, there can be an inconsistent behavior during a failover or compatibility issues between the DB instances in your DB cluster.

For more information, see [Working with parameter groups](#).

Report columns

- Status
- Region
- Resource
- Recommended Value
- Engine Name
- Last Updated Time

Amazon RDS Enhanced Monitoring is turned off

Description

Your database resources don't have Enhanced Monitoring turned on. Enhanced Monitoring provides real-time operating system metrics for monitoring and troubleshooting.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt004

Alert Criteria

Yellow: Amazon RDS resources don't have Enhanced Monitoring turned on.

Recommended Action

Turn on Enhanced Monitoring.

Additional Resources

Enhanced Monitoring for Amazon RDS provides additional visibility on the health of your DB instances. We recommend that you turn on Enhanced Monitoring. When the Enhanced Monitoring option is turned on for your DB instance, it collects vital operating system metrics and process information.

For more information, see [Monitoring OS metrics with Enhanced Monitoring](#).

Report columns

- Status
- Region
- Resource
- Recommended Value
- Engine Name
- Last Updated Time

Amazon RDS Performance Insights is turned off

Description

Amazon RDS Performance Insights monitors your DB instance load to help you analyze and resolve database performance issues. We recommend that you turn on Performance Insights.

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

 Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**.

If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt012

Alert Criteria

Yellow: Amazon RDS resources don't have Performance Insights turned on.

Recommended Action

Turn on Performance Insights.

Additional Resources

Performance Insights uses a lightweight data collection method that doesn't impact the performance of your applications. Performance Insights helps you assess the database load quickly.

For more information, see [Monitoring DB load with Performance Insights on Amazon RDS](#).

Report columns

- Status
- Region
- Resource
- Recommended Value

- Engine Name
- Last Updated Time

Amazon RDS `track_counts` parameter is turned off

Description

When the **track_counts** parameter is turned off, the database doesn't collect the database activity statistics. Autovacuum requires these statistics to work correctly.

We recommend that you set **track_counts** parameter to 1

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Note

When a DB instance or DB cluster is stopped, you can view the Amazon RDS recommendations in Trusted Advisor for 3 to 5 days. After five days, the recommendations are not available in Trusted Advisor. To view the recommendations, open the Amazon RDS console, and then choose **Recommendations**. If you delete a DB instance or DB cluster, then recommendations associated with those instances or clusters are not available in Trusted Advisor or the Amazon RDS management console.

Check ID

c1qf5bt027

Alert Criteria

Yellow: DB parameter groups have **track_counts** parameter turned off.

Recommended Action

Set **track_counts** parameter to 1

Additional Resources

When **track_counts** parameter is turned off, it disables the collection of database activity statistics. The autovacuum daemon requires the collected statistics to identify the tables for autovacuum and autoanalyze.

For more information, see [Run-time Statistics for PostgreSQL](#) on the PostgreSQL documentation website.

Report columns

- Status
- Region
- Resource
- Parameter Value
- Recommended Value
- Last Updated Time

Amazon Redshift cluster audit logging

Description

Checks if your Amazon Redshift clusters have database audit logging turned on. Amazon Redshift logs information about connections and user activities in your database.

You can specify your desired logging Amazon S3 bucket name to match in the **bucketNames** parameter of your AWS Config rules.

For more information, see [Database audit logging](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz134

Source

AWS Config Managed Rule: redshift-audit-logging-enabled

Alert Criteria

Yellow: An Amazon Redshift cluster has database audit logging disabled

Recommended Action

Turn on logging and monitoring for your Amazon Redshift clusters.

For more information, see [Configuring auditing using the console](#).

Additional Resources

[Logging and monitoring in Amazon Redshift](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon S3 Access Logs Enabled

Description

Checks the logging configuration of Amazon Simple Storage Service buckets.

Activating server access logging delivers detailed hourly access logs to a specified Amazon S3 bucket. Access logs contain request details including type, specified resources, and processing

time/date. Logging is turned off by default. Customers should activate access logging to perform security audits or analyze user behavior and usage patterns.

When logging is initially activated, the configuration is automatically validated. However, future modifications can result in logging failures. Note that currently this check doesn't examine Amazon S3 bucket write permissions.

Check ID

c1fd6b9614

Alert Criteria

- Yellow: The bucket does not have server access logging enabled.
- Yellow: The target bucket permissions do not include the root account, so Trusted Advisor cannot check it.
- Red: The target bucket does not exist.
- Red: The target bucket and the source bucket have different owners.
- Green: Bucket has server access logging enabled, the target exists, and permissions to write to target exists

Recommended Action

Activate server access logging for all relevant Amazon S3 buckets. Server access logs provide an audit trail that can be used to understand bucket access patterns and investigate suspicious activity. Activating logging on all applicable buckets will improve visibility into access events across your Amazon S3 environment. See [Enabling Logging Using the Console](#) and [Enabling Logging Programmatically](#).

If the target bucket permissions do not include the root account and you want Trusted Advisor to check the logging status, add the root account as a grantee. See [Editing Bucket Permissions](#).

If the target bucket does not exist, select an existing bucket as a target or create a new one and select it. See [Managing Bucket Logging](#).

If the target and source have different owners, change the target bucket to one that has the same owner as the source bucket. See [Managing Bucket Logging](#).

Additional Resources

[Working with buckets](#)

[Server access logging](#)

[Server access log format](#)

[Deleting log files](#)

Report columns

- Status
- Region
- Resource ARN
- Bucket Name
- Target Name
- Target Exists
- Same Owner
- Write Enabled
- Reason
- Last Updated Time

Amazon S3 does not have Event Notifications enabled

Description

Checks if Amazon S3 Event Notifications is enabled or is correctly configured with the desired destination or types.

The Amazon S3 Event Notifications feature sends notifications when certain events happen in your Amazon S3 buckets. Amazon S3 can send notification messages to Amazon SQS queues, Amazon SNS topics, and AWS Lambda functions.

You can specify your desired destination and event types using the **destinationArn** and **eventTypes** parameters of your AWS Config rules.

For more information, see [Amazon S3 Event Notifications](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz163

Source

AWS Config Managed Rule: s3-event-notifications-enabled

Alert Criteria

Yellow: Amazon S3 does not have Event Notifications enabled, or not configured with the desired destination or types.

Recommended Action

Configure Amazon S3 Event Notifications for object and bucket events.

For more information, see [Enabling and configuring event notifications using the Amazon S3 console](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon SNS Topics Not Logging Message Delivery Status

Description

Checks if Amazon SNS topics have message delivery status logging turned on.

Configure Amazon SNS topics for logging message delivery status to help provide better operational insights. For example, message delivery logging verifies if a message was delivered

to a particular Amazon SNS endpoint. And, it also helps identify the response sent from the endpoint.

For more information, see [Amazon SNS message delivery status](#).

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz121

Source

AWS Config Managed Rule: sns-topic-message-delivery-notification-enabled

Alert Criteria

Yellow: Message delivery status logging is not turned on for an Amazon SNS topic.

Recommended Action

Turn on message delivery status logging for your SNS topics.

For more information, see [Configuring delivery status logging using the AWS Management Console](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Amazon VPC Without Flow Logs

Description

Checks if Amazon Virtual Private Cloud Flow Logs are created for a VPC.

You can specify the traffic type using the **trafficType** parameter in your AWS Config rules.

For more information, see [Logging IP traffic using VPC Flow Logs](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz122

Source

AWS Config Managed Rule: vpc-flow-logs-enabled

Alert Criteria

Yellow: VPCs do not have Amazon VPC Flow Logs.

Recommended Action

Create VPC Flow Logs for each of your VPCs.

For more information, see [Create a flow log](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters

- Last Updated Time

Application Load Balancers and Classic Load Balancers Without Access Logs Enabled

Description

Checks if Application Load Balancers and Classic Load Balancers have access logging enabled.

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and troubleshoot issues.

Access logs are an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logs for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify.

You can specify the access log Amazon S3 bucket that you want to check using the **s3BucketNames** parameter in your AWS Config rules.

For more information, see [Access logs for your Application Load Balancer](#) or [Access logs for your Classic Load Balancer](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz167

Source

AWS Config Managed Rule: elb-logging-enabled

Alert Criteria

Yellow: Access logs feature not enabled for an Application Load Balancer or Classic Load Balancer.

Recommended Action

Enable access logs for your Application Load Balancers and Classic Load Balancers.

For more information, see [Enable access logs for your Application Load Balancer](#) or [Enable access logs for your Classic Load Balancer](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

CloudFormation Stack Notification

Description

Checks if all of your CloudFormation stacks use Amazon SNS to receive notifications when an event occurs.

You can configure this check to look for specific Amazon SNS topic ARNs using parameters in your AWS Config rules.

For more information, see [Setting CloudFormationstack options](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz111

Source

AWS Config Managed Rule: `cloudformation-stack-notification-check`

Alert Criteria

Yellow: Amazon SNS event notifications for your CloudFormation stacks are not turned on.

Recommended Action

Make sure that your CloudFormation stacks use Amazon SNS to receive notifications when an event occurs.

Monitoring stack events helps you to respond quickly to unauthorized actions that might alter your AWS environment.

Additional Resources

[How can I receive an email alert when my AWS CloudFormation stack enters ROLLBACK_IN_PROGRESS status?](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

AWS CloudTrail data events logging for objects in an S3 bucket

Description

Checks if at least one AWS CloudTrail trail logs Amazon S3 data events for all of your Amazon S3 buckets.

For more information, see [Logging Amazon S3 API calls using AWS CloudTrail](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz166

Source

AWS Config Managed Rule: cloudtrail-s3-dataevents-enabled

Alert Criteria

Yellow: AWS CloudTrail event logging for Amazon S3 buckets is not configured

Recommended Action

Enable CloudTrail event logging for Amazon S3 buckets and objects to track requests for target bucket access.

For more information, see [Enabling CloudTrail event logging for S3 buckets and objects](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

AWS CodeBuild Project Logging

Description

Checks if the AWS CodeBuild project environment uses logging. Logging options can be logs in Amazon CloudWatch Logs, or built in a specified Amazon S3 bucket, or both. Enabling logging in a CodeBuild project can provide several benefits such as debugging and auditing.

You can specify the name of the Amazon S3 bucket or CloudWatch Logs group for storing the logs, using the **s3BucketNames** or **cloudWatchGroupNames** parameter in your AWS Config rules.

For more information, see [Monitoring AWS CodeBuild](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz113

Source

AWS Config Managed Rule: codebuild-project-logging-enabled

Alert Criteria

Yellow: AWS CodeBuild project logging is not enabled.

Recommended Action

Make sure that logging is turned on in your AWS CodeBuild project. This check can't be excluded from view in the AWS Trusted Advisor console.

For more information, see [Logging and monitoring in AWS CodeBuild](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

AWS CodeDeploy Auto Rollback and Monitor Enabled

Description

Checks if the deployment group is configured with automatic deployment rollback and deployment monitoring with alarms attached. If something goes wrong during a deployment, it is automatically rolled back, and your application remains in a stable state

For more information, see [Redeploy and roll back a deployment with CodeDeploy](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz114

Source

AWS Config Managed Rule: codedeploy-auto-rollback-monitor-enabled

Alert Criteria

Yellow: AWS CodeDeploy automatic deployment rollback and deployment monitoring are not enabled.

Recommended Action

Configure a deployment group or deployment to automatically roll back when a deployment fails or when a monitoring threshold you specify is met.

Configure alarm to monitor various metrics, such as CPU usage, memory usage, or network traffic, during the deployment process. If any of these metrics exceed certain thresholds, the alarms trigger, and the deployment is stopped or rolled back.

For information on setting up automatic rollbacks and configuring alarms for your deployment groups, see [Configure advanced options for a deployment group](#).

Additional Resources

[What is CodeDeploy?](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

AWS CodeDeploy Lambda is using all-at-once deployment configuration

Description

Checks if the AWS CodeDeploy deployment group for AWS Lambda compute platform is using all-at-once deployment configuration.

To reduce the risk of deployment failures of your Lambda functions in CodeDeploy, it's a best practice to use the canary or linear deployment configuration instead of the default option where all traffic is shifted from the original Lambda function to the updated function at once.

For more information, see [Lambda function versions](#) and [Deployment configuration](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz115

Source

AWS Config Managed Rule: codedeploy-lambda-allatonce-traffic-shift-disabled

Alert Criteria

Yellow: AWS CodeDeploy Lambda deployment uses the all-at-once deployment configuration to shift all traffic to the updated Lambda functions at once.

Recommended Action

Use the Canary or Linear deployment configuration of CodeDeploy deployment group for the Lambda compute platform.

Additional Resources

[Deployment configuration](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

AWS Elastic Beanstalk Enhanced Health Reporting is not Configured

Description

Checks if an AWS Elastic Beanstalk environment is configured for enhanced health reporting.

Elastic Beanstalk enhanced health reporting provides detailed performance metrics, such as CPU usage, memory usage, network traffic, and infrastructure health information, such as number of instances and load balancer status.

For more information, see [Enhanced health reporting and monitoring](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz108

Source

AWS Config Managed Rule: beanstalk-enhanced-health-reporting-enabled

Alert Criteria

Yellow: Elastic Beanstalk environment is not configured for enhanced health reporting

Recommended Action

Make sure that an Elastic Beanstalk environment is configured for enhanced health reporting.

For more information, see [Enabling enhanced health reporting using the Elastic Beanstalk console](#).

Additional Resources

- [Enabling Elastic Beanstalk enhanced health reporting](#)

- [Enhanced health reporting and monitoring](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

AWS Elastic Beanstalk with Managed Platform Updates Disabled

Description

Checks if managed platform updates in Elastic Beanstalk environments and configuration templates are enabled.

AWS Elastic Beanstalk regularly releases platform updates to provide fixes, software updates, and new features. With managed platform updates, Elastic Beanstalk can automatically perform platform updates for new patch and minor platform versions.

You can specify your desired update level in the **UpdateLevel** parameters of your AWS Config rules.

For more information, see [Updating your Elastic Beanstalk environment's platform version](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz177

Source

AWS Config Managed Rule: elastic-beanstalk-managed-updates-enabled

Alert Criteria

Yellow: AWS Elastic Beanstalk managed platform updates is not configured at all, including at a minor or patch level.

Recommended Action

Enable managed platform updates in your Elastic Beanstalk environments, or configure it at a minor or update level.

For more information, see [Managed platform updates](#).

Additional Resources

- [Enabling Elastic Beanstalk enhanced health reporting](#)
- [Enhanced health reporting and monitoring](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

AWS Fargate platform version is not latest

Description

Checks if Amazon ECS is running the latest platform version of AWS Fargate. The Fargate platform version refers to a specific runtime environment for Fargate task infrastructure. It's a combination of the kernel and container runtime versions. New platform versions are released as runtime environment evolves. For example, if there are kernel or operating system updates, new features, bug fixes, or security updates.

For more information, see [Fargate task maintenance](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz174

Source

AWS Config Managed Rule: ecs-fargate-latest-platform-version

Alert Criteria

Yellow: Amazon ECS is not running on the latest version of the Fargate platform.

Recommended Action

Update to the latest Fargate platform version.

For more information, see [Fargate task maintenance](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

AWS Systems Manager State Manager Association in Non-compliant Status**Description**

Checks if the status of the AWS Systems Manager association compliance is COMPLIANT or NON_COMPLIANT after the association execution on the instance.

State Manager, a capability of AWS Systems Manager, is a secure and scalable configuration management service that automates the process of keeping your managed nodes and other AWS resources in a state that you define. A State Manager association is a configuration that you assign to your AWS resources. The configuration defines the state that you want to maintain on your resources, so it helps you to achieve the target, such as avoidance of configuration drifts across your Amazon EC2 instances.

For more information, see [AWS Systems Manager State Manager](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear. For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz147

Source

AWS Config Managed Rule: ec2-managedinstance-association-compliance-status-check

Alert Criteria

Yellow: The status of the AWS Systems Manager association compliance is NON_COMPLIANT.

Recommended Action

Validate the status of the State Manager associations, and then take any needed actions to return the status back to COMPLIANT.

For more information, see [About State Manager](#).

Additional Resources

[AWS Systems Manager State Manager](#)

Report columns

- Status

- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

CloudTrail trails are not configured with Amazon CloudWatch Logs

Description

Checks if AWS CloudTrail trails are configured to send logs to CloudWatch Logs.

Monitor CloudTrail Log files with CloudWatch Logs to trigger an automated response when critical events are captured in AWS CloudTrail.

For more information, see [Monitoring CloudTrail Log Files with CloudWatch Logs](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz164

Source

AWS Config Managed Rule: `cloud-trail-cloud-watch-logs-enabled`

Alert Criteria

Yellow: AWS CloudTrail is not set up with CloudWatch Logs integration.

Recommended Action

Configure CloudTrail trails to send log events to CloudWatch Logs.

For more information, see [Creating CloudWatch alarms for CloudTrail events: examples](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Elastic Load Balancing Deletion Protection Not Enabled for Load Balancers

Description

Checks if deletion protection is turned on for your load balancers.

Elastic Load Balancing supports deletion protection for your Application Load Balancers, Network Load Balancers, and Gateway Load Balancers. Turn on deletion protection to prevent your load balancer from accidental deletion. Deletion protection is turned off by default when you create a load balancer. If your load balancers are part of a production environment, then consider turning on deletion protection.

Access logs are an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logs for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify.

For more information, see [Application Load Balancer Deletion protection](#), [Network Load Balancers Deletion protection](#), or [Gateway Load Balancers Deletion protection](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz168

Source

AWS Config Managed Rule: `elb-deletion-protection-enabled`

Alert Criteria

Yellow: Deletion protection is not enabled for a load balancer.

Recommended Action

Turn on deletion protection for your Application Load Balancers, Network Load Balancers, and Gateway Load Balancers.

For more information, see [Application Load Balancer Deletion protection](#), [Network Load Balancers Deletion protection](#), or [Gateway Load Balancers Deletion protection](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

NitroV4 Amazon EC2 instances running outdated ENA driver

Description

Checks for Amazon Elastic Compute Cloud (Amazon EC2) instances running on Nitro system version 4 using outdated Elastic Network Adapter (ENA) drivers. To take advantage of the Nitro performance improvements using the accelerated path, you must update your ENA driver to a version compatible with NitroV4, as outlined in the [Amazon EC2 Instance Types guide](#). To make this change directly, update your driver, or you can update your operating system to a version that is configured with a compatible driver.

For Nitro system version 5 and later based instances, the minimum driver versions in the [Amazon EC2 Instance Types guide](#) are required. AWS recommends upgrading to the latest

ENA driver version to benefit from performance improvements, bug fixes, and security enhancements. Starting with Nitro system version 5 or later, outdated ENA drivers fail to bind to instances after launch, which prevents network communication.

 Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

Check ID

c17gc3v72d

Alert Criteria

Yellow: Amazon EC2 instances on Nitro system version 4 are using ENA driver versions older than Linux 2.2.9 or Windows 2.2.3.

Recommended Action

Update to [newer ENA drivers](#) (Linux 2.2.9+, Windows 2.2.3+, DPDK ENA PMD v2.2.0+) to access accelerated path features that are unavailable in older versions. DPDK users can either move to DPDK 20.11+ (which includes DPDK ENA PMD v2.2.0+) or apply DPDK ENA PMD backports directly from [GitHub](#) without changing their DPDK release.

For existing instances, perform this update when it's convenient to reboot your instance. Update your operating system or ENA driver to the latest supported version for your operating system, and then reboot your instance. For Linux instances, update to ENA driver version 2.2.9 or later. For Windows instances, update to version 2.2.3 or later.

For new instance launches, build and use updated AMIs with the latest ENA drivers preinstalled using Amazon EC2 Image Builder.

Additional Resources

- [ENA driver performance on Nitro-based instances](#)
- [Monitoring network performance for ENA](#)
- [Amazon EC2 Image Builder for automated AMI creation](#)
- [Installing and upgrading the ENA driver on Linux and Windows](#)

Report columns

- Status
- Region
- Instance ID
- Interface ID
- Last Updated Time

RDS DB Cluster Deletion Protection Check

Description

Checks if your Amazon RDS DB clusters have deletion protection enabled.

When a cluster is configured with deletion protection, the database cannot be deleted by any user.

Deletion protection is available for Amazon Aurora and RDS for MySQL, RDS for MariaDB, RDS for Oracle, RDS for PostgreSQL, and RDS for SQL Server database instances in all AWS Regions.

For more information, see [Deletion protection for Aurora clusters](#).

Check ID

c18d2gz160

Source

AWS Config Managed Rule: `rds-cluster-deletion-protection-enabled`

Alert Criteria

Yellow: You have Amazon RDS DB clusters that don't have deletion protection enabled.

Recommended Action

Turn on deletion protection when you create an Amazon RDS DB cluster.

You can only delete clusters that don't have deletion protection enabled. Enabling deletion protection adds an extra layer of protection and avoids data loss from accidental or non-accidental deletion of a database instance. Deletion protection also helps meet regulatory compliance requirements and ensure business continuity.

For more information, see [Deletion protection for Aurora clusters](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Additional Resources

[Deletion protection for Aurora clusters](#)

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

RDS DB Instance Automatic Minor Version Upgrade Check

Description

Checks if Amazon RDS DB instances have automatic minor version upgrades configured.

Turn on automatic minor version upgrades for an Amazon RDS instance to make sure that the database is always running the latest secure and stable version. Minor upgrades provide security updates, bug fixes, performance improvements, and maintain compatibility with existing applications.

For more information, see [Upgrading a DB instance engine version](#).

Note

Results for this check are automatically refreshed several times daily, and refresh requests are not allowed. It might take a few hours for changes to appear.

For AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan customers, you can use the [BatchUpdateRecommendationResourceExclusion](#) API to include or exclude one or more resources from your Trusted Advisor results.

Check ID

c18d2gz155

Source

AWS Config Managed Rule: `rds-automatic-minor-version-upgrade-enabled`

Alert Criteria

Yellow: RDS DB instance does not have automatic minor version upgrades turned on.

Recommended Action

Turn on automatic minor version upgrades when you create a Amazon RDS DB instance.

When you turn on minor version upgrade, the database version automatically upgrades if it is running a minor version of the DB engine that is lower than the [Manually upgrading the engine version](#).

Report columns

- Status
- Region
- Resource
- AWS Config Rule
- Input Parameters
- Last Updated Time

Change log for AWS Trusted Advisor

See the following topic for recent changes to Trusted Advisor checks.

Note

If you use the Trusted Advisor console or the AWS Support API, deprecated checks won't appear in check results. If you use a deprecated check, such as specifying the check ID in

an AWS Support API operation or your code, then you receive API call errors. Remove these checks to avoid errors.

For more information about the available checks, see the [AWS Trusted Advisor check reference](#).

Change date	Check name	Change description
September 19, 2026	Updated c1dvkm4z6b: Amazon ECS AWSLogs driver in blocking mode	Updated the Description to indicate that the check evaluates only Amazon ECS task definitions that are actively in use by running tasks or that are among the 5 most recent revisions per task definition family.
September 8, 2026	New NitroV4 Amazon EC2 instances running outdated ENA driver	Added a new check that identifies Amazon EC2 instances running on Nitro system version 4 that use outdated Elastic Network Adapter (ENA) drivers.
August 27, 2026	Updated Security Hub CSPM checks in Trusted Advisor	Trusted Advisor supports only AWS Security Hub CSPM controls available through July 8, 2026. For more information, see Viewing AWS Security Hub CSPM controls in AWS Trusted Advisor .
February 19, 2026	New Idle NAT gateways	Added a new check to identify idle or underutilized NAT gateways that might be candidates for cost optimization.

Change date	Check name	Change description
January 9, 2026	Deprecated 5 Cost Optimization Trusted Advisor checks	The following checks were deprecated: • Savings Plan - vZ2c2W1srf • Amazon Relational Database Service (RDS) Reserved Instance Optimization - 1qazXsw23e • Amazon Redshift Reserved Node Optimization - 1qw23er45t • Amazon ElastiCache Reserved Node Optimization - h3L1otH3re • Amazon OpenSearch Service Reserved Instance Optimization - 7ujm6yhn5t
December 18, 2025	Updated Amazon S3 Bucket Versioning	Added a new Alert criteria: • Yellow: Trusted Advisor doesn't have access to validate versioning
December 17, 2025	Updated Amazon S3 Bucket Permissions	Updated the Alert criteria section.
November 21, 2025	Updated Application Load Balancer security group	Updated the Application Load Balancer security group alerts and recommendations.

Change date	Check name	Change description
November 17, 2025	Updated AWS STS global endpoint usage across AWS Regions check description	Updated the AWS STS global endpoint usage across AWS Regions check description to clarify when check results are refreshed.
October 15, 2025	Updated multiple check descriptions	A note was added to multiple check descriptions to indicate that the check reports all resources that are flagged by the criteria and the total number of resources evaluated, including OK resources. The resources table lists only the flagged resources.
October 10, 2025	Updated check reference	Updated the check reference to show all available checks.
September 11, 2025	L4dfs2Q4C5: AWS Lambda functions using deprecated runtimes	Updated Yellow alert criterion to indicate that runtimes deprecating within at least 180 are included.
August 19, 2025	Pfx0RwqBli: Amazon S3 Bucket Permissions	Alert criteria updated: <i>Trusted Advisor does not have permission to check the policy or ACL, or the policy or ACL could not be evaluated for other reasons</i> changed from Yellow to Red .

Change date	Check name	Change description
July 22, 2025	- vjafUGJ9H0: AWS CloudTrail Logging - BueAdJ7NrP: Amazon S3 Bucket Logging	These checks are deprecated.
July 03, 2025	Pfx0RwqBli: Amazon S3 Bucket Permissions	The Alert criteria is updated to reflect all Yellow and Red criteria.
July 03, 2025	c1dfprch15: Amazon EC2 instances with Ubuntu LTS end of standard support	Updated the note to indicate that this check refreshes at least once daily.
July 02, 2025	c1dvkm4z6b: Amazon ECS AWSLogs driver in blocking mode	Amazon ECS changed the default setting for awslogs driver logging configuration parameter mode from blocking to non-blocking. The Yellow status description has been updated to reflect this change.
July 02, 2025	7DAFEemoDos: MFA on root account	Added information indicating that member account root user credentials can be deleted centrally, removing the need to manage MFA on root user credentials.
June 9, 2025	c1z7kmr17n: Amazon Aurora cost optimization recommendations for DB cluster storage	New check
June 09, 2025	c15m0mgld3: AWS STS global endpoint usage across AWS Regions	Updated check: This check is now available for all AWS Support plans.

Change date	Check name	Change description
June 02, 2025	c15m0mgld3: AWS STS global endpoint usage across AWS Regions	New check

Change date	Check name	Change description
May 30, 2025	• c1z7kmr15n: Amazon DynamoDB reserved capacity purchase recommendations • c1z7kmr02n: Amazon EBS cost optimization recommendations for volumes • c1z7kmr01n: Amazon EC2 cost optimization recommendations for Amazon EC2 Auto Scaling groups • c1z7kmr00n: Amazon EC2 cost optimization recommendations for instances • c1z7kmr13n: Amazon ElastiCache reserved node purchase recommendations • c1z7kmr16n: Amazon MemoryDB reserved node purchase recommendations • c1z7kmr14n: Amazon OpenSearch Service Reserved Instance purchase recommendations • c1z7kmr04n: Amazon RDS cost optimization recommendations for DB instance storage • c1z7kmr03n: Amazon RDS cost optimization	New Cost Optimization Hub checks

Change date	Check name	Change description
	recommendations for DB instances • c1z7kmr11n: Amazon RDS Reserved Instance purchase recommendations • c1z7kmr12n: Amazon Redshift reserved node purchase recommendations • c1z7kmr06n: AWS Fargate cost optimization recommendations for Amazon ECS • c1z7kmr05n: AWS Lambda cost optimization recommendations for functions • c1z7kmr08n: AWS Savings Plans purchase recommendations for Amazon SageMaker AI • c1z7kmr09n: AWS Savings Plans purchase recommendations for compute	
April 30, 2025	• N420c450f2: CloudFront Alternate Domain Names • N425c450f2: CloudFront Custom SSL Certificates in the IAM Certificate Store	Added a note indicating that this check applies to classic Amazon CloudFront distributions.
April 30, 2025	N415c450f2: CloudFront Header Forwarding and Cache Hit Ratio	Added a note indicating that this check applies to classic Amazon CloudFront distributions.

Change date	Check name	Change description
April 02, 2025	c1dfprch02: Amazon EFS Throughput Mode Optimization	The description of this check has changed. For more information, see Amazon EC2 instances with Microsoft Windows Server end of support .
April 02, 2025	Qsdfp3A4L4: Amazon EC2 instances with Microsoft Windows Server end of support	The description of this check has changed. For more information, see Amazon EFS Throughput Mode Optimization .

Older updates

The following AWS Security Hub CSPM checks are deprecated:

Check name	Check ID
S3.10 - S3 general purpose buckets with versioning enabled should have lifecycle configurations	Hs4Ma3G211
S3.11 - S3 general purpose buckets should have event notifications enabled	Hs4Ma3G212
CodeBuild.5 - CodeBuild project environments should not have privileged mode enabled	Hs4Ma3G218
CloudFormation.1 - CloudFormation stacks should be integrated with Amazon Simple Notification Service (SNS)	Hs4Ma3G245

Check name	Check ID
SNS.2 - Logging of delivery status should be enabled for notification messages sent to a topic	Hs4Ma3G263
Athena.1 - Athena workgroups should be encrypted at rest	Hs4Ma3G294

New check: Amazon RDS Continuous Backup Not Enabled

Trusted Advisor added the following check on December 23, 2024.

Check name	Check category	Check ID
Amazon RDS Continuous Backup Not Enabled	Fault tolerance	44fde09ab5

Checks if an Amazon RDS instance is enabled with automated backups using Amazon RDS or with continuous backups of AWS Backup. Continuous backups reduce the risk of unexpected data loss and allow for point-in-time recovery.

For more information, see [Amazon RDS Continuous Backup Not Enabled](#).

New check: AWS CloudTrail Management Events Logging

Trusted Advisor added the following check on December 23, 2024.

Check name	Check category	Check ID
AWS CloudTrail Management Events Logging	Security	c25hn9x03v

Checks your use of AWS CloudTrail.

For more information, see [AWS CloudTrail Management Event Logging](#).

Updated the Auto Scaling Group Resources check

Trusted Advisor updated the following check on December 23, 2024.

Check name	Check category	Check ID
Auto Scaling Group Resources	Fault tolerance	8CNsS11I5v

The description of this check is updated to include launch configurations and launch templates.

A new alert criteria, Red: A launch template is associated with a deleted Amazon Machine Image (AMI). was added.

For more information, see [Auto Scaling Group Resources](#).

Updated the IAM Access Analyzer External Access check

Trusted Advisor updated the following check on December 23, 2024.

Check name	Check category	Check ID
IAM Access Analyzer External Access	Security	07602fcad6

The description of this check is updated to indicate that it analyzes IAM access at the account level.

For more information, see [IAM Access Analyzer External Access](#).

Added 1 new check

Trusted Advisor added 1 new check on November 22, 2024:

- 8604e947f2 - [Application Load Balancer Security Groups](#)

Updated 3 checks

Trusted Advisor updated 3 checks on November 7, 2024:

- b92b83d667 - [ELB Target Imbalance](#)
- 8CNsSIII5v - [Auto Scaling Group Resources](#)
- wuy7G1zxql - [Amazon EC2 Availability Zone Balance](#)

Added 4 checks

Trusted Advisor added 4 new checks on October 11, 2024:

- 07602fcad6 - IAM Access Analyzer - external access
- 528d6f5ee7 - GWLB - Endpoint AZ
- c2vlfgojp6 - Inactive VPC interface endpoints
- c2vlfgo0k35 - Inactive Gateway Load Balancer endpoints

Updated 3 checks

Trusted Advisor updated 3 checks on October 2, 2024:

- Check ID 7040ea389a moved from Cost Optimization pillar to the Fault Tolerance pillar
- Updated Check ID 7DAFEemoDos
- Updated Check ID Cmsvnj8db2

Added 9 new checks

Trusted Advisor added 9 new checks on August 23, 2024:

- c2vlfgo0p86 - [IAM] - SAML 2.0 Identity Provider
- 7040ea389a - Network Firewall endpoint Cross-AZ Data Transfer
- c2vlfgo0b0w - Low utilization Network Firewall
- c2vlfgo0gqd - Network Firewall Multi-AZ
- c2vlfgo0p1w - Application Load Balancer Target Groups encrypted protocol
- c2vlfgo022t - [NAT Gateway] - Underutilized Resource
- c243hjzrh - AWS Outposts Single Rack deployment
- b92b83d667 - ELB Target Imbalance
- 90046ff5b5 - MSK availability is limited to two zones

For more information, see the [AWS Trusted Advisor check reference](#).

Updated 1 Security check and added 1 Security check

Trusted Advisor updated 1 Operational Excellence checks on August 22, 2024:

- c1fd6b96l4

Trusted Advisor added 1 Security checks on August 22, 2024:

- c2vlfG0f4h

For more information, see the [AWS Trusted Advisor check reference](#).

Updated 6 Security checks

Trusted Advisor updated 6 Security checks on August 20, 2024:

- nNauJisYIT
- c9D319e7sG
- a2sEc6ILx
- HCP4007jGY
- 1iG5NDGVre
- Yw2K9puPzl

For more information, see the [AWS Trusted Advisor check reference](#).

Updated 1 fault tolerance checks

Trusted Advisor updated the 1 fault tolerance check and 1 security on August 12, 2024:

- VPN Tunnel Redundancy
- Amazon RDS engine minor version upgrade is required

For more information, see the [AWS Trusted Advisor check reference](#).

Updated 9 checks

Trusted Advisor updated the 9 checks on July 21, 2024:

- 7qGXsKIUw
- ZRxQIPsb6c
- N425c450f2
- 7DAFEemoDos
- Pfx0RwqBli
- H7lgTzjTYb
- C056F80cR3
- Yw2K9puPzl
- xSqX82fQu

For more information, see the [AWS Trusted Advisor check reference](#).

Removed 5 checks and added 1 check

Trusted Advisor deprecated 3 Fault Tolerance checks, 1 Performance check, and 1 Security check on May 15, 2024:

- IAM Use
- ELB Cross-Zone Load Balancing
- Overutilized Amazon EBS Magnetic Volumes
- Large Number of EC2 Security Group Rules Applied to an Instance
- Large Number of Rules in an EC2 Security Group

Trusted Advisor added 1 new security check on May 15, 2024:

- Amazon S3 Server Access Logs Enabled

For more information, see the [AWS Trusted Advisor check reference](#).

Removed fault tolerance checks

Trusted Advisor deprecated 3 Fault Tolerance check on April 25, 2024:

- Direct Connect Connection Redundancy
- Direct Connect Location Redundancy
- Direct Connect Virtual Interface Redundancy

For more information, see the [AWS Trusted Advisor check reference](#).

New fault tolerance check

Trusted Advisor added 1 Fault Tolerance check on February 29, 2024:

- NLB - Internet-facing resource in private subnet

For more information, see the [AWS Trusted Advisor check reference](#).

Updated fault tolerance and security checks

Trusted Advisor added 1 new Fault Tolerance check and amended 1 existing Fault tolerance and 1 Security check on March 28 2024:

- Added AWS Resilience Hub Application Component check
- Updated AWS Lambda VPC-enabled Functions without Multi-AZ Redundancy
- Updated AWS Lambda Functions Using Deprecated Runtimes

For more information, see the [AWS Trusted Advisor check reference](#).

New fault tolerance check

Trusted Advisor added 1 Fault Tolerance check on January 31, 2024:

- Direct Connect Location Resiliency

For more information, see the [AWS Trusted Advisor check reference](#).

Updated fault tolerance check

Trusted Advisor amended 1 Fault Tolerance check on January 08, 2024:

- Amazon RDS innodb_flush_log_at_trx_commit parameter is not 1

For more information, see the [AWS Trusted Advisor check reference](#).

Updated security check

Trusted Advisor amended 1 Security check on December 21, 2023:

- AWS Lambda Functions Using Deprecated Runtimes

For more information, see the [AWS Trusted Advisor check reference](#).

New security and performance checks

Trusted Advisor added 2 new Security checks and 2 new Performance checks on December 20, 2023:

- Amazon EFS clients not using data-in-transit encryption
- Amazon Aurora DB cluster under-provisioned for read workload
- Amazon RDS instance under-provisioned for system capacity
- Amazon EC2 instances with Ubuntu LTS end of standard support

For more information, see the [AWS Trusted Advisor check reference](#).

New security check

Trusted Advisor added 1 new Security check on December 15, 2023:

- Amazon Route 53 mismatching CNAME records pointing directly to S3 buckets

For more information, see the [AWS Trusted Advisor check reference](#).

New fault tolerance and cost optimization checks

Trusted Advisor added 2 new Fault Tolerance checks and 1 new Cost Optimization check on December 07, 2023:

- Amazon DocumentDB Single-AZ clusters
- Amazon S3 Incomplete Multipart Upload Abort Configuration
- Amazon ECS AWSLogs driver in blocking mode

For more information, see the [AWS Trusted Advisor check reference](#).

New fault tolerance checks

Trusted Advisor added 3 new fault tolerance checks on November 17, 2023:

- ALB Multi-AZ
- NLB Multi-AZ
- VPC interface endpoint network interfaces in multiple AZs

For more information, see the [AWS Trusted Advisor check reference](#).

New checks for Amazon RDS

Trusted Advisor added 37 new checks for Amazon RDS on November 15, 2023.

For more information, see the [AWS Trusted Advisor check reference](#).

New AWS Trusted Advisor API

AWS Trusted Advisor introduces new APIs to enable you to programmatically access Trusted Advisor best practice checks, recommendations, and prioritized recommendations. Trusted Advisor APIs enable you to programmatically integrate Trusted Advisor with your preferred operational tool to automate and optimize your workloads at scale. Available to Business, Enterprise On-Ramp, or Enterprise Support customers, the new APIs provide access to Trusted Advisor recommendations for your account or all the linked accounts within a payer account. Enterprise Support customers with access to management or delegated administrator accounts can additionally programmatically retrieve prioritized recommendations across their organization.

The new Trusted Advisor APIs will replace the 3 functionalities previously offered through AWS Support API (SAPI). SAPI will continue to offer case and other support information.

Trusted Advisor APIs are generally available in the US East (Ohio), US East (N. Virginia), US West (Oregon), Asia Pacific (Seoul), Asia Pacific (Sydney), and Europe (Ireland) Regions.

To learn more, please visit the [AWS Trusted Advisor API page](#).

Trusted Advisor check removal

Trusted Advisor removed the following checks on November 9, 2023.

Check name	Check category	Check ID
EBS volumes should be attached to EC2 instances	Security	Hs4Ma3G119
S3 buckets should have server-side encryption enabled	Security	Hs4Ma3G167
CloudFront distributions should have origin access identity enabled	Security	Hs4Ma3G195

Integration of AWS Config checks into Trusted Advisor

Trusted Advisor added 64 new checks powered by AWS Config on October 30, 2023.

For more information, see the [View AWS Trusted Advisor checks powered by AWS Config](#).

New fault tolerance checks

Trusted Advisor added the following checks on October 12, 2023.

- Amazon RDS ReplicaLag
- Amazon RDS FreeStorageSpace
- Amazon RDS DiskQueueDepth
- Amazon Route 53 Resolver Endpoint Availability Zone Redundancy

- Auto Scaling available IPs in Subnets
- Amazon MSK brokers hosting too many partitions

For more information, see the [Fault tolerance](#) category.

New service limits check

Trusted Advisor added the following check on August 17, 2023.

- Lambda Code Storage Usage

For more information, see the [Service limits](#) category.

New fault tolerance check

Trusted Advisor added the following check on August 3, 2023.

- AWS Lambda On Failure Event Destinations

For more information, see the [Fault tolerance](#) category.

New fault tolerance and performance checks

Trusted Advisor added the following checks on June 1, 2023.

- Amazon EFS No Mount Target Redundancy
- Amazon EFS Throughput Mode Optimization
- ActiveMQ Availability Zone Redundancy
- RabbitMQ Availability Zone Redundancy

For more information, see the [Fault tolerance](#) category and [Performance](#) category.

New fault tolerance checks

Trusted Advisor added the following checks on May 16, 2023.

- NAT Gateway AZ Independence

- Single AZ Application Check

For more information, see the [Fault tolerance](#) category.

New fault tolerance checks

Trusted Advisor added the following checks on April 27, 2023.

- Number of AWS Regions in an Incident Manager replication set
- AWS Resilience Hub assessment age

For more information, see the [Fault tolerance](#) category.

Region Expansion of Amazon ECS Fault Tolerance Checks

Trusted Advisor expanded the following checks into additional regions on April 27, 2023. Trusted Advisor checks for Amazon ECS are now available in all regions where Amazon ECS is generally available.

- Amazon ECS service using a single AZ
- Amazon ECS Multi-AZ placement strategy

Regions expanded into include Africa (Cape Town), Asia Pacific (Hong Kong), Asia Pacific (Hyderabad), Asia Pacific (Jakarta), Asia Pacific (Melbourne), Europe (Milan), Europe (Spain), Europe (Zurich), Middle East (Bahrain), Middle East (UAE).

New fault tolerance checks

Trusted Advisor added the following checks on March 30, 2023.

- Amazon ECS service using a single AZ
- Amazon ECS Multi-AZ placement strategy

For more information, see the [Fault tolerance](#) category.

New fault tolerance checks

Trusted Advisor added the following checks on December 15, 2022.

- AWS CloudHSM clusters running HSM instances in a single AZ
- Amazon ElastiCache Multi-AZ clusters
- Amazon MemoryDB Multi-AZ clusters

To receive results in Trusted Advisor for your AWS CloudHSM, ElastiCache, and MemoryDB clusters, you must have clusters in your Availability Zones. For more information, see the following documentation:

- [AWS CloudHSM User Guide](#)
- [Amazon MemoryDB Developer Guide](#)

Trusted Advisor updated the following check information on December 15, 2022.

- AWS Resilience Hub policy breached – App Name was updated to Application Name
- AWS Resilience Hub resilience scores – App Name and App Resilience Score were updated to Application Name and Application Resilience Score

For more information, see the [Fault tolerance](#) category.

Updates to the Trusted Advisor integration with AWS Security Hub CSPM

Trusted Advisor made the following update on November 17, 2022.

If you disable Security Hub CSPM or AWS Config for an AWS Region, Trusted Advisor now removes your control findings for that AWS Region within 7-9 days. Previously, the time frame to remove your Security Hub CSPM data from Trusted Advisor was 90 days.

For more information, see the following sections in the [Troubleshooting](#) topic:

- [I turned off Security Hub CSPM or AWS Config in a Region](#)
- [My control is archived in Security Hub CSPM, but I still see the findings in Trusted Advisor](#)

New fault tolerance checks for AWS Resilience Hub

Trusted Advisor added the following checks on November 17, 2022.

- AWS Resilience Hub policy breached
- AWS Resilience Hub resilience scores

You can use these checks to view the latest resilience policy status and resilience score for your applications. Resilience Hub provides you with a central place to define, track, and manage the resiliency and availability of your applications.

To receive results in Trusted Advisor for your Resilience Hub applications, you must deploy an AWS application and use Resilience Hub to track the resiliency posture of the application. For more information, see the [AWS Resilience Hub User Guide](#).

To receive results in Trusted Advisor for your ElastiCache and MemoryDB clusters, you must have clusters in your Availability Zones. For more information, see the following documentation:

[Amazon MemoryDB Developer Guide](#)

For more information, see the [Fault tolerance](#) category.

Update to the Trusted Advisor console

Trusted Advisor added the following change on November 16, 2022.

The Trusted Advisor Dashboard in the console is now Trusted Advisor Recommendations. The Trusted Advisor Recommendations page still shows the check results and the available checks for each category for your AWS account.

This name change only updates the Trusted Advisor console. You can continue to use the Trusted Advisor console and the Trusted Advisor operations in the Support API as usual.

For more information, see [Get started with Trusted Advisor Recommendations](#).

New checks for Amazon EC2

Trusted Advisor added the following check on September 1, 2022.

- Amazon EC2 instances with Microsoft Windows Server end of support

For more information, see the [Security](#) category.

Added Security Hub CSPM checks to Trusted Advisor

As of June 23, 2022, Trusted Advisor only supports Security Hub CSPM controls available through April 7, 2022. This release supports all controls in the AWS Foundational Security Best Practices security standard except for controls in the Category: Recover > Resilience. For more information, see [Viewing AWS Security Hub CSPM controls in AWS Trusted Advisor](#).

For a list of supported controls, see [AWS Foundational Security Best Practices controls](#) in the *AWS Security Hub CSPM User Guide*.

Added checks from AWS Compute Optimizer

Trusted Advisor added the following checks on May 4, 2022.

Check name	Check category	Check ID
Amazon EBS over-provisioned volumes	Cost optimization	C0r6dfpM03
Amazon EBS under-provisioned volumes	Performance	C0r6dfpM04
AWS Lambda over-provisioned functions for memory size	Cost optimization	C0r6dfpM05
AWS Lambda under-provisioned functions for memory size	Performance	C0r6dfpM06

You must opt in your AWS account for Compute Optimizer so that these checks can receive data from your Lambda and Amazon EBS resources. For more information, see [Opt in AWS Compute Optimizer for Trusted Advisor checks](#).

Updates to the Exposed Access Keys check

Trusted Advisor updated the following check on April 25, 2022.

Check name	Check category	Check ID
Exposed Access Keys	Security	12Fnkp18Y5

Trusted Advisor now refreshes this check for you automatically. This check can't be refreshed manually from the Trusted Advisor console or the AWS Support API. If your application or code refreshes this check for your AWS account, we recommend that you update it to no longer refresh this check. Otherwise, you will receive the `InvalidParameterValue` error.

Any access keys that you excluded before this update will no longer be excluded and will appear as affected resources. You can't exclude access keys from your check results. For more information, see [Exposed Access Keys](#).

Note

If you created your AWS account after April 25, 2022, the check results for Exposed Access Keys initially shows the gray icon



even for unexposed access keys. This means that Trusted Advisor hasn't identified any changes to the check.

If Trusted Advisor identifies a resource at risk, the status changes to the action recommended icon



After you fix or delete the resource, the check result shows the check mark icon



Updated checks for AWS Direct Connect

Trusted Advisor updated the following checks on March 29, 2022.

Check name	Check category	Check ID
AWS Direct Connect Connection Redundancy	Fault tolerance	0t121N1Ty3

Check name	Check category	Check ID
AWS Direct Connect Location Redundancy	Fault tolerance	8M012Ph3U5
AWS Direct Connect Virtual Interface Redundancy	Fault tolerance	4g3Nt5M1Th

- The value for the **Region** column now shows the AWS Region code instead of the full name. For example, resources in US East (N. Virginia) will now have the `us-east-1` value.
- The value for the **Time Stamp** column now appears in the RFC 3339 format, such as `2022-03-30T01:02:27.000Z`.
- Resources that don't have any detected problems will now appear in the check table. These resources will have a check mark icon (✓) next to them.

Previously, only resources that Trusted Advisor recommended that you investigate appeared in the table. These resources have a warning icon (⚠) next to them.

AWS Security Hub CSPM controls added to the AWS Trusted Advisor console

AWS Trusted Advisor added 111 Security Hub CSPM controls to the **Security** category on January 18, 2022.

You can view your findings for Security Hub CSPM controls from the AWS Foundational Security Best Practices security standard. This integration doesn't include controls that have the **Category: Recover > Resilience**.

For more information about this feature, see [Viewing AWS Security Hub CSPM controls in AWS Trusted Advisor](#).

New checks for Amazon EC2 and AWS Well-Architected

Trusted Advisor added the following checks on December 20, 2021.

- Amazon EC2 instances consolidation for Microsoft SQL Server
- Amazon EC2 instances over-provisioned for Microsoft SQL Server
- Amazon EC2 instances with Microsoft SQL Server end of support
- AWS Well-Architected high risk issues for cost optimization
- AWS Well-Architected high risk issues for performance
- AWS Well-Architected high risk issues for security
- AWS Well-Architected high risk issues for reliability

For more information, see the [AWS Trusted Advisor check reference](#).

Updated check name for Amazon OpenSearch Service

Trusted Advisor updated the name for the Amazon OpenSearch Service Reserved Instance Optimization check on September 8, 2021.

The check recommendations, category, and ID are the same.

Check name	Check category	Check ID
Amazon OpenSearch Service Reserved Instance Optimization	Cost optimization	7ujm6yhn5t

Note

If you use Trusted Advisor for Amazon CloudWatch metrics, the metric name for this check is also updated. For more information, see [Creating Amazon CloudWatch alarms to monitor AWS Trusted Advisor metrics](#).

Added checks for Amazon Elastic Block Store volume storage

Trusted Advisor added the following checks on June 8, 2021.

Check name	Check category	Check ID
EBS General Purpose SSD (gp3) Volume Storage	Service limits	dH7RR016J3
EBS Provisioned IOPS SSD (io2) Volume Storage	Service limits	gI7MM017J2

Added checks for AWS Lambda

Trusted Advisor added the following checks on March 8, 2021.

Check name	Check category	Check ID
AWS Lambda Functions with Excessive Timeouts	Cost optimization	L4dfs2Q3C3
AWS Lambda Functions with High Error Rates	Cost optimization	L4dfs2Q3C2
AWS Lambda Functions Using Deprecated Runtimes	Security	L4dfs2Q4C5
AWS Lambda VPC-enabled Functions without Multi-AZ Redundancy	Fault tolerance	L4dfs2Q4C6

For more information about how to use these checks with Lambda, see [Example AWS Trusted Advisor workflow to view recommendations](#) in the *AWS Lambda Developer Guide*.

Trusted Advisor check removal

Trusted Advisor removed the following check for the AWS GovCloud (US) Region on March 8, 2021.

Check name	Check category	Check ID
EC2 Elastic IP Addresses	Service limits	aW9HH018J6

Updated checks for Amazon Elastic Block Store

Trusted Advisor updated the unit of Amazon EBS volume from gibibyte (GiB) to tebibyte (TiB) for the following checks on March 5, 2021.

Note

If you use Trusted Advisor for Amazon CloudWatch metrics, the metric names for these five checks are also updated. For more information, see [Creating Amazon CloudWatch alarms to monitor AWS Trusted Advisor metrics](#).

Check name	Check category	Check ID	Updated CloudWatch metric for ServiceLimit
EBS Cold HDD (sc1) Volume Storage	Service limits	gH5CC0e3J9	Cold HDD (sc1) volume storage (TiB)
EBS General Purpose SSD (gp2) Volume Storage	Service limits	dH7RR016J9	General Purpose SSD (gp2) volume storage (TiB)
EBS Magnetic (standard) Volume Storage	Service limits	cG7HH017J9	Magnetic (standard) volume storage (TiB)
EBS Provisioned IOPS SSD (io1) Volume Storage	Service limits	gI7MM017J9	Provisioned IOPS (SSD) storage (TiB)

Check name	Check category	Check ID	Updated CloudWatch metric for ServiceLimit
EBS Throughput Optimized HDD (st1) Volume Storage	Service limits	wH7DD013J9	Throughput Optimized HDD (st1) volume storage (TiB)

Trusted Advisor check removal

Note

Trusted Advisor removed the following checks on November 18, 2020.

Checks removed on November 18, 2020	Check category	Check ID
EC2Config Service for EC2 Windows Instances	Fault tolerance	V77i0L1Bqz
ENA Driver Version for EC2 Windows Instances	Fault tolerance	TyfdMXG69d
NVMe Driver Version for EC2 Windows Instances	Fault tolerance	yHAGQJV9K5
PV Driver Version for EC2 Windows Instances	Fault tolerance	Wnwm9I15bG
EBS Active Volumes	Service limits	fH7LL017J9

Amazon Elastic Block Store no longer has a limit on the number of volumes that you can provision.

You can monitor your Amazon EC2 instances and verify they are up to date by using [AWS Systems Manager Distributor](#), other third-party tools, or write your own scripts to return driver information for Windows Management Instrumentation (WMI).

Trusted Advisor check removal

Trusted Advisor removed the following check on February 18, 2020.

Check name	Check category	Check ID
Service Limits	Performance	eW7HH017J9

AWS Support App in Slack

You can use the AWS Support App to manage your AWS support cases in Slack. Invite your team members to chat channels, respond to case updates, and chat directly with support agents. Use the AWS Support App to manage support cases quickly in Slack.

Use the AWS Support App to do the following:

- Create, update, search for, and resolve support cases in Slack channels
- Attach files to support cases
- Share support case details with your team without leaving the Slack channel
- Start a live chat session with support agents

When you create, update, or resolve a support case in the AWS Support App, the case is also updated in the AWS Support Center Console. You don't need to sign in to the Support Center Console to manage your support cases separately.

Notes

- The response times for support cases are the same, whether you created the case from Slack or from the Support Center Console.
- You can create a support case for account and billing support, and technical support.

Topics

- [Prerequisites](#)
- [Authorize a Slack workspace](#)
- [Configuring a Slack channel](#)
- [Creating support cases in a Slack channel](#)
- [Replying to support cases in Slack](#)
- [Join a live chat session with Support](#)
- [Searching for support cases in Slack](#)
- [Resolving a support case in Slack](#)
- [Reopening a support case in Slack](#)

- [Deleting a Slack channel configuration from the AWS Support App](#)
- [Deleting a Slack workspace configuration from the AWS Support App](#)
- [AWS Support App in Slack commands](#)
- [View AWS Support App correspondences in the AWS Support Center Console](#)
- [Creating AWS Support App in Slack resources with AWS CloudFormation](#)

Prerequisites

You must meet the following requirements to use the AWS Support App in Slack:

- You have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan. You can find your support plan from the AWS Support Center Console or from the [Support plans](#) page. For more information, see [Compare AWS Support plans](#).
- You have a [Slack](#) workspace and channel for your organization. You must be a Slack workspace administrator, or have permission to add apps to that Slack workspace. For more information, see the [Slack Help Center](#).
- You sign in to the AWS account as an AWS Identity and Access Management (IAM) user or role with the required permissions. For more information, see [Managing access to the AWS Support App widget](#).
- You will need to create an IAM role that has the required permissions to perform actions for you. The AWS Support App uses this role to make API calls to different services. For more information, see [Managing access to the AWS Support App](#).

Topics

- [Managing access to the AWS Support App widget](#)
- [Managing access to the AWS Support App](#)

Managing access to the AWS Support App widget

You can attach an AWS Identity and Access Management (IAM) policy to grant an IAM user permission to configure the AWS Support App widget in the AWS Support Center Console.

For more information about how to add a policy to an IAM entity, see [Adding IAM identity permissions \(console\)](#) in the IAM User Guide.

Note

You can also sign in as the root user in your AWS account, but we don't recommend that you do this. For more information about root user access, see [Safeguard your root user credentials and don't use them for everyday tasks](#) in the *IAM User Guide*.

Example IAM policy

You can attach the following policy to an entity, such as an IAM user or group. This policy allows a user to authorize a Slack workspace and configure Slack channels in the Support Center Console.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "supportapp:GetSlackOAuthParameters",
        "supportapp:RedeemSlackOAuthCode",
        "supportapp:DescribeSlackChannels",
        "supportapp:ListSlackWorkspaceConfigurations",
        "supportapp:ListSlackChannelConfigurations",
        "supportapp:CreateSlackChannelConfiguration",
        "supportapp>DeleteSlackChannelConfiguration",
        "supportapp>DeleteSlackWorkspaceConfiguration",
        "supportapp:GetAccountAlias",
        "supportapp:PutAccountAlias",
        "supportapp>DeleteAccountAlias",
        "supportapp:UpdateSlackChannelConfiguration",
        "iam:ListRoles"
      ],
      "Resource": "*"
    }
  ]
}
```

Permissions required to connect the AWS Support App to Slack

The AWS Support App includes permission-only actions that don't directly correspond to an API operation. These actions are indicated in the [Service Authorization Reference](#) with [permission only].

The AWS Support App uses the following API actions to connect to Slack and then lists your *public* Slack channels in the AWS Support Center Console:

- `supportapp:GetSlackOAuthParameters`
- `supportapp:RedeemSlackOAuthCode`
- `supportapp:DescribeSlackChannels`

These API actions are not intended to be called by your code. Therefore, these API actions are not included in the AWS CLI and AWS SDKs.

Managing access to the AWS Support App

After you have permissions to the AWS Support App widget, you must also create an AWS Identity and Access Management (IAM) role. This role performs actions from other AWS services for you, such as the AWS Support API and Service Quotas.

You then attach an IAM policy to this role so that the role has the required permissions to complete these actions. You choose this role when you create your Slack channel configuration in the Support Center Console.

Users in your Slack channel have the same permissions that you grant to the IAM role. For example, if you specify read-only access to your support cases, then users in your Slack channel can view your support cases, but can't update them.

Important

When you request a live chat with a support agent and choose new private channel as your live chat channel preference, the AWS Support App creates a separate Slack channel. This Slack channel has the same permissions as the channel where you created the case or initiated the chat.

If you change the IAM role or the IAM policy, your changes apply to the Slack channel that you configured and to any new live chat Slack channels that the AWS Support App creates for you.

Follow these procedures to create your IAM role and policy.

Topics

- [Use an AWS managed policy or create a customer managed policy](#)
- [Create an IAM role](#)
- [Troubleshooting](#)

Use an AWS managed policy or create a customer managed policy

To grant your role permissions, you can use either an AWS managed policy or a customer managed policy.

Tip

If you don't want to create a policy manually, we recommend that you use an AWS managed policy instead and skip this procedure. Managed policies automatically have the required permissions for the AWS Support App. You don't need to update the policies manually. For more information, see [AWS managed policies for AWS Support App in Slack](#).

Follow this procedure to create a customer managed policy for your role. This procedure uses the JSON policy editor in the IAM console.

To create a customer managed policy for the AWS Support App

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 **Policies**.
3. Choose **Create policy**.
4. Choose the **JSON** tab.
5. Enter your JSON, and then replace the default JSON in the editor. You can use the [example policy](#).

6. Choose **Next: Tags**.
7. (Optional) You can use tags as key–value pairs to add metadata to the policy.
8. Choose **Next: Review**.
9. On the **Review policy** page, enter a **Name**, such as *AWSSupportAppRolePolicy*, and a **Description** (optional).
10. Review the **Summary** page to see the permissions that the policy allows and then choose **Create policy**.

This policy defines the actions that the role can take. For more information, see [Creating IAM policies \(console\)](#) in the *IAM User Guide*.

Example IAM policy

You can attach the `AWSSupportAppFullAccess` policy to your IAM role. This policy allows the role to have full permissions to all required actions for the AWS Support App. After you configure a Slack channel with the role, any user in your channel has the same permissions.

To view the permissions for the `AWSSupportAppFullAccess` policy, see [AWSSupportAppFullAccess](#) in the *AWS Managed Policy Reference*.

Note

For a list of AWS managed policies, see [AWS managed policies for AWS Support App in Slack](#).

You can update the policy to remove a permission from the AWS Support App.

For descriptions for each action, see the following topics in the *Service Authorization Reference*:

- [Actions, resources, and condition keys for AWS Support](#)
- [Actions, resources, and condition keys for Service Quotas](#)
- [Actions, resources, and condition keys for AWS Identity and Access Management](#)

Create an IAM role

After you have your policy, you must create an IAM role, and then attach the policy to that role. You choose this role when you create a Slack channel configuration in the Support Center Console.

Important

To create or update a Slack channel configuration, you must have the `iam:PassRole` permission for the role that you choose. This permission allows you to pass the role to the AWS Support App. If you don't have this permission, you receive an error similar to the following when you create or update a channel configuration:

An error occurred when we tried to process your request. User: *user* is not authorized to perform: supportapp:CreateSlackChannelConfiguration.

For more information, see [Granting a user permissions to pass a role to an AWS service](#) in the *IAM User Guide*.

To create a role for the AWS Support App

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**, and then choose **Create role**.
3. For **Select trusted entity**, choose **AWS service**.
4. Choose **AWS Support App**.
5. Choose **Next: Permissions**.
6. Enter the policy name. You can choose the AWS managed policy or choose a customer managed policy that you created, such as *AWSSupportAppRolePolicy*. Then select the check box next to the policy.
7. Choose **Next: Tags**.
8. (Optional) You can use tags as key–value pairs to add metadata to the role.
9. Choose **Next: Review**.
10. For **Role name**, enter a name, such as *AWSSupportAppRole*.
11. (Optional) For **Role description**, enter a description for the role.
12. Review the role and then choose **Create role**. You can now choose this role when you configure a Slack channel in the Support Center Console. See [Configuring a Slack channel](#).

For more information, see [Creating a role for an AWS service](#) in the *IAM User Guide*.

Troubleshooting

See the following topics to manage access to the AWS Support App.

Contents

- [I want to restrict specific users in my Slack channel from specific actions](#)
- [When I configure a Slack channel, I don't see the IAM role that I created](#)
- [My IAM role is missing a permission](#)
- [A Slack error says that my IAM role isn't valid](#)
- [The AWS Support App says that I'm missing an IAM role for Service Quotas](#)

I want to restrict specific users in my Slack channel from specific actions

By default, users in your Slack channel have the same permissions specified in the IAM policy that you attach to the IAM role that you create. This means anyone in the channel has read or write access to your support cases, whether or not they have an AWS account or an IAM user.

We recommend the following best practices:

- Configure private Slack channels with the AWS Support App
- Only invite users to your channel who need access to your support cases
- Use an IAM policy that has the minimum required permissions to the AWS Support App. See [AWS managed policies for AWS Support App in Slack](#).

When I configure a Slack channel, I don't see the IAM role that I created

If your IAM role doesn't appear in the **IAM role for the AWS Support App** list, this means that the role doesn't have the AWS Support App as a trusted entity, or that the role was deleted. You can update the existing role, or create another one. See [Create an IAM role](#).

My IAM role is missing a permission

The IAM role that you create for your Slack channel needs permissions to perform the actions that you want. For example, if you want your users in Slack to create support cases, the role must have the `support:CreateCase` permission. The AWS Support App assumes this role to perform these actions for you.

If you receive an error about a missing permission from the AWS Support App, verify that the policy attached to your role has the required permission.

See the previous [Example IAM policy](#).

A Slack error says that my IAM role isn't valid

Verify that you chose the correct role for your channel configuration.

To verify your role

1. Sign in to the AWS Support Center Console at <https://console.aws.amazon.com/support/app#/config> page.
2. Choose the channel that you configured with the AWS Support App.
3. From the **Permissions** section, find the IAM role name that you chose.
 - To change the role, choose **Edit**, choose another role, and then choose **Save**.
 - To update the role or the policy attached to the role, sign in to the [IAM console](#).

The AWS Support App says that I'm missing an IAM role for Service Quotas

You must have the `AWSServiceRoleForServiceQuotas` role in your account to request quota increases from Service Quotas. If you receive an error about a missing resource, complete one of the following steps:

- Use the [Service Quotas](#) console to request a quota increase. After you make a successful request, Service Quotas creates this role for you automatically. Then, you can use the AWS Support App to request quota increases in Slack. For more information, see [Requesting a quota increase](#).
- Update the IAM policy attached to your role. This grants the role permission to Service Quotas. The following section in the [Example IAM policy](#) allows the AWS Support App to create the Service Quotas role for you.

```
{
  "Effect": "Allow",
  "Action": "iam:CreateServiceLinkedRole",
  "Resource": "*",
  "Condition": {
    "StringEquals": {"iam:AWSServiceName": "servicequotas.amazonaws.com"}
  }
}
```

If you delete the IAM role that you configure for your channel, you must manually create the role or update the IAM policy to allow the AWS Support App to create one for you.

Authorize a Slack workspace

After you authorize your workspace and give the AWS Support App permission to access it, you then need an AWS Identity and Access Management (IAM) role for your AWS account. The AWS Support App uses this role to call API operations from [AWS Support](#) and [Service Quotas](#) for you. For example, the AWS Support App uses the role to call the `CreateCase` operation to create a support case for you in Slack.

Notes

- Your Slack channel inherits permissions from the IAM role. All users in the Slack channel have the same permissions that are specified in the IAM policy attached to the role. For example, if your IAM policy grants full read and write permissions to your support cases, anyone in your Slack channel can create, update, and resolve your support cases. If your IAM policy allows the role read-only permissions, then users in your Slack channel can only view your support cases.
- We recommend that you add the Slack workspaces and channels that you need to manage your support operations. We recommend that you configure private channels and only invite required users.

You must authorize each Slack workspace that you want to use for your AWS account. If you have multiple AWS accounts, you must sign in to each account and repeat the following procedure to authorize the workspace. If your account belongs to an organization in AWS Organizations and you want to authorize multiple accounts, skip to [Authorize multiple accounts](#).

To authorize the Slack workspace for your AWS account

1. Sign in to the [AWS Support Center Console](#) and choose **Slack configuration**.
2. On the **Getting started** page, choose **Authorize workspace**.
3. If you're not already signed in to Slack, on the **Sign in to your workspace** page, enter your workspace name, and then choose **Continue**.
4. On the **AWS Support is requesting permission to access the your-workspace-name Slack** page, choose **Allow**.

 Note

If you can't allow Slack to access your workspace, make sure that you have permissions from your Slack administrator to add the AWS Support App to the workspace. See [Prerequisites](#).

On the **Slack configuration** page, your workspace name appears under **Workspaces**.

5. (Optional) To add more workspaces, choose **Authorize workspace** and repeat steps 3-4. You can add up to five workspaces to your account.
6. (Optional) By default, your AWS account ID number appears as the account name in your Slack channel. To change this value, under **Account name**, choose **Edit**, enter your account name, and then choose **Save**.

 Tip

Use a name that you and your team can easily recognize. The AWS Support App uses this name to identify your account in the Slack channel. You can update this name at any time.

Edit account name ✕

Choose an account name that you can easily recognize in Slack. This name won't appear in your AWS account settings.

Account name

Maximum 30 characters (5 remaining)

Example Usage:

Account name being used by Support Slack App Bot

- **AWS account:** aws-administrator-account (ID: 123456789012)

Cancel Save

Your workspace and account name appear on the **Slack configuration** page.

Slack configuration

Workspaces

Delete Authorize workspace Add multiple accounts ↺

Workspace

troubleshooting

Account name

Delete Edit

Name used in Slack

aws-administrator-account

Authorize multiple accounts

If your account belongs to an organization in AWS Organizations, you can authorize multiple AWS accounts to use the same Slack workspace. You can use AWS CloudFormation templates to automate this process, or you can manually configure each account using the AWS Support App API.

Use AWS CloudFormation templates

You can use [AWS CloudFormation](#) or [Terraform](#) to create your AWS Support App resources across multiple accounts. This is the recommended approach for managing multiple accounts at scale.

For more information, see [Creating AWS Support App in Slack resources with AWS CloudFormation](#).

Use the AWS Support App API

If you prefer not to use AWS CloudFormation templates, you can manually authorize multiple accounts using the AWS Support App API. Complete the following steps for each member account that you want to add to the Slack workspace.

To manually authorize multiple accounts

1. In the management account (the account where you created the Slack workspace), authorize the Slack workspace. For more information, see the procedure in [Authorize a Slack workspace](#).
2. In each member account, call the [RegisterSlackWorkspaceForOrganization](#) API operation to register the Slack workspace for that account.

This operation adds the member account to the Slack workspace that was authorized by the management account.

3. In each member account, create an IAM role with the required permissions for the AWS Support App. For more information, see [Managing access to the AWS Support App](#).
4. In each member account, call the [CreateSlackChannelConfiguration](#) API operation to configure a Slack channel for that account.

This operation associates the Slack channel with the member account and specifies the IAM role that the AWS Support App uses to call AWS Support and Service Quotas operations.

Note

Each member account must be part of the same organization in AWS Organizations as the management account that authorized the Slack workspace.

Configuring a Slack channel

After you authorize your Slack workspace, you can configure your Slack channels to use the AWS Support App.

In the channel where you invite and add the AWS Support App, you can create and search for cases, and receive case notifications. This channel displays case updates, including newly created or resolved cases, added correspondence, and shared case details.

Your Slack channel inherits permissions from the IAM role. All users in the Slack channel have the same permissions that are specified in the IAM policy attached to the role. For example, if your IAM policy grants full read and write permissions to your support cases, anyone in your Slack channel can create, update, and resolve your support cases. If your IAM policy allows the role read-only permissions, then users in your Slack channel can only view your support cases.

You can add up to 20 channels per AWS account. Each Slack channel supports up to 100 AWS accounts. This means that only 100 accounts can add the same Slack channel to the AWS Support App. To minimize distractions, we recommend adding only the accounts necessary for managing your organization's support cases.

Each AWS account must configure a Slack channel separately in the AWS Support App. This way, the AWS Support App can access the support cases in that AWS account. If another AWS account in your organization already invited the AWS Support App to that Slack channel, skip to step 3.

Note

You can configure channels that are part of [Slack Connect](#) and channels that are shared with multiple workspaces. However, only the first workspace that configured the shared channel for an AWS account can use the AWS Support App. The AWS Support App returns an error message if you try to configure the same Slack channel for another workspace.


To configure a Slack channel

1. From your Slack application, choose the Slack channel that you want to use with the AWS Support App.
2. Complete the following steps to invite the AWS Support App to your channel:

- a. Open the context (right-click) menu on the channel name, and then choose **View channel details**.
- b. Choose the **Integrations** tab, and then choose **Add an App**.
- c. To search for the app, enter **AWS Support App**.
- d. Choose **Add** next to the **AWS Support App**.



3. Sign in to the [Support Center Console](#) and choose **Slack configuration**.
4. Choose **Add channel**.
5. On the **Add channel** page, under **Workspace**, choose the workspace name that you previously authorized. You can choose the refresh icon if the workspace name doesn't appear in the list.



6. Under **Slack channel**, for **Channel type**, choose one of the following:
 - **Public** – Under **Public channel**, choose the Slack channel that you invited the AWS Support App to (step 2). If your channel doesn't appear in the list, choose the refresh icon and try again.
 - **Private** – Under **Channel ID**, enter the ID or the URL of the Slack channel that you invited the AWS Support App to.

 Tip

To find the channel ID, open the context (right-click) menu for the channel name in Slack, and then choose **Copy**, and then choose **Copy link**. Your channel ID is the value that looks like **C01234A5BCD**.

7. Under **Channel configuration name**, enter a name that easily identifies your Slack channel configuration for the AWS Support App. This name appears only in your AWS account and doesn't appear in Slack. You can rename your channel configuration later.

Your Slack channel type might look like the following example.

▼ **Slack channel**

Channel Type

☐ Public
Choose a public channel from the list.

☒ Private
A channel member must invite a user to join or view.

Channel ID

C01234A5BCD

Channel configuration name

Choose a name that you can easily identify. You can change the name at any time.

MyTroubleshootingChannel

 **Tip**
Tip To find the channel ID, right-click your channel name in Slack, choose **Copy** and then choose **Copy link**. Your channel ID is the value that looks like **C01234A5BCD**.

8. Under **Permissions**, for **IAM role for the AWS Support App in Slack**, choose a role that you created for the AWS Support App. Only roles that have the AWS Support App as a trusted entity appear in the list.

▼ **Permissions**

IAM role for the AWS Support App

Choosing another IAM role for this Slack channel configuration can affect the permissions for any chat channels created from this troubleshooting channel. You can verify that your role has the required permissions. [Learn more](#)

▼ 

 Note

If you haven't created a role or don't see your role in the list, see [Managing access to the AWS Support App](#).

9. Under **Notifications**, specify how to get notified for cases.
 - **All cases** – Get notified for all case updates.
 - **High-severity cases** – Get notified for only cases that affect a production system or higher. For more information, see [Choosing an initial support case severity level](#).
 - **None** – Don't get notified for case updates.
10. (Optional) If you choose **All cases** or **High-severity cases**, you must select at least one of the following options:
 - **New and reopened cases**
 - **Case correspondences**
 - **Resolved cases**

The following channel receives case notifications for all case updates in Slack.

▼ Notifications

Additional case notifications
Choose when to get notified for cases created and updated.

☒ All cases

☐ High-severity cases

☐ None

Notification types
Get notified for the following types of cases that are created.

☒ New and reopened cases

☒ Case correspondences

☒ Resolved cases

Note: You will receive notifications in your Slack channel for all case updates for this account.

11. Review your configuration and choose **Add channel**. Your channel appears in the **Slack configuration** page.

Update your Slack channel configuration

After you configured your Slack channel, you can update them later to change the IAM role or case notification.

To update your Slack channel configuration

1. Sign in to the [Support Center Console](#) and choose **Slack configuration**.
2. Under **Channels**, choose the channel configuration that you want.
3. On the **channelName** page, you can do the following tasks:
 - Choose **Rename** to update your channel configuration name. This name only appears in your AWS account and won't appear in Slack.
 - Choose **Delete** to delete the channel configuration from the AWS Support App. See [Deleting a Slack channel configuration from the AWS Support App](#).
 - Choose **Open in Slack** to open the Slack channel in your browser.
 - Choose **Edit** to change the IAM role or notifications.

Creating support cases in a Slack channel

After you authorize your Slack workspace and add your Slack channel, you can create a support case in your Slack channel.

To create a support case in Slack

1. In your Slack channel, enter the following command:

```
/awssupport create
```

2. In the **Create a support case** dialog box, do the following:
 - a. If you configured more than one account for this Slack channel, for **AWS account**, choose the account ID. If you created an account name, this value appears next to the account ID. For more information, see [Authorize a Slack workspace](#).
 - b. For **Subject**, enter a title for the support case.
 - c. For **Description**, describe the support case. Provide details, such as how you're using an AWS service and what troubleshooting steps you tried.

aws **Create a support case**  

Step 1 of 3

You can create a case with AWS Support for technical and account-related issues.

AWS account

dev-ops-production (ID:123456789012) 

Subject

AWS resources issue

Description

I can't find my resource in my AWS account. 2457 

Note: You can add attachments after step 3 when you confirm the case.

Cancel **Next**

3. Choose **Next**.
4. On the **Create a support case** dialog box, specify the following options:
 - a. Choose the **Issue type**.
 - b. Choose the **Service**.
 - c. Choose the **Category**.
 - d. Choose the **Severity**.
 - e. Review your case details and choose **Next**.

The following example shows a technical support case for Alexa Services.

The screenshot shows a modal window titled "Create a support case" with the AWS logo. It is at "Step 2 of 3". The form contains the following fields:

- Issue type:** A dropdown menu with "Technical support" selected.
- Service:** A dropdown menu with "Alexa Services" selected.
- Category:** A dropdown menu with "APIs" selected.
- Severity:** A dropdown menu with "General guidance" selected.

At the bottom right, there are two buttons: "Back" (outlined) and "Next" (solid green).

- For **Contact language**, choose your preferred language for your support case.

Note

Japanese language support isn't available for live chat in Slack for account and billing cases.

- For **Contact method**, choose **Email and Slack notifications** or **Live chat in Slack**.

The following example shows how to choose a live chat in Slack.

 **Create a support case** ×

Step 3 of 3

Contact language

English ▼

Contact method

☒ Live chat in Slack

☐ Email and Slack notifications

Live chat channel preference

New private channel ▼

 A new channel will be created for your live chat session, and anyone who is invited to the channel can see previous chat history.

Additional chat members (optional)

Add chat members

You will be added to the live chat automatically.

Back Review

- a. If you choose **Live chat in Slack**, choose **New private channel** or **Current channel** as your **Live chat channel preference**. **New private channel** will create a separate private channel for you to chat with the AWS Support agent, and **Current channel** will use a thread in the current channel for you to chat with the AWS Support agent.
- b. (Optional) If you choose **Live chat in Slack**, you can enter the names of other Slack members. For **New private channel**, the AWS Support App will automatically add you and selected members to the new channel. For **Current channel**, the AWS Support App will automatically tag you and selected members in the chat thread when the AWS Support agent joins.

⚠ Important

- We recommend that you only add chat members that you want to have access to your support case details and chat history.
- If you start a new live chat session for an existing support case, the AWS Support App uses the same chat channel or thread that was used for a previous live chat. The AWS Support App also uses the same live chat channel preference that was used previously.
- The **Current channel** option is only available if the chat is requested from a private channel. We recommend that you only use this option if you want all channel members to have access to your chat.

7. (Optional) For **Additional contacts to notify**, enter email addresses to also receive updates about this support case. You can add up to 10 email addresses.
8. Choose **Review**.
9. In the Slack channel, review the case details. You can do the following:
 - Choose **Edit** to change the case details.
 - Add a file to your case. To do so, follow these steps:
 - a. Choose **Attach file**, choose the **+** icon in Slack, and choose **Your computer**.
 - b. Navigate to and choose your file.
 - c. In the **Upload a file** dialog box, enter @awssupport, and press the send message  icon.

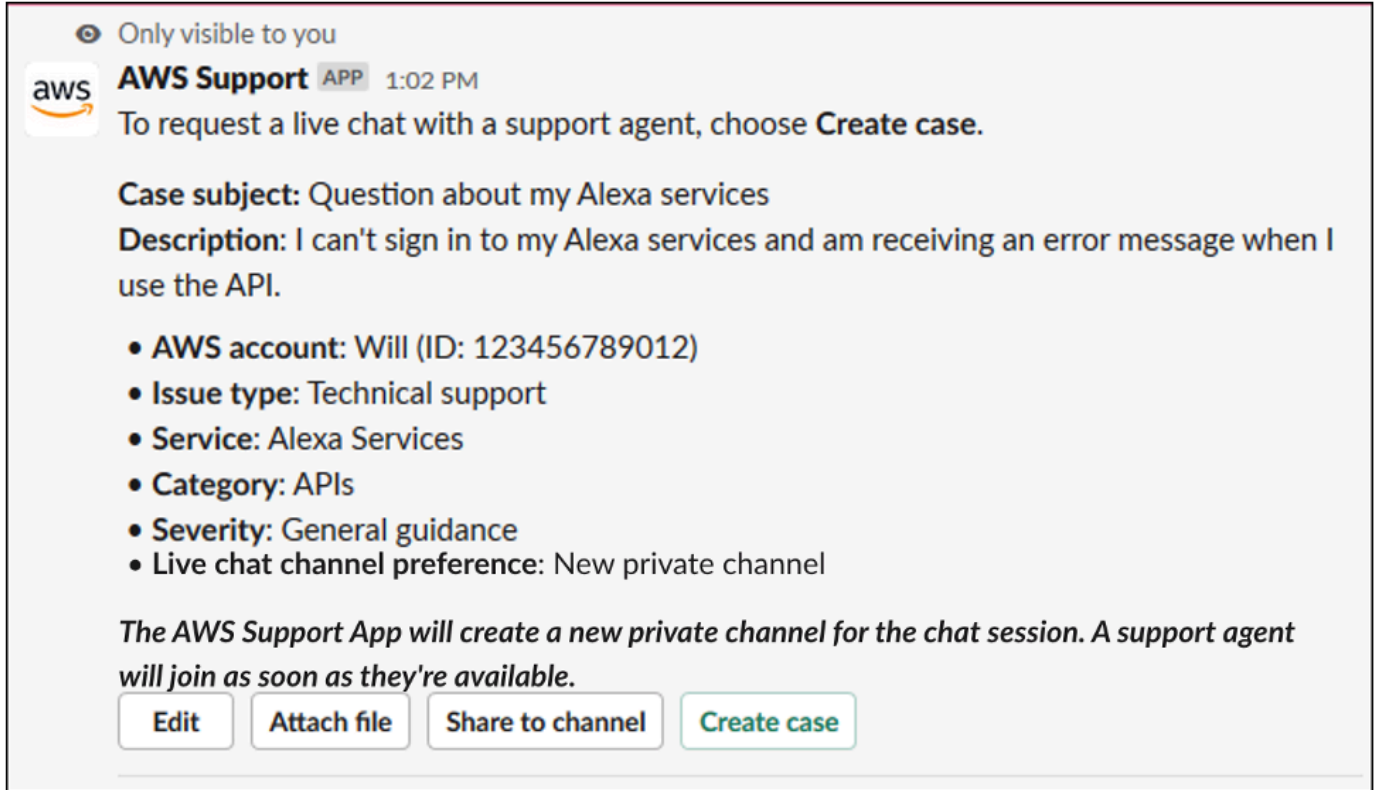
** Notes**

- You can attach up to three files. Each file can be up to 5 MB.
- If you attach a file to your support case, you must submit your case within 1 hour. If you don't, you must add the files again.

- Choose **Share to channel** to share the case details with others in the Slack channel. You can use this option to share the case details with your team before you create the case.

10. Review your case details, and then choose **Create case**.

The following example shows a technical support case for Alexa Services.



Only visible to you

 **AWS Support** APP 1:02 PM

To request a live chat with a support agent, choose **Create case**.

Case subject: Question about my Alexa services

Description: I can't sign in to my Alexa services and am receiving an error message when I use the API.

- **AWS account:** Will (ID: 123456789012)
- **Issue type:** Technical support
- **Service:** Alexa Services
- **Category:** APIs
- **Severity:** General guidance
- **Live chat channel preference:** New private channel

The AWS Support App will create a new private channel for the chat session. A support agent will join as soon as they're available.

[Edit](#) [Attach file](#) [Share to channel](#) [Create case](#)

After you create a support case, it might take a few minutes for your case details to appear.

11. When your support case is updated, you can choose **See details** to view your case information. You can then do the following:

- Choose **Share to channel** to share the case details with others in the Slack channel.
- Choose **Reply** to add a correspondence.
- Choose **Resolve case**.

 Note

If you didn't choose to receive automatic case updates in Slack, you can search for the support case to find the **See details** option.

Replying to support cases in Slack

You can add updates to your case such as case details and attachments, and reply to responses from the support agent.

Note


- You can also use the AWS Support Center Console to reply to support agents. For more information, see [Legacy experience: Updating, resolving, and reopening your case](#).
- You cannot add correspondences to cases from chat channels created by the AWS Support App. Live chat channels only send messages to agents during the live chat.

To reply to a support case in Slack

1. In your Slack channel, choose the case that you want to respond to. You can enter `/awssupport search` to find your support case.
2. Choose **See details** next to the case that you want.
3. At the bottom of the case details, choose **Reply**.



4. In the **Reply to case** dialog box, enter a brief description of the issue in the **Message** field. Then choose **Next**.



The screenshot shows a dialog box titled "aws Reply to case" with a close button (X) and a share icon. It is labeled "Step 1 of 2". The "Case subject" is "AWS resources issue". The "Message" field contains the text: "I'm attaching a file to this case that can help you troubleshoot the issue." Below the message field is a note: "Note: You can add attachments after step 2 when you confirm the message." At the bottom right are "Cancel" and "Next" buttons.

5. Choose your contact method. The available contact methods depend on your case type and support plan.
6. (Optional) For **Additional contacts to notify**, enter additional email addresses that you want to receive updates about this support case. You can add up to 10 email addresses.
7. Choose **Review**. You can then choose if you want to edit your reply, attach files, or share to the channel.
8. When you're ready to reply, choose **Send message**.
9. (Optional) To view previous correspondence for your case, choose **Previous correspondence**. To view shortened messages, choose **Show full message**.

Example: Reply to a case in Slack

Only visible to you

 **AWS Support** APP 10:53 AM

To respond to this case, review and then choose **Send message**.

Case subject: AWS resources issue

Message: I'm attaching a file to this case that can help you troubleshoot the issue.

We will contact you by email and Slack notifications within 24 hours.

Additional contacts to notify: None

Attachments: error-log

✓ You successfully attached 1 file. Choose **Create case** within 1 hour to include the file with your case.

Join a live chat session with Support

When you request a live chat for your case, you choose to either use a new chat channel or a thread in the current channel for you and the AWS Support agent. Use this chat channel or thread to communicate with the support agent and any others that you invited to the live chat.

Important

Anyone who joins a channel with a live chat can view details about the specific support case and the chat history. Its a best practice to add only users that require access to your support cases. Any member of a chat channel or thread can also participate in an active chat.

Note

Live chat channels and threads also receive notifications when a correspondence is added to the case outside of the live chat session. This occurs before, during, and after a chat session, so you can use a chat channel or thread to monitor all updates for a case. If you

chose to use a new chat channel, use the configuration channel that you invited the AWS Support App to reply to these correspondences.

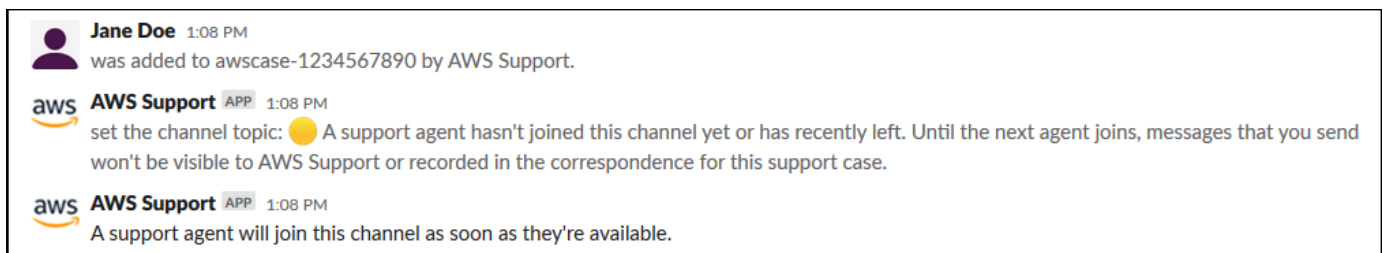
To join a live chat session with Support in a new channel

1. In the Slack application, navigate to the channel that the AWS Support App creates for you. The channel name includes your support case ID, such as *awscase-1234567890*.

Note

The AWS Support App adds a pinned message to the live chat channel that contains details about your support case. From the pinned message, you can end the chat or resolve the case. You can find all pinned messages in this channel under the channel name.

2. When the support agent joins the channel, you can chat about your support case. Until a support agent joins the channel, the agent won't see messages in that chat and the messages don't appear in your case correspondence.



The screenshot shows a Slack channel interface. At the top, a message from Jane Doe (1:08 PM) states: "was added to awscase-1234567890 by AWS Support." Below this, a message from the AWS Support App (1:08 PM) says: "set the channel topic: 🟡 A support agent hasn't joined this channel yet or has recently left. Until the next agent joins, messages that you send won't be visible to AWS Support or recorded in the correspondence for this support case." At the bottom, another message from the AWS Support App (1:08 PM) states: "A support agent will join this channel as soon as they're available."

3. (Optional) Add other members to the chat channel. By default, chat channels are private.
4. After the support agent joins the chat, the chat channel is active and the AWS Support App records the chat.

You can chat with the agent about your support case and upload any file attachments to the channel. The AWS Support App automatically saves your files and chat log to your case correspondence.

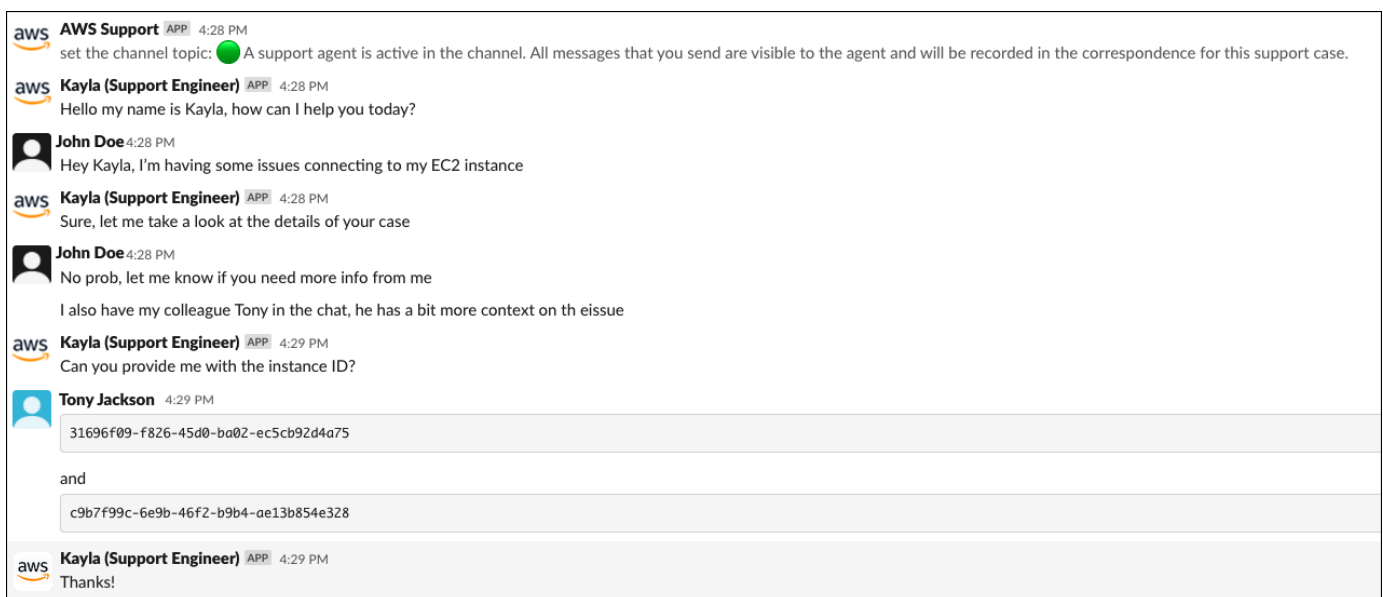
Note

When you chat with a support agent, note the following differences in Slack for the AWS Support App:

- Support agents can't view shared messages or threads. To share text from a message or thread, enter the text as a new message.
- If you edit or delete a message, the agent still sees the original message. You must enter your new message again to show the revision.

Example: Live chat session

The following is an example of a live chat session with a support agent to fix a connectivity issue for two Amazon Elastic Compute Cloud (Amazon EC2) instances.



5. (Optional) To stop the live chat, choose **End chat**. The support agent leaves the channel and the AWS Support App stops recording the live chat. You can find the chat history attached to the case correspondence for this support case.
6. If the issue is resolved, you can choose **Resolve case** from the pinned message or enter `/awssupport resolve`.

Example: End a live chat

The following pinned message shows the case details about an Amazon EC2 instance. You can find the pinned messages under the Slack channel name.

★ Pinned by AWS Support

 **AWS Support** APP 2:33 PM

This is a live chat channel for the following case.

Case subject: Cannot connect to ec2 instance (Case ID: 6887208841)

Description: The ec2 instance i-09f00da444 was unable to lookup our dns region. We had full access yesterday. Now we get "Access denied" message.

Case created by Jane Doe (in Slack)

- **Status:** Unassigned
- **Created:** 02/16/2021, 2:33PM PST
- **AWS account:** Instance Management (ID: 111122223333)
- **Issue type:** Technical support
- **Service:** Elastic Compute Cloud (EC2-Linux)
- **Category:** SSH Issue
- **Severity:** Production system impaired

[End chat](#) [Resolve case](#)

Example: Correspondence notification in chat channel

The following is an example of a live chat channel receiving a notification when the another collaborator adds an update after the chat has ended.

 **AWS Support** APP 3:28 PM

A correspondence was added to the case after the live chat ended.

Correspondence: Can you link me the article one more time? *Correspondence added by [redacted] (in Slack)*

Status: Unassigned

To reply to this correspondence, go to this [thread](#) or sign in to the AWS Support Center. [Learn more](#)

 **AWS Support**

The following case was created for account [redacted] (ID: [redacted]).

[redacted] (Case ID: [redacted])

[View original message](#)

Thread in # [redacted] Jan 23rd | [View message](#)

 **docs.aws.amazon.com**

[Replying to support cases in Slack - AWS Support](#)

Use the AWS Support App to reply to your support cases in Slack.

The notification will indicate the chat status (requested, in progress, or ended) and whether the correspondence was added by an agent or by another collaborator. The Support App will also attempt to link back to the original Slack thread or channel where this chat was requested. You can [reply to this case](#) from that channel, or any other channel with access to this case.

To join a live chat session with Support in the current channel

1. In the Slack application, navigate to the thread in the current channel that the AWS Support App uses for the chat. In most cases, this will be the thread that started when the case was first created.
2. When the support agent joins the thread, you can chat about your support case. Until a support agent joins the thread, the agent won't see messages in that thread, and the messages won't appear in your case correspondence when the chat ends.

Note

Messages sent to this channel outside of the chat thread are never seen by Support, even while a chat is active.

Thread aws-support-communications**AWS Support** APP < 1 minute ago

The following case was created for account [REDACTED].

Question about my Alexa services (Case ID: [REDACTED])



A support agent hasn't joined this chat session yet or has recently left

Get updates

See details

End chat

Reply

Resolve case

7 replies

**AWS Support** APP < 1 minute ago

@Jane Doe requested a chat for this case.

Question about my Alexa services (Case ID: [REDACTED])

**AWS Support** APP < 1 minute ago

A support agent will join this chat session as soon as they're available.

**Tip:** Editing and deleting messages is not supported during the chat session. Support agents will still see original messages.

3. (Optional) Tag other channel members to notify them on the chat thread.
4. After the support agent joins the chat, the chat thread is active and the AWS Support App records the chat. Similar to the new chat channel option, you can chat with the agent about your support case and upload any file attachments to the thread. The AWS Support App automatically saves your files and chat log to your case correspondence.
5. (Optional) To stop the live chat, choose End chat from the initial message for this thread. The support agent leaves the thread and the AWS Support App stops recording the live chat. You can find the chat history attached to the case correspondence for this support case.
6. If the issue is resolved, you can choose Resolve case from the initial message for this thread.

Thread  aws-support-communications

 **AWS Support** APP < 1 minute ago

The following case was created for account .

Question about my Alexa services (Case ID: )

 A support agent hasn't joined this chat session yet or has recently left

[Get updates](#) [See details](#) [End chat](#) [Reply](#) [Resolve case](#)

7 replies

Searching for support cases in Slack

From your Slack channel, you can search for support cases from your AWS account and from other accounts that configured the same channel and workspace. For example, if your account (123456789012) and your coworker's account (111122223333) have configured the same workspace and channels in the AWS Support Center Console, you can use the AWS Support App to search for each other's support cases.

To filter your search results, you can use the following options:

- Account ID
- Case ID
- Case status
- Contact language
- Date range

Example: Search for cases in Slack

The following example shows how to search by **Filter options** for a single account by specifying the date range, case status, and contact language.

👁 Only visible to you

 **AWS Support** APP 1:07 PM

Search for cases created by account **aws-administrator-account** (ID: 123456789012).

I want to search for cases by:

☒ Filter options

☐ Case ID

Date range:

 10/01/2022 

 Today 

Case status:

All cases 

Case created in:

English 

Search

To search for a support case in Slack

1. In the Slack channel, enter the following command:

```
/awssupport search
```

2. For the **I want to search for cases by:** option, choose one of the following:

A. **Filter options** – You can filter cases with the following options:

- **AWS account** – This list only appears if you have multiple accounts in the channel.
- **Date range** – The date the case was created.
- **Case status** – The current case status, such as **All open cases** or **Resolved**.
- **Case created in** – The contact language for the case.

- B. **Case ID** – Enter the case ID. You can only enter one case ID at a time. If you have multiple accounts in the channel, choose the AWS account to search for the case.
3. Choose **Search**. Your search results appear in Slack.

Use your search results

The following example returns three support cases from one AWS account.

 Only visible to you

 **AWS Support** APP 1:51 PM

3 results found for cases created from 10/01/2022 to 12/28/2022 with AWS account aws-administrator-account (ID:123456789012).

Case subject: Can't retrieve info about my certificate (Case ID: 1234567890)
Created: 10/25/2022, 10:30 PM UTC
Status: Resolved

See details

Case subject: Question about my AWS account bill (Case ID: 4445556660)
Created: 10/14/2022, 7:35 PM UTC
Status: Resolved

See details

Case subject: Technical support for EC2 instances (Case ID: 9087654321)
Created: 10/13/2022, 2:28 PM UTC
Status: In progress

See details

Edit Search

Share to channel

After you receive your search results, you can do the following:

To use your search results

1. Choose **Edit Search** to change your previous filter options or case ID.
2. Choose **Share to channel** to share the search results with the channel.
3. Choose **See details** for more information about a case. You can choose **Show full message** to view the rest of the latest correspondence.

4. If you searched by **Filter options**, search results can return multiple cases. Choose **Next 5 results** or **Previous 5 results** to view the next or previous 5 cases.

Example: Resolved support case

The following example shows a resolved support case for an account and billing issue after choosing **See details**.

 Only visible to you

This case was created on 10/14/2022, 10:30 PM UTC.

Case subject: Question about my AWS account bill (Case ID: 4445556660)

Description: I have a question about a charge for my last statement

- **Status:** Resolved
- **AWS account:** aws-administrator-account (ID: 123456789012)
- **Issue type:** Account and billing support
- **Service:** Academy
- **Category:** Account/Lab access issue
- **Severity:** General question
- **Language:** English

Correspondence:

Amazon Web Services, 10/25/2022, 10:30 PM UTC

This case has been resolved. Please contact us again if you need further assistance.

Share to channel


Reopen case

Resolving a support case in Slack

If you don't need your support case anymore, or you fixed the issue, you can resolve a support case directly in Slack. This also resolves the case in the AWS Support Center Console. After you resolve a case, you can reopen the case later.

To resolve a support case in Slack

1. In your Slack channel, navigate to the support case. See [Searching for support cases in Slack](#).
2. Choose **See details** for the case.
3. Choose **Resolve case**.
4. In the **Resolve case** dialog box, choose **Resolve case**. You can reopen a case in the Slack channel or from the Support Center Console.



Reopening a support case in Slack

After you resolve a support case, you can reopen the case from Slack.

To reopen a support case in Slack

1. Find the support case to reopen in Slack. See [Searching for support cases in Slack](#).
2. Choose **See details**.
3. Choose **Reopen case**.
4. In the **Reopen case** dialog box, enter a brief description of the issue in the **Message** field.
5. Choose **Next**.

aws **Reopen case** X

Step 1 of 2

Case subject: Question about my AWS bill

Message

I still see this issue in my account. 2463

Note: You can add attachments after step 2 when you confirm the message.

Cancel Next

6. (Optional) Enter additional contacts.
7. Choose **Review**.
8. Review your case details, and then choose **Send message**. Your case reopens. If you requested a new live chat with a support agent, Slack uses the same chat channel or thread as the one that was used for a previous live chat. If you requested a live chat in a new channel and you haven't had one so far, a new chat channel opens. If you requested a live chat in the current channel and you haven't had one so far, a thread in the current channel is used.

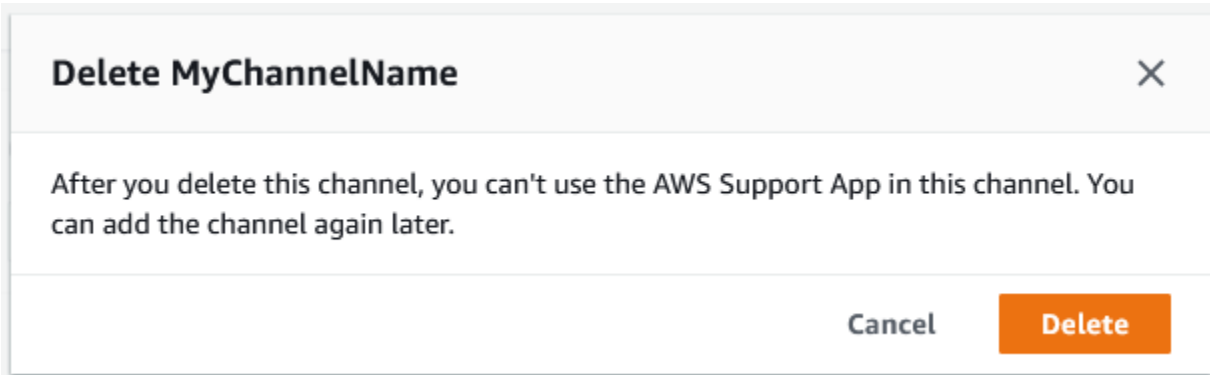
Deleting a Slack channel configuration from the AWS Support App

You can delete a channel configuration from the AWS Support App if you don't need it. This action only removes the channel from the AWS Support App and the AWS Support Center Console. Your channel isn't deleted from Slack.

You can add up to 20 channels for your AWS account. If you already reached this quota, you must delete a channel before you can add another one.

To delete a Slack channel configuration

1. Sign in to the [Support Center Console](#) and choose **Slack configuration**.
2. On the **Slack configuration** page, under **Channels**, choose the channel name, and then choose **Delete**.
3. In the **Delete channel name** dialog box, choose **Delete**. You can add this channel to the AWS Support App again later.



Deleting a Slack workspace configuration from the AWS Support App

You can delete a workspace configuration from the AWS Support App if you don't need it. This action only removes the workspace from the AWS Support App and the AWS Support Center Console. Your workspace isn't deleted from Slack.

You can add up to 5 workspaces for your AWS account. If you already reached this quota, you must delete a Slack workspace before you can add another one.


Note

If you added channels from this workspace to the AWS Support App, you must first delete these channels before you can delete the workspace. See [Deleting a Slack channel configuration from the AWS Support App](#).

To delete a Slack workspace configuration

1. Sign in to the [AWS Support Center Console](#) and choose **Slack configuration**.

2. On the **Slack configuration** page, under **Slack workspaces**, choose **Delete a workspace**.
3. In the **Delete Slack workspace** dialog box, choose the Slack workspace name, and then choose **Delete**. You can add the workspace to your AWS account again later.



AWS Support App in Slack commands

Slack channel commands

You can enter the following commands in the Slack channel where you invited the AWS Support App. This Slack channel name also appears as a configured channel in the AWS Support Center Console.

`/awssupport create` or `/awssupport create-case`

Create a support case.

`/awssupport search` or `/awssupport search-case`

Search for cases. You can search for support cases for the AWS accounts that configured the AWS Support App for the same Slack channel.

`/awssupport quota` or `/awssupport service-quota-increase`

Request a service quota increase.

Live chat channel commands

You can enter the following commands in the live chat channel. This is the channel that the AWS Support App creates for you if you choose a new channel for your chat with Support. Chat channels include your support case ID, such as *aws-case-1234567890*.

Note

The following commands are not available when using a thread in the current channel for a live chat. Instead, use the buttons attached to the initial thread message to end a chat, invite a new agent, or resolve the case.

`/awssupport endchat`

Remove the support agent and end the live chat session.

`/awssupport invite`

Invite a new support agent to this channel.

`/awssupport resolve`

Resolve this support case.

View AWS Support App correspondences in the AWS Support Center Console

When you create, update, or resolve support cases for your account in the Slack channel, you can also sign in to the Support Center Console to view your cases. You can view the case correspondences to determine whether the case was updated in the Slack channel, view the chat history with a support agent, and find any attachments that you uploaded from Slack.

To view case correspondences from Slack

1. Sign in to the [AWS Support Center Console](#) for your account.
2. Choose your support case.
3. In the **Correspondence**, you can view whether the case was created and updated from the Slack channel.

Example: Support case

In the following screenshot, Jane Doe reopened a support case in Slack. This correspondence appears for the support case in the Support Center Console.

Correspondence	
MyIAMRole (Role) Thu Feb 24 2022 09:09:33 GMT-0800 (Pacific Standard Time)	I am having difficulty retrieving information about my certificates. _Case created by JaneDoe (in Slack)_

Creating AWS Support App in Slack resources with AWS CloudFormation

AWS Support App in Slack is integrated with AWS CloudFormation, a service that helps you to model and set up your AWS resources so that you can spend less time creating and managing your resources and infrastructure. You create a template that describes all the AWS resources that you want (such as your AccountAlias and SlackChannelConfiguration), and CloudFormation provisions and configures those resources for you.

When you use CloudFormation, you can reuse your template to set up your AWS Support App resources consistently and repeatedly. Describe your resources once, and then provision the same resources over and over in multiple AWS accounts and Regions.

AWS Support App and CloudFormation templates

To provision and configure resources for AWS Support App and related services, you must understand [CloudFormation templates](#). Templates are formatted text files in JSON or YAML. These templates describe the resources that you want to provision in your CloudFormation stacks. If you're unfamiliar with JSON or YAML, you can use CloudFormation Designer to help you get started with CloudFormation templates. For more information, see [What is CloudFormation Designer?](#) in the *AWS CloudFormation User Guide*.

AWS Support App supports creating your AccountAlias and SlackChannelConfiguration in CloudFormation. For more information, including examples of JSON and YAML templates for the

AccountAlias and SlackChannelConfiguration resources, see the [AWS Support App resource type reference](#) in the *AWS CloudFormation User Guide*.

Create Slack configuration resources for your organization

You can use CloudFormation templates to create the resources that you need for the AWS Support App. If you're the management account for your organization, you can use the templates to create these resources for your member accounts in AWS Organizations.

For example, you might use a template to create the same Slack workspace configuration for all accounts in the organization, but then use separate templates to create different Slack channel configurations for specific AWS accounts or organizational units (OUs). You can also use a template to create a Slack workspace configuration so that member accounts can then configure the Slack channels that they want for their AWS accounts.

You can choose whether to use CloudFormation templates or not. If you don't use CloudFormation templates, you can complete the following manual steps instead:

- Create the AWS Support App resources in the AWS Support Center Console.
- Call the [RegisterSlackWorkspaceForOrganization](#) API operation to register a Slack workspace for your account. The CloudFormation stack calls this API operation for you.
- Create an IAM role with the required permissions for each member account.
- Call the [CreateSlackChannelConfiguration](#) API operation to configure a Slack channel for each member account.

For detailed manual steps, see [Authorize multiple accounts](#).

Follow these procedures to upload the CloudFormation template to your organization. You can use the example templates from the [AWS Support App resource type reference](#) page.

The templates tell CloudFormation to create the following resources:

- A [Slack channel configuration](#).
- A [Slack workspace configuration](#).
- An [IAM role](#) with the AWSSupportSlackAppCFNRole name. The AWSSupportAppFullAccess AWS managed policy is attached.

Contents

- [Update your CloudFormation templates for Slack](#)
- [Create a stack for the management account](#)
- [Create a stack set for your organization](#)

Update your CloudFormation templates for Slack

To get started, use the following templates to create your stack. You must replace the templates with valid values for your Slack workspace and channel.

Note

We don't recommend the use of the template to create an [AccountAlias](#) resource for your organization. The AccountAlias resource uniquely identifies an AWS account in the AWS Support App. Your member accounts can enter an account name in the Support Center Console. For more information, see [Authorize a Slack workspace](#).

To update your CloudFormation templates for Slack

1. If you're the management account for an organization, you must manually authorize a Slack workspace for your account before your member accounts can use CloudFormation to create the resources. If you haven't already done so, see [Authorize a Slack workspace](#).
2. From the [AWS Support App resource type reference](#) page, copy the JSON or YAML template for the resource that you want.
3. In a text editor, paste the template into a new file.
4. In the template, specify the parameters that you want. At a minimum, replace the values for the following fields:
 - TeamId with your Slack workspace ID
 - ChannelId with the Slack channel ID
 - ChannelName with a name to identify the Slack channel configuration

Tip

To find the workspace and channel IDs, open your Slack channel in a browser. In the URL, your workspace ID is the first identifier and the channel ID is the second.

For example, in `https://app.slack.com/client/T012ABCDEFGH/C01234A5BCD`, T012ABCDEFGH is the workspace ID and C01234A5BCD is the channel ID.

5. Save the file as either a JSON or YAML file.

Create a stack for the management account

Next, you must create a stack for the management account in the organization. This step calls the [RegisterSlackWorkspaceForOrganization](#) API operation for you and authorizes the workspace with Slack.

Note

We recommend that you upload the Slack workspace configuration template that you updated in the previous procedure for the management account. You don't need to upload the Slack channel configuration template unless you're also configuring the management account to use the AWS Support App.

To create a stack for the management account

1. Sign in to the AWS Management Console as the management account for your organization.
2. Open the CloudFormation console at <https://console.aws.amazon.com/cloudformation>.
3. If you haven't already, in the **Region selector**, choose one of the following AWS Regions:
 - Europe (Frankfurt)
 - Europe (Ireland)
 - Europe (London)
 - US East (N. Virginia)
 - US East (Ohio)
 - US West (Oregon)
 - Asia Pacific (Singapore)
 - Asia Pacific (Tokyo)
 - Canada (Central)
4. Follow the procedure to create a stack. For more information, see [Creating a stack on the CloudFormation console](#).

After CloudFormation successfully creates the stack, you can use the same template to create a stack set for your organization.

Create a stack set for your organization

Next, use the same template for the Slack workspace configuration to create a stack set with service-managed permissions. You can use stack sets to create the stack for your entire organization or specify the OUs that you want. For more information, see [Create a stack set](#).

This procedure also calls the [RegisterSlackWorkspaceForOrganization](#) API operation for you. This API operation authorizes the workspace with Slack for the member accounts.

To create a stack set for your organization

1. Sign in to the AWS Management Console as the management account for your organization.
2. Open the CloudFormation console at <https://console.aws.amazon.com/cloudformation>.
3. If you haven't already, in the **Region selector**, choose the same AWS Region that you used in the previous procedure.
4. In the navigation pane, choose **StackSets**.
5. Choose **Create StackSet**.
6. On the **Choose a template** page, keep the default options for the following options:
 - For **Permissions**, keep **Service-managed permissions**.
 - For **Prerequisite - Prepare template**, keep **Template is ready**.
7. Under **Specify template**, choose **Upload a template file**, and then choose **Choose file**.
8. Choose the file and then choose **Next**.
9. On the **Specify StackSet details** page, enter a stack name such as **support-app-slack-workspace**, enter a description, and then choose **Next**.
10. On the **Configure StackSet options** page, keep the default options and then choose **Next**.
11. On the **Set deployment options** page, for **Add stacks to stack set**, keep the default **Deploy new stacks** option.
12. For **Deployment targets**, choose if you want to create the stack for the entire organization or specific OUs. If you choose an OU, enter the OU ID.
13. For **Specify regions**, enter only *one* of the following AWS Regions:

- Europe (Frankfurt)
- Europe (Ireland)
- Europe (London)
- US East (N. Virginia)
- US East (Ohio)
- US West (Oregon)
- Asia Pacific (Singapore)
- Asia Pacific (Tokyo)
- Canada (Central)

 Notes:

- To streamline your workflow, we recommend that you use the same AWS Region that you chose in step 3.
- Choosing more than one AWS Region can cause conflicts with creating your stack.

14. For **Deployment options**, for **Failure tolerance - optional**, enter the number of accounts where the stacks can fail before CloudFormation stops the operation. We recommend that you enter the number of accounts that you want to add, minus one. For example, if your specified OU has 10 member accounts, enter 9. This means that even if CloudFormation fails the operation 9 times, at least one account will succeed.
15. Choose **Next**.
16. On the **Review** page, review your options, and then choose **Submit**. You can check the status of your stack on the **Stack instances** tab.
17. (Optional) Repeat this procedure to upload a template for a Slack channel configuration. The example template also creates the IAM role and attaches an AWS managed policy. This role has the required permissions to access other services for you. For more information, see [Managing access to the AWS Support App](#).

If you don't create a stack set to create the Slack channel configuration, your member accounts can manually configure the Slack channel. For more information, see [Configuring a Slack channel](#).

After CloudFormation creates the stacks, each member account can sign in to the Support Center Console and find their configured Slack workspaces and channels. They can then use the AWS Support App for their AWS account. See [Creating support cases in a Slack channel](#).

Tip

If you need to upload a new template, we recommend that you use the same AWS Region that you specified before.

Learn more about CloudFormation

To learn more about CloudFormation, see the following resources:

- [AWS CloudFormation](#)
- [AWS CloudFormation User Guide](#)
- [CloudFormation API Reference](#)
- [AWS CloudFormation Command Line Interface User Guide](#)

Create AWS Support App resources by using Terraform

You can also use [Terraform](#) to create the AWS Support App resources for your AWS account. Terraform is an infrastructure-as-code tool that you can use for your cloud applications. You can use Terraform to create AWS Support App resources instead of deploying a CloudFormation stack to an account.

After you install Terraform, you can specify the AWS Support App resources that you want. Terraform calls the [RegisterSlackWorkspaceForOrganization](#) API operation to register a Slack workspace for you and creates your resources. You can then sign in the Support Center Console and find your configured Slack workspaces and channels.

Notes

- If you're the management account for an organization, you must manually authorize a Slack workspace for your account before your member accounts can use Terraform to create the resources. If you haven't already done so, see [Authorize a Slack workspace](#).

- Unlike CloudFormation stack sets, you can't use Terraform to create the AWS Support App resources for an OU in your organization.
- You can also find the event history for these updates from Terraform in AWS CloudTrail. The eventSource for these events will be `cloudcontrolapi.amazonaws.com` and `supportapp.amazonaws.com`. For more information, see [Logging AWS Support App in Slack API calls using AWS CloudTrail](#).

Learn more

To learn more about Terraform, see the following topics:

- [Terraform installation](#)
- [Terraform tutorial: Build infrastructure for AWS](#)
- [awscc_support_app_account_alias](#)
- [awscc_supportapp_slack_workspace_configuration](#)
- [awscc_supportapp_slack_channel_configuration](#)

Security in AWS Support

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. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS compliance programs](#). To learn about the compliance programs that apply to AWS Support, 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 company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Support. The following topics show you how to configure Support to meet your security and compliance objectives. You also learn how to use other Amazon Web Services that help you to monitor and secure your Support resources.

Topics

- [Data protection in AWS Support](#)
- [Security for your AWS Support cases](#)
- [AWS Support authorization](#)
- [Identity and access management for AWS Support](#)
- [Incident response](#)
- [Logging and monitoring in AWS Support and AWS Trusted Advisor](#)
- [Compliance validation for AWS Support](#)
- [Resilience in AWS Support](#)
- [Infrastructure security in AWS Support](#)
- [Configuration and vulnerability analysis in Support](#)

Data protection in AWS Support

The AWS [shared responsibility model](#) applies to data protection in Support. 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 AWS Support 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.

⚠ Important

In the case correspondence, never share sensitive information, such as credentials, credit cards, signed URLs, or personally identifiable information.

Security for your AWS Support cases

When you create a support case, you own the information that you include in your support case. AWS doesn't access your AWS account data without your permission. AWS doesn't share your information with third parties.

When you create a support case, note the following:

- AWS Support uses the permissions defined in the `AWSServiceRoleForSupport` service-linked role to call other AWS services that troubleshoot customer issues for you. For more information, see [Using service-linked roles for AWS Support](#) and [AWS managed policy: `AWSSupportServiceRolePolicy`](#).
- You can view API calls to AWS Support that occurred in your AWS account. For example, you can view log information when someone in your account creates or resolves a support case. For more information, see [Logging AWS Support API calls with AWS CloudTrail](#).
- You can use the AWS Support API to call the `DescribeCases` API. This API returns support case information, such as the case ID, the create and resolve date, and correspondences with the support agent. You can view case details for up to 24 months after the case was created. For more information, see [DescribeCases](#) in the *AWS Support API Reference*.
- Your support cases follow [Compliance validation for AWS Support](#).
- When you create a support case, AWS doesn't gain access your account. If necessary, support agents use a screen-sharing tool to view your screen remotely and identify and troubleshoot problems. This tool is view-only. AWS Support can't act for you during the screen-share session. You must give consent to share a screen with a support agent. For more information, see the [AWS Support FAQs](#).
- You can change your AWS Support plan to get the help that you need for your account. For more information, see [Compare AWS Support Plans](#) and [Changing your AWS Support plan](#).

AWS Support authorization

AWS Support authorization gives you auditable control over AWS Support access to information about your services that might contain your data during support case resolution.

A support permit is a customer-managed resource that defines the scope and conditions of authorized access. Each support permit is cryptographically signed by using a customer-managed AWS Key Management Service (AWS KMS) key that you control. AWS CloudTrail logs all actions taken on your resources under your authorization during the granted access period, giving you a complete audit trail.

Important

The information obtained through AWS Support authorization might be accessed from outside the AWS Region where your data is stored.

AWS Support authorization supports two authorization scenarios to fit your operational needs.

Proactive authorization

You create support permits in advance to pre-authorize access within a defined scope. When AWS Support needs access to resolve an issue, a signed authorization is automatically issued if a matching support permit exists. This reduces resolution time during incidents.

Reactive authorization

AWS Support submits a support permit request in the Support Center when access to information about your services is needed. You review the request and either approve it by creating a scoped support permit or reject it. If you don't approve the request, AWS Support can't access the requested information. This gives you per-case control over access decisions.

Topics

- [AWS Support authorization concepts](#)
- [Getting started with AWS Support authorization](#)
- [Configuring AWS KMS keys for AWS Support authorization](#)
- [Managing support permits](#)

- [Managing support permit requests](#)
- [Discovering support actions](#)
- [Monitoring AWS Support authorization with AWS CloudTrail](#)
- [AWS Support authorization quotas](#)

AWS Support authorization concepts

The following concepts help you understand how AWS Support authorization works.

Support permits

A support permit is a resource that defines the scope and conditions under which AWS Support can access information about your services that might contain your data during support case resolution. Without a support permit, AWS Support can't access this information. Each support permit specifies which actions AWS Support can take and the resources on which AWS Support can take those actions, along with optional time-based conditions.

Support permits can't be modified after creation. To change the scope of an existing authorization, delete the support permit and create a new one.

Support permit scope

The support permit scope defines what AWS Support is authorized to do. A support permit scope contains three elements:

Actions

The specific operations that AWS Support can perform. You can permit all actions or specify individual actions.

Resources

The AWS resources for which AWS Support can access information about your services that might contain your data. You can permit all resources in the Region or specify individual resource ARNs.

Conditions

Optional time-based constraints. You can set an `allowAfter` time (the earliest time AWS Support is able to use the support permit to authorize an action) and an `allowBefore` time

(the latest time AWS Support is able to use the support permit to authorize an action) to create a time window. A support permit supports a maximum of 2 conditions.

Signing keys

The signing key is a customer-managed AWS KMS asymmetric key that provides verification for support permits. The key must meet the following requirements:

- Key spec: ECC_NIST_P384
- Key usage: SIGN_VERIFY
- Location: Same Region as the support permit.

Your private key material never leaves AWS KMS. Your KMS key is called to perform the signing operation and receives only the resulting signature.

Support permit requests

A support permit request is a notification from AWS Support requesting access to information about your services that might contain your data for a specific resource to help resolve a support case. If you don't approve the request, AWS Support can't access the requested information and might not be able to proceed with the investigation.

When AWS Support submits a support permit request, you receive notifications through Amazon EventBridge events and AWS CloudTrail log entries. You can then review the requested scope and approve or reject the request.

Permit states

A support permit transitions through the following states:

ACTIVE

The support permit is active. AWS Support can receive a signed authorization—a cryptographically signed token that permits a specific action, for actions within the support permit scope.

INACTIVE

New signed authorizations aren't issued for INACTIVE support permits. A support permit transitions to INACTIVE automatically when the linked support case closes. Support permits can

also transition to INACTIVE if the linked support case closes while an account is isolated or a Region is opted out. To re-authorize access after these transitions, create a new support permit.

DELETING

Deletion initiated. New signed authorizations stop being issued. Due to eventual consistency, there might be a brief delay before the change takes effect. The support permit is then removed. To help prevent further signing operations, revoke the grant on the associated AWS KMS key. For more information, see [the section called “Revoking AWS Support authorization access”](#).

Permit request states

A support permit request transitions through the following states:

PENDING

The request is awaiting your review.

ACCEPTED

You created a support permit that satisfies the request.

REJECTED

You rejected the request by calling `RejectSupportPermitRequest`.

CANCELLED

AWS Support cancelled the request.

Actions and resources

Actions are service-specific operations that AWS Support can perform during support investigations, such as reading diagnostic logs or viewing operational metrics. Each action has a descriptive name, such as `rds:ReadClusterData` or `rds:ViewQueryLogsWithParameters`. You can view available actions for each service by using the `ListActions` and `GetAction` API operations. For more information, see [Discovering support actions](#).

Resources are identified by their ARN. When you scope a support permit, you specify the resource ARNs for which AWS Support can access service information, or you can permit access to service information for all resources in the Region.

Getting started with AWS Support authorization

This tutorial walks you through setting up AWS Support authorization and creating your first support permit. You complete the following steps:

This tutorial takes approximately 10 minutes to complete.

1. Create an AWS KMS signing key
2. Set up IAM permissions
3. Create your first support permit
4. Review support permit requests from AWS Support

Prerequisites

Before you begin, verify that you have the following:

- An active AWS account
- Permissions to create AWS KMS keys in your account
- Permissions to create IAM policies and attach them to users or roles

Step 1: Create an AWS KMS signing key

AWS Support authorization requires a AWS KMS key with key spec `ECC_NIST_P384` and key usage `SIGN_VERIFY`. The key must be in the same Region as your support permit.

Note

If you already have a AWS KMS key with key spec `ECC_NIST_P384` and key usage `SIGN_VERIFY` in the same Region, you can skip this step and use that key.

To create the key, run the following AWS CLI command:

```
aws kms create-key \  
  --key-usage SIGN_VERIFY \  
  --key-spec ECC_NIST_P384 \  
  --description "SupportAuthZ signing key" \  
  --tags TagKey=supportauthz:managed,TagValue=true \
```

```
--region us-east-1
```

Note the key ARN from the output. You use this value when you create support permits.

Include the required AWS Support authorization policy statements in the create-key call. For the required statements, see [Configuring AWS KMS keys for AWS Support authorization](#).

Example output

```
{
  "KeyMetadata": {
    "KeyId": "1234abcd-12ab-34cd-56ef-1234567890ab",
    "Arn": "arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab",
    "KeySpec": "ECC_NIST_P384",
    "KeyUsage": "SIGN_VERIFY"
  }
}
```

Step 2: Set up IAM permissions

Attach an IAM policy that grants permissions for AWS Support authorization API operations. The following example policy grants access to all AWS Support authorization actions.

The `supportauthz:RegisterKey` action lets you associate AWS KMS keys with support permits. This permission-only action must be present in your IAM policy for AWS Support authorization to function.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "supportauthz:CreateSupportPermit",
        "supportauthz:RegisterKey",
        "supportauthz>DeleteSupportPermit",
        "supportauthz:GetSupportPermit",
        "supportauthz:ListSupportPermits",
        "supportauthz:ListSupportPermitRequests",
        "supportauthz:RejectSupportPermitRequest",
        "supportauthz:GetAction",
        "supportauthz:ListActions"
      ]
    }
  ]
}
```

```

    ],
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "kms:DescribeKey"
    ],
    "Resource": "arn:aws:kms:us-
east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab",
    "Condition": {
      "StringEquals": {
        "kms:ViaService": "supportauthz.us-east-1.amazonaws.com"
      }
    }
  },
  {
    "Effect": "Allow",
    "Action": [
      "kms:CreateGrant"
    ],
    "Resource": "arn:aws:kms:us-
east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab",
    "Condition": {
      "StringEquals": {
        "kms:ViaService": "supportauthz.us-east-1.amazonaws.com",
        "kms:RetiringServicePrincipal": "us-
east-1.supportauthz.amazonaws.com",
        "kms:GranteeServicePrincipal": "us-
east-1.supportauthz.amazonaws.com"
      },
      "ForAllValues:StringEquals": {
        "kms:GrantOperations": [
          "DescribeKey",
          "GetPublicKey",
          "Sign"
        ]
      }
    }
  }
]
}

```

To attach this policy to an IAM user or role, see [Adding and removing IAM identity permissions](#) in the *AWS Identity and Access Management User Guide*.

Step 3: Create your first support permit

Create a support permit that authorizes AWS Support to access information about your services for specific resources.

Console

1. Sign in to the [AWS Support Center Console](#).
2. In the navigation pane, choose **Support authorization**.
3. In the **Granted access** section, choose **Preconfigure access**.
4. For **Access type**, choose one of the following:
 - **All supported resources in this Region** – Applies broad access to your entire account in the selected Region.
 - **Choose resource ARN** – Limits access to a specific resource that you identify by ARN.
5. (Optional) For **Details**, enter a **Name** and **Description** for the support permit.
6. For **Access duration**, choose one of the following:
 - **No expiration** – Access remains active until you revoke it.
 - **Custom duration** – Set a specific time window during which the support permit is valid.
7. For **Actions**, choose the actions that AWS Support is permitted to perform:
 - Select the **All actions** check box to permit all available actions on the covered resources. This removes the 10-action limit.
 - To choose specific actions, clear the **All actions** check box. Select a service from the **Services** list, and then select up to 10 individual actions from the actions table.
8. For **Signing key**, choose an AWS KMS key from the dropdown list. To create a new key, choose **Create KMS signing key**.
9. Choose **Submit**.

AWS CLI

Run the following command to create a time-bounded support permit for a specific resource:

```
aws supportauthz create-support-permit \
  --name "MyFirstPermit" \
  --signing-key-info '{"kmsKey": "arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab"}' \
  --permit '{
    "actions": {"allActions": {}},
    "resources": {"resources": ["arn:aws:rds:us-east-1:111122223333:cluster:my-aurora-cluster"]},
    "conditions": [
      {"allowAfter": "2026-06-01T00:00:00Z"},
      {"allowBefore": "2026-07-01T00:00:00Z"}
    ]
  }'
```

This support permit authorizes AWS Support to perform all available actions on the specified Amazon Aurora cluster during June 2026.

Step 4: Review support permit requests from AWS Support

When AWS Support needs access to information about your services and no matching support permit exists, AWS Support submits a support permit request. You can view pending requests and approve or reject them.

Console

1. Sign in to the [AWS Support Center Console](#).
2. In the navigation pane, choose **Support authorization**.
3. In the **Pending access requests** section, review the list of requests. Each request shows the associated support case, service, Support Permit Request ARN, status, and request access date.
4. (Optional) To filter requests, use the **Status** and **Service** dropdown filters, or search by resource in the search field.
5. To approve or reject a request, choose **Manage support access**.

AWS CLI

To list pending requests, run the following command:


```
aws supportauthz list-support-permit-requests
```

Example output

```
{
  "supportPermitRequests": [
    {
      "arn": "arn:aws:supportauthz:us-east-1:111122223333:supportpermitrequest/req-1234abcd",
      "status": "PENDING",
      "supportCaseDisplayId": "case-12345678",
      "requestedActions": ["rds:ReadClusterData"],
      "requestedResources": ["arn:aws:rds:us-east-1:111122223333:cluster:my-aurora-cluster"],
      "createdAt": "2026-06-01T12:00:00Z"
    }
  ]
}
```

To approve this request, create a support permit that covers the requested scope:

```
aws supportauthz create-support-permit \
  --name "ApproveCase12345678" \
  --signing-key-info '{"kmsKey": "arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab"}' \
  --support-case-display-id "case-12345678" \
  --permit '{
    "actions": {"actions": ["rds:ReadClusterData"]},
    "resources": {"resources": ["arn:aws:rds:us-east-1:111122223333:cluster:my-aurora-cluster"]},
    "conditions": [
      {"allowBefore": "2026-06-15T00:00:00Z"}
    ]
  }'
```

To reject a request, run the following command. This notifies AWS Support that you don't want to accept the request. Rejecting a request doesn't prevent you from creating a support permit for the same actions and resources later.

```
aws supportauthz reject-support-permit-request \
```

```
--request-arn "arn:aws:supportauthz:us-east-1:111122223333:supportpermitrequest/  
request-id"
```

Review each request and decide whether to approve or reject it:

- **To approve:** Create a support permit that covers the requested scope. Include the `supportCaseDisplayId` parameter to automatically deactivate the support permit when the case closes.
- **To reject:** Call `RejectSupportPermitRequest` to notify AWS Support that you don't want to accept the request. Rejecting a request doesn't prevent you from creating a support permit for the same actions and resources later.

Next steps

After you complete this tutorial, see the following topics:

- To learn more about scoping support permits, see [Managing support permits](#).
- To monitor support actions on your resources, see [Monitoring AWS Support authorization with AWS CloudTrail](#).

Configuring AWS KMS keys for AWS Support authorization

Key requirements

A customer-managed AWS KMS key is required for cryptographic signing of support permits. Your key must meet the following requirements:

- Key spec: `ECC_NIST_P384`
- Key usage: `SIGN_VERIFY`
- The key must be in the same AWS Region as the support permit.

Your private key material never leaves AWS KMS. You create a grant on your key that allows it to call `DescribeKey`, `GetPublicKey`, and `Sign`. The grant is scoped to the support permit and retires when the support permit is deleted or deactivated.

How AWS Support authorization uses your AWS KMS key

When you create a support permit, your AWS KMS key is used in the following way:

1. Uses [forward access sessions](#) to call `DescribeKey` on your behalf, verifying that the key meets the required spec (ECC_NIST_P384) and usage (SIGN_VERIFY).
2. Creates a grant on the key by calling `CreateGrant` on your behalf. The grant allows `DescribeKey`, `GetPublicKey`, and `Sign` operations.
3. Verifies authorization when an AWS Support request matches an existing support permit by using the grant to call `Sign` on the key. The resulting signature verifies that AWS Support is authorized to perform the action.

Required key policy statements

Add the following policy statements to your AWS KMS key policy. These are the minimum permissions required for AWS Support authorization to use your key.

```
{
  "Sid": "Allows access to CreateGrant through SupportAuthZ",
  "Effect": "Allow",
  "Principal": {"AWS": "arn:aws:iam::111122223333:role/ExampleRole"},
  "Action": ["kms:CreateGrant"],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:ViaService": "supportauthz.us-east-1.amazonaws.com",
      "kms:GranteeServicePrincipal": "us-east-1.supportauthz.amazonaws.com",
      "kms:RetiringServicePrincipal": "us-east-1.supportauthz.amazonaws.com"
    },
    "ForAllValues:StringEquals": {
      "kms:GrantOperations": ["DescribeKey", "GetPublicKey", "Sign"]
    },
    "ArnLike": {
      "kms:GrantConstraintSourceArn": "arn:aws:supportauthz:us-east-1:111122223333:supportpermit/*"
    }
  }
},
{
  "Sid": "Allows access to DescribeKey through SupportAuthZ",
  "Effect": "Allow",
```

```
"Principal": {"AWS": "arn:aws:iam::111122223333:role/ExampleRole"},
"Action": ["kms:DescribeKey"],
"Resource": "*",
"Condition": {
  "StringEquals": {
    "kms:ViaService": "supportauthz.us-east-1.amazonaws.com"
  }
}
```

These statements grant the following permissions:

CreateGrant

Allows the creation of grants for `DescribeKey`, `GetPublicKey`, and `Sign` operations. The conditions restrict grant creation to requests that are made through the AWS Support authorization service and scoped to support permit resources in your account. AWS Support authorization uses [forward access sessions](#) to call `CreateGrant` as part of creating a support permit.

DescribeKey

Allows verification that the key meets the required spec and usage type when you create a support permit using forward access sessions.

Note

Replace 111122223333 with your AWS account ID and us-east-1 with the AWS Region where you create your support permits.

Revoking AWS Support authorization access

The recommended way to revoke signing permissions on your key is to revoke the grant. This prevents further signing operations without affecting other uses of the key.

To list and revoke grants for AWS Support authorization:

1. List grants on the key to find the grant ID:

```
aws kms list-grants \
```

```
--key-id 1234abcd-12ab-34cd-56ef-1234567890ab
```

Identify the grant by its name or the `GrantConstraints` source ARN that references your support permit.

2. Revoke the grant:

```
aws kms revoke-grant \  
  --key-id 1234abcd-12ab-34cd-56ef-1234567890ab \  
  --grant-id grant-id-from-list-grants
```

After you revoke the grant, new signed authorizations can no longer be issued for any support permit that uses this key. Because the AWS KMS API follows an eventual consistency model, there might be a brief delay before the change is available throughout AWS KMS.

Note

Revoking a grant doesn't delete the associated support permit. The support permit remains in `ACTIVE` state, but authorizations can't be signed for it. Grants are specific to each support permit. Creating a new support permit with the same AWS KMS key doesn't restore AWS Support authorization's ability to use the key for the original support permit.

Disabling or deleting the key

You can also disable or schedule deletion of the AWS KMS key to prevent AWS Support authorization from issuing signed authorizations. However, this affects all uses of the key, not just AWS Support authorization.

Important

Both AWS KMS and AWS Support authorization are eventually consistent. After you disable or schedule deletion of a key, it can take up to several minutes for the change to take full effect. During this period, signed authorizations might continue to be issued by using the key. Grant revocation is also eventually consistent.

If you disable a key, then you can re-enable it in the AWS KMS console or by calling `EnableKey`. If you delete a key, it can't be recovered.

Monitoring key usage

When your AWS KMS key is used to sign an authorization, AWS CloudTrail logs the signing operation in the key's event history. You can use these events to audit when and how often authorizations are signed on your behalf.

Managing support permits

Support permits define the scope and conditions under which AWS Support can access information about your services. This topic describes how to create, view, list, and delete support permits.

Creating a support permit

Create a support permit to define the scope and conditions under which AWS Support can access information about your services.

Console

1. Sign in to the [AWS Support Center Console](#).
2. In the navigation pane, choose **Support authorization**.
3. In the **Granted access** section, choose **Preconfigure access**.
4. For **Access type**, choose one of the following:
 - **All supported resources in this Region** – Applies broad access to your entire account in the selected Region.
 - **Choose resource ARN** – Limits access to a specific resource that you identify by ARN.
5. (Optional) For **Details**, enter a **Name** and **Description** for the support permit.
6. For **Access duration**, choose one of the following:
 - **No expiration** – Access remains active until you revoke it.
 - **Custom duration** – Set a specific time window during which the support permit is valid.
7. For **Actions**, choose the actions that AWS Support is permitted to perform:
 - Select the **All actions** check box to permit all available actions on the covered resources. This removes the 10-action limit.

- To choose specific actions, clear the **All actions** check box. Select a service from the **Services** list, and then select up to 10 individual actions from the actions table.
8. For **Signing key**, choose an AWS KMS key from the dropdown list. To create a new key, choose **Create KMS signing key**.
 9. Choose **Submit**.

AWS CLI

Use `CreateSupportPermit` to create a new support permit. The following parameters are available:

Required parameters: `name`, a unique name for the support permit (1–256 alphanumeric characters).

Optional parameters: `description`, a human-readable description (maximum 1,024 characters).

Create a support permit with specific actions and resources

```
aws supportauthz create-support-permit \
  --name "AuroraDBPermit" \
  --description "Access to Aurora cluster diagnostic data" \
  --signing-key-info '{"kmsKey": "arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab"}' \
  --support-case-display-id "case-12345678" \
  --permit '{
    "actions": {"actions": ["rds:ReadClusterData",
    "rds:ViewQueryLogsWithParameters"]},
    "resources": {"resources": ["arn:aws:rds:us-east-1:111122223333:cluster:my-cluster"]},
    "conditions": [
      {"allowAfter": "2026-06-01T00:00:00Z"},
      {"allowBefore": "2026-06-15T00:00:00Z"}
    ]
  }'
```

Scoping a support permit

You control the scope of a support permit by specifying actions, resources, and conditions.

Actions

You can permit actions in two ways:

- **allActions**: Permits all support actions for the specified resources.
- **actions**: A list of up to 10 specific action names. Use `ListActions` to view available actions.

Resources

You can specify resources in two ways:

- **allResourcesInRegion**: Permits access to service information for all resources in the Region.
- **resources**: A list of up to 10 specific resource ARNs for which AWS Support can access service information.

Conditions

You can define up to 2 time-based conditions:

- **allowAfter**: The support permit is not valid before this time (ISO 8601 format).
- **allowBefore**: The support permit is not valid after this time (ISO 8601 format).

Use both conditions together to create a time window during which the support permit is valid.

Viewing a support permit

Retrieve the configuration and status of a support permit.

Console

1. Sign in to the [AWS Support Center Console](#).
2. In the navigation pane, choose **Support authorization**.
3. In the **Granted access** section, locate the support permit. The table displays the support case, service, Support Permit ARN, name, status, creation date, and duration for each support permit.

AWS CLI

Use `GetSupportPermit` to retrieve the configuration and status of a support permit. You can identify the support permit by name or ARN.

```
aws supportauthz get-support-permit \
  --support-permit-identifier "AuroraDBPermit"
```

Listing support permits

View all support permits in your account and Region.

Console

1. Sign in to the [AWS Support Center Console](#).
2. In the navigation pane, choose **Support authorization**.
3. In the **Granted access** section, view the list of support permits. Use the **Status** and **Service** dropdown filters to narrow the results, or search by resource in the search field.

AWS CLI

Use `ListSupportPermits` to list support permits in your account and Region. You can filter results by status.

```
aws supportauthz list-support-permits \
  --support-permit-statuses '["ACTIVE"]'
```

The response includes up to 100 results per page. Use the `nextToken` value in subsequent requests to retrieve additional pages.

Revoking a support permit

Revoke a support permit to terminate active access for the selected resources. This action can't be undone. New authorizations stop being issued for the revoked support permit.

Console

1. Sign in to the [AWS Support Center Console](#).

2. In the navigation pane, choose **Support authorization**.
3. In the **Granted access** section, select the support permit that you want to revoke.
4. Choose **Revoke access**.

 Important

Revoking access terminates active access for the selected resources. This action can't be undone.

AWS CLI

Use `DeleteSupportPermit` to delete a support permit.

```
aws supportauthz delete-support-permit \  
  --arn "arn:aws:supportauthz:us-east-1:111122223333:supportpermit/AuroraDBPermit"
```

To prevent new signing operations, you can revoke the grant on the associated AWS KMS key. For more information, see [the section called “Revoking AWS Support authorization access”](#). Alternatively, you can disable or schedule deletion of the key, but both AWS KMS and AWS Support authorization are eventually consistent. After you disable a key, it can take up to several minutes for the change to take effect. During this period, signed authorizations might continue to be issued by using the key.

Best practices

Follow these recommendations when you create support permits:

- **Use time-based conditions:** Set `allowAfter` and `allowBefore` conditions to limit support permit validity to specific time windows.
- **Apply least-privilege scoping:** Specify only the actions and resources that AWS Support needs for the investigation. Avoid using `allActions` and `allResourcesInRegion` together unless necessary.
- **Link support permits to support cases:** Use the `supportCaseDisplayId` parameter to automatically deactivate the support permit when the support case closes.

- **Audit regularly:** Use `ListSupportPermits` to review active support permits and delete support permits that are no longer needed.

Managing support permit requests

In the reactive authorization scenario, AWS Support submits a support permit request when access to information about your services is needed for support case resolution. You review each request and decide whether to approve or reject it.

How reactive authorization works

The reactive authorization workflow follows these steps:

1. AWS Support identifies an issue with your resource that requires your permission to resolve.
2. AWS Support submits a support permit request, which is visible to you through AWS Support authorization. A `SupportPermitRequestCreated` event appears in your AWS CloudTrail logs.
3. You receive notifications through Amazon EventBridge events.
4. You review the request details, including the requested actions, resources, and support case context.
5. You approve the request by calling `CreateSupportPermit` with a support permit that covers the requested scope, or reject it by calling `RejectSupportPermitRequest`.

Viewing support permit requests

View pending permit requests from AWS Support.

Console

1. Sign in to the [AWS Support Center Console](#).
2. In the navigation pane, choose **Support authorization**.
3. In the **Pending access requests** section, review the list of requests. Each request shows the associated support case, service, Support Permit Request ARN, status, and request access date.
4. (Optional) Use the **Status** and **Service** dropdown filters to narrow the results, or search by resource in the search field.

AWS CLI

Use `ListSupportPermitRequests` to view permit requests. You can filter results by support case ID.

List all support permit requests

```
aws supportauthz list-support-permit-requests
```

List requests for a specific support case

```
aws supportauthz list-support-permit-requests \  
  --support-case-display-id "case-12345678"
```

The response includes up to 100 results per page. Use the `nextToken` value in subsequent requests to retrieve additional pages.

Approving a support permit request

To approve a support permit request, create a support permit that satisfies the requested scope.

Console

1. Sign in to the [AWS Support Center Console](#).
2. In the navigation pane, choose **Support authorization**.
3. In the **Pending access requests** section, choose **Manage support access**.
4. Review the **Request summary**, including the Support Permit Request ARN, service, status, and associated support case.
5. Review the **Access configuration**, which shows the name, description, requested actions, and access duration for the request.
6. In the **Resource access** section, choose **Approve** for each resource that you want to authorize. To approve all resources, choose **Approve all**.
7. For **Signing key**, choose an AWS KMS key from the dropdown list. To create a new key, choose **Create KMS signing key**.
8. Choose **Submit**.

AWS CLI

To approve a support permit request, create a support permit that satisfies the requested scope. Include the `supportCaseDisplayId` from the request so that the support permit automatically transitions to `INACTIVE` when the case closes.

```
aws supportauthz create-support-permit \
  --name "ApprovedRequest123" \
  --signing-key-info '{"kmsKey": "arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab"}' \
  --support-case-display-id "case-12345678" \
  --permit '{
    "actions": {"actions": ["rds:ReadClusterData"]},
    "resources": {"resources": ["arn:aws:rds:us-east-1:111122223333:cluster:my-cluster"]},
    "conditions": [
      {"allowBefore": "2026-06-15T00:00:00Z"}
    ]
  }'
```

When you create a support permit that covers the requested scope, the request status transitions to `ACCEPTED`.

Rejecting a support permit request

Reject a support permit request to decline an access request from AWS Support.

Console

1. Sign in to the [AWS Support Center Console](#).
2. In the navigation pane, choose **Support authorization**.
3. In the **Pending access requests** section, choose **Manage support access**.
4. In the **Resource access** section, choose **Deny** for each resource that you want to reject. To deny all resources, choose **Deny all**.
5. Choose **Submit**.

AWS CLI

Use `RejectSupportPermitRequest` to decline an access request. The request transitions to `REJECTED` status and AWS Support receives notification through the support case.

```
aws supportauthz reject-support-permit-request \
  --request-arn "arn:aws:supportauthz:us-east-1:111122223333:supportpermitrequest/
  request-id"
```

You can only reject requests that are in `PENDING` status.

Note

If you reject a support permit request or don't create a support permit that satisfies the request, AWS Support can't access the requested information and might not be able to proceed with the investigation or resolution of your support case.

Permit request state transitions

A support permit request can transition to the following states from `PENDING`:

- **ACCEPTED:** You created a support permit that satisfies the request.
- **REJECTED:** You rejected the request by calling `RejectSupportPermitRequest`.
- **CANCELLED:** AWS Support cancelled the request.

Discovering support actions

Support actions are service-specific operations that AWS Support can perform on your resources. Each action has a descriptive name and description that explains what the operation does. Actions use a namespaced format that identifies the AWS service and the specific operation. Use the `ListActions` API to view the actions available for a service before you create a scoped support permit.

Listing actions for a service

Use `ListActions` to view the actions available for a specific service. The `service` parameter is required.

```
aws supportauthz list-actions \  
  --service "rds"
```

Example output

```
{  
  "actionSummaries": [  
    {  
      "action": "rds:ReadClusterData",  
      "service": "rds",  
      "description": "Read diagnostic data from the database cluster."  
    },  
    {  
      "action": "rds:ViewQueryLogsWithParameters",  
      "service": "rds",  
      "description": "View query logs including query parameters."  
    }  
  ]  
}
```

The response includes up to 100 results per page. Use the `nextToken` value in subsequent requests to retrieve additional pages.

Getting action details

Use `GetAction` to retrieve details about a specific action.

```
aws supportauthz get-action \  
  --action "rds:ReadClusterData"
```

Example output

```
{  
  "action": "rds:ReadClusterData",  
  "service": "rds",  
  "description": "Read diagnostic data from the database cluster."  
}
```

Using actions in support permits

After you identify the available actions, use specific action names in your support permit to apply least-privilege scoping. Specify individual actions in the actions list rather than using `allActions` when you want to limit what AWS Support can do on your resources.

Create a support permit with specific actions

```
aws supportauthz create-support-permit \
  --name "ScopedPermit" \
  --signing-key-info '{"kmsKey": "arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab"}' \
  --permit '{
    "actions": {"actions": ["rds:ReadClusterData"]},
    "resources": {"resources": ["arn:aws:rds:us-east-1:111122223333:cluster:my-cluster"]},
    "conditions": [
      {"allowBefore": "2026-06-15T00:00:00Z"}
    ]
  }'
```

Monitoring AWS Support authorization with AWS CloudTrail

AWS CloudTrail logs all AWS Support authorization API calls and support operations as management events. These events are logged by default without additional configuration.

AWS service events

When AWS Support performs operations related to your support permits, AWS CloudTrail delivers the following events as `AwsServiceEvent` types. The event source for all events is `supportauthz.amazonaws.com`.

SupportPermitRequestCreated

AWS CloudTrail logs this event when AWS Support requests access to your resources. It indicates that a support permit request was created.

```
{
  "eventVersion": "1.11",
  "userIdentity": {
```



```

    "accountId": "111122223333",
    "invokedBy": "supportauthz.amazonaws.com"
  },
  "eventTime": "2026-06-16T20:04:51Z",
  "eventSource": "supportauthz.amazonaws.com",
  "eventName": "SupportPermitRequestCreated",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "supportauthz.amazonaws.com",
  "userAgent": "supportauthz.amazonaws.com",
  "requestParameters": null,
  "responseElements": null,
  "eventID": "a1b2c3d4-e5f6-7890-abcd-ef1234567890",
  "readOnly": false,
  "resources": [
    {
      "accountId": "111122223333",
      "type": "AWS::SupportAuthZ::SupportPermitRequest",
      "ARN": "arn:aws:supportauthz:us-east-1:111122223333:supportpermitrequest/
e5f6a7b8-c9d0-1234-efab-345678901234"
    }
  ],
  "eventType": "AwsServiceEvent",
  "managementEvent": true,
  "recipientAccountId": "111122223333",
  "serviceEventDetails": {
    "supportPermitRequestArn": "arn:aws:supportauthz:us-
east-1:111122223333:supportpermitrequest/e5f6a7b8-c9d0-1234-efab-345678901234",
    "supportCaseDisplayId": "1234567890",
    "resources": ["arn:aws:example:us-east-1:111122223333:resourcetype/resource-
id"],
    "actions": ["example:ExampleAction"]
  },
  "eventCategory": "Management"
}

```

RetrieveSupportPermit

AWS CloudTrail logs this event when AWS Support retrieves a signed authorization to perform a support action. This indicates that your AWS KMS key was used to sign the authorization.

```

{
  "eventVersion": "1.11",
  "userIdentity": {

```

```

    "accountId": "111122223333",
    "invokedBy": "supportauthz.amazonaws.com"
  },
  "eventTime": "2026-06-16T20:22:42Z",
  "eventSource": "supportauthz.amazonaws.com",
  "eventName": "RetrieveSupportPermit",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "supportauthz.amazonaws.com",
  "userAgent": "supportauthz.amazonaws.com",
  "requestParameters": null,
  "responseElements": null,
  "eventID": "b2c3d4e5-f6a7-8901-bcde-f12345678901",
  "readOnly": true,
  "eventType": "AwsServiceEvent",
  "managementEvent": true,
  "recipientAccountId": "111122223333",
  "serviceEventDetails": {
    "supportCaseDisplayId": "1234567890",
    "resource": "arn:aws:example:us-east-1:111122223333:resourcetype/resource-id",
    "action": "example:ExampleAction",
    "supportPermitArn": "arn:aws:supportauthz:us-east-1:111122223333:supportpermit/f6a7b8c9-d0e1-2345-fabc-456789012345"
  },
  "eventCategory": "Management"
}

```

StartSupportAction

AWS CloudTrail logs this event when AWS Support performs an authorized action on your resources.

```

{
  "eventVersion": "1.11",
  "userIdentity": {
    "accountId": "111122223333",
    "invokedBy": "supportauthz.amazonaws.com"
  },
  "eventTime": "2026-06-17T21:38:37Z",
  "eventSource": "supportauthz.amazonaws.com",
  "eventName": "StartSupportAction",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "supportauthz.amazonaws.com",
  "userAgent": "supportauthz.amazonaws.com",

```

```

    "requestParameters": null,
    "responseElements": null,
    "eventID": "c3d4e5f6-a7b8-9012-cdef-123456789012",
    "readOnly": false,
    "resources": [
      {
        "accountId": "111122223333",
        "type": "AWS::SupportAuthZ::SupportPermit",
        "ARN": "arn:aws:supportauthz:us-west-2:111122223333:supportpermit/example-permit"
      }
    ],
    "eventType": "AwsServiceEvent",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "serviceEventDetails": {
      "supportCaseDisplayId": "1234567890",
      "resourceIdentifiers": ["arn:aws:example:us-east-1:111122223333:resourcetype/resource-id"],
      "action": "example:ExampleAction",
      "supportPermitArn": "arn:aws:supportauthz:us-east-1:111122223333:supportpermit/f6a7b8c9-d0e1-2345-fabc-456789012345"
    },
    "eventCategory": "Management"
  }

```

Note

A single signed authorization can result in multiple `StartSupportAction` events. Each event represents one action performed on one resource. For example, if AWS Support reads diagnostic data from a cluster multiple times during a support case, each access generates a separate `StartSupportAction` event referencing the same `customerPermitArn`.

CancelSupportPermitRequest

AWS CloudTrail logs this event when AWS Support cancels a support permit request.

```

{
  "eventVersion": "1.11",
  "userIdentity": {
    "accountId": "111122223333",

```

```

    "invokedBy": "supportauthz.amazonaws.com"
  },
  "eventTime": "2026-06-16T20:04:42Z",
  "eventSource": "supportauthz.amazonaws.com",
  "eventName": "CancelSupportPermitRequest",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "supportauthz.amazonaws.com",
  "userAgent": "supportauthz.amazonaws.com",
  "requestParameters": null,
  "responseElements": null,
  "eventID": "d4e5f6a7-b8c9-0123-defa-234567890123",
  "readOnly": false,
  "resources": [
    {
      "accountId": "111122223333",
      "type": "AWS::SupportAuthZ::SupportPermitRequest",
      "ARN": "arn:aws:supportauthz:us-east-1:111122223333:supportpermitrequest/a7b8c9d0-e1f2-3456-abcd-567890123456"
    }
  ],
  "eventType": "AwsServiceEvent",
  "managementEvent": true,
  "recipientAccountId": "111122223333",
  "serviceEventDetails": {
    "supportPermitRequestArn": "arn:aws:supportauthz:us-east-1:111122223333:supportpermitrequest/a7b8c9d0-e1f2-3456-abcd-567890123456"
  },
  "eventCategory": "Management"
}

```

API events

AWS CloudTrail also generates standard management events for your own API calls to AWS Support authorization, including `CreateSupportPermit`, `DeleteSupportPermit`, `GetSupportPermit`, `ListSupportPermits`, `ListSupportPermitRequests`, and `RejectSupportPermitRequest`.

Setting up alerts

You can create Amazon CloudWatch alarms based on AWS CloudTrail events to alert you when AWS Support requests access or performs actions on your resources. To set up alerts:

1. Create an AWS CloudTrail trail that delivers events to Amazon CloudWatch Logs.

2. Create a metric filter for the AWS CloudTrail log group that matches on eventSource value supportauthz.amazonaws.com and the desired eventName (for example, SupportPermitRequestCreated or StartSupportAction).
3. Create a CloudWatch alarm on the metric filter to notify you through Amazon SNS.

Amazon EventBridge events

Events are sent to Amazon EventBridge when a support permit request is created. You can use these events to automate responses, such as automatically creating a support permit for low-risk actions or notifying your security team for review.

The following is an example event for a new support permit request:

```
{
  "version": "0",
  "id": "a1b2c3d4-e5f6-7890-abcd-ef1234567890",
  "detail-type": "Support Permit Request Created",
  "source": "aws.supportauthz",
  "account": "111122223333",
  "time": "2026-06-01T12:00:00Z",
  "region": "us-east-1",
  "detail": {
    "requestArn": "arn:aws:supportauthz:us-east-1:111122223333:supportpermitrequest/req-1234abcd",
    "supportCaseDisplayId": "case-12345678",
    "requestedActions": ["rds:ReadClusterData"],
    "requestedResources": ["arn:aws:rds:us-east-1:111122223333:cluster:my-aurora-cluster"],
    "status": "PENDING"
  }
}
```

To create an Amazon EventBridge rule that matches support permit request events, use the following event pattern:

```
{
  "source": ["aws.supportauthz"],
  "detail-type": ["Support Permit Request Created"]
}
```

You can route these events to targets such as Lambda functions, Amazon SNS topics, or Amazon SQS queues to build automated approval workflows.

AWS Support authorization quotas

The following tables list AWS Support authorization quotas.

Resource quotas

Resource	Quota	Adjustable
Support permits per account	200	No

Request rate quotas

Operation	Throttle rate	Throttle burst
CreateSupportPermit	3 tps	5 tps
DeleteSupportPermit	3 tps	5 tps
GetSupportPermit	3 tps	5 tps
ListSupportPermits	10 tps	15 tps
ListSupportPermitRequests	10 tps	15 tps
RejectSupportPermitRequest	3 tps	5 tps
GetAction	3 tps	5 tps
ListActions	10 tps	15 tps
ListTagsForResource	10 tps	15 tps
TagResource	3 tps	5 tps
UntagResource	3 tps	5 tps

Identity and access management for AWS Support

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 Support resources. IAM is an AWS service that you can use with no additional charge.

Topics

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)
- [How AWS Support works with IAM](#)
- [AWS Support identity-based policy examples](#)
- [Using service-linked roles](#)
- [AWS managed policies for AWS Support](#)
- [Manage access to AWS Support Center](#)
- [Manage access to AWS Support Plans](#)
- [Manage access to AWS Trusted Advisor](#)
- [Troubleshooting AWS Support identity and access](#)

Audience

How you use AWS Identity and Access Management (IAM) differs based on your role:

- **Service user** - request permissions from your administrator if you cannot access features (see [Troubleshooting AWS Support identity and access](#))
- **Service administrator** - determine user access and submit permission requests (see [How AWS Support works with IAM](#))
- **IAM administrator** - write policies to manage access (see [AWS Support identity-based policy examples](#))

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*.

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*.

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*.

How AWS Support works with IAM

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

For information about how to manage access for Support using IAM, see [Manage access for Support](#).

Topics

- [Support identity-based policies](#)
- [Support IAM roles](#)

Support identity-based policies

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. Support supports specific actions. To learn about the elements that you use in a JSON policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

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 Support use the following prefix before the action: `support:`. 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. Support 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": [  
    "ec2:action1",  
    "ec2:action2"
```

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

```
"Action": "ec2:Describe*"
```

To see a list of Support actions, see [Actions Defined by AWS Support](#) in the *IAM User Guide*.

Examples

To view examples of Support identity-based policies, see [AWS Support identity-based policy examples](#).

Support IAM roles

An [IAM role](#) is an entity within your AWS account that has specific permissions.

Using temporary credentials with Support

You can use temporary credentials to sign in with federation, assume an IAM role, or to assume a cross-account role. You obtain temporary security credentials by calling AWS STS API operations such as [AssumeRole](#) or [GetFederationToken](#).

Support supports using temporary credentials.

Service-linked roles

[Service-linked roles](#) allow AWS services to access resources in other services to complete an action on your behalf. Service-linked roles appear in your IAM account and are owned by the service. An IAM administrator can view but not edit the permissions for service-linked roles.

Support supports service-linked roles. For details about creating or managing Support service-linked roles, see [Using service-linked roles for AWS Support](#).

Service roles

This feature allows a service to assume a [service role](#) on your behalf. This role allows the service to access resources in other services to complete an action on your behalf. Service roles appear in your IAM account and are owned by the account. This means that an IAM administrator can change the permissions for this role. However, doing so might break the functionality of the service.

Support supports service roles.

AWS Support identity-based policy examples

By default, IAM users and roles don't have permission to create or modify Support resources. They also can't perform tasks using the AWS Management Console, AWS CLI, or AWS API. An IAM administrator must create IAM policies that grant users and roles permission to perform specific API operations on the specified resources they need. The administrator must then attach those policies to the IAM users or groups that require those permissions.

To learn how to create an IAM identity-based policy using these example JSON policy documents, see [Creating policies on the JSON tab](#) in the *IAM User Guide*.

Topics

- [Policy best practices](#)
- [Using the Support console](#)
- [Allow users to view their own permissions](#)

Policy best practices

Identity-based policies are very powerful. They determine whether someone can create, access, or delete Support resources in your account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get Started Using AWS Managed Policies** – To start using Support quickly, use AWS managed policies to give your employees the permissions they need. These policies are already available in your account and are maintained and updated by AWS. For more information, see [Get started using permissions with AWS managed policies](#) in the *IAM User Guide*.
- **Grant Least Privilege** – When you create custom policies, grant only the permissions required to perform a task. Start with a minimum set of permissions and grant additional permissions as necessary. Doing so is more secure than starting with permissions that are too lenient and then trying to tighten them later. For more information, see [Grant least privilege](#) in the *IAM User Guide*.
- **Enable MFA for Sensitive Operations** – For extra security, require IAM users to use multi-factor authentication (MFA) to access sensitive resources or API operations. For more information, see [Using multi-factor authentication \(MFA\) in AWS](#) in the *IAM User Guide*.
- **Use Policy Conditions for Extra Security** – To the extent that it's practical, define the conditions under which your identity-based policies allow access to a resource. For example, you can write conditions to specify a range of allowable IP addresses that a request must come from. You can also write conditions to allow requests only within a specified date or time range, or to require the use of SSL or MFA. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.

Using the Support console

To access the AWS Support console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Support resources in your AWS account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (IAM users or roles) with that policy.

To be sure that those entities can still use the Support console, also attach the following AWS managed policy to the entities. For more information, see [Adding permissions to a user](#) in the *IAM User Guide*:

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that you're trying to perform.

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewOwnUserInfo",
      "Effect": "Allow",
      "Action": [
        "iam:GetUserPolicy",
        "iam:ListGroupsForUser",
        "iam:ListAttachedUserPolicies",
        "iam:ListUserPolicies",
        "iam:GetUser"
      ],
      "Resource": ["arn:aws:iam::*:user/${aws:username}"]
    },
    {
      "Sid": "NavigateInConsole",
      "Effect": "Allow",
      "Action": [
        "iam:GetGroupPolicy",
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedGroupPolicies",
        "iam:ListGroupPolicies",
        "iam:ListPolicyVersions",
        "iam:ListPolicies",
        "iam:ListUsers"
      ],
      "Resource": "*"
    }
  ]
}
```

Using service-linked roles

AWS Support and AWS Trusted Advisor use AWS Identity and Access Management (IAM) [service-linked roles](#). A service-linked role is a unique IAM role that is linked directly to Support and Trusted Advisor. In each case, the service-linked role is a predefined role. This role includes all the permissions that Support or Trusted Advisor require to call other AWS services on your behalf. The following topics explain what service-linked roles do and how to work with them in Support and Trusted Advisor.

Topics

- [Using service-linked roles for AWS Support](#)
- [Using service-linked roles for AWS Support Plans](#)
- [Using service-linked roles for Trusted Advisor](#)

Using service-linked roles for AWS Support

AWS Support tools gather information about your AWS resources through API calls to provide customer service and technical support. To increase the transparency and auditability of support activities, Support uses an AWS Identity and Access Management (IAM) [service-linked role](#).

The `AWSServiceRoleForSupport` service-linked role is a unique IAM role that is linked directly to Support. This service-linked role is predefined, and it includes the permissions that Support requires to call other AWS services on your behalf.

The `AWSServiceRoleForSupport` service-linked role trusts the `support.amazonaws.com` service to assume the role.

To provide these services, the role's predefined permissions give Support access to resource metadata, not customer data. Only Support tools can assume this role, which exists within your AWS account.

We redact fields that could contain customer data. For example, the Input and Output fields of the [GetExecutionHistory](#) for the AWS Step Functions API call aren't visible to Support. We use AWS KMS keys to encrypt sensitive fields. These fields are redacted in the API response and aren't visible to AWS Support agents.

Note

AWS Trusted Advisor uses a separate IAM service-linked role to access AWS resources for your account to provide best practice recommendations and checks. For more information, see [Using service-linked roles for Trusted Advisor](#).

The `AWSServiceRoleForSupport` service-linked role enables all AWS Support API calls to be visible to customers through AWS CloudTrail. This helps with monitoring and auditing requirements, because it provides a transparent way to understand the actions that Support performs on your behalf. For information about CloudTrail, see the [AWS CloudTrail User Guide](#).

Service-linked role permissions for Support

This role uses the `AWSSupportServiceRolePolicy` AWS managed policy. This managed policy is attached to the role and allows the role permission to complete actions on your behalf.

These actions might include the following:

- **Billing, administrative, support, and other customer services** – AWS customer service uses the permissions granted by the managed policy to perform a number of services as part of your support plan. These include investigating and answering account and billing questions, providing administrative support for your account, increasing service quotas, and offering additional customer support.
- **Processing of service attributes and usage data for your AWS account** – Support might use the permissions granted by the managed policy to access service attributes and usage data for your AWS account. This policy allows Support to provide billing, administrative, and technical support for your account. Service attributes include your account's resource identifiers, metadata tags, roles, and permissions. Usage data includes usage policies, usage statistics, and analytics.
- **Maintaining the operational health of your account and its resources** – Support uses automated tools to perform actions related to operational and technical support.

For more information about the allowed services and actions, see the [AWSSupportServiceRolePolicy](#) policy in the IAM console.

 Note

AWS Support automatically updates the `AWSSupportServiceRolePolicy` policy once per month to add permissions for new AWS services and actions.

For more information, see [AWS managed policies for AWS Support](#).

Creating a service-linked role for Support

You don't need to manually create the `AWSServiceRoleForSupport` role. When you create an AWS account, this role is automatically created and configured for you.

 Important

If you used Support before it began supporting service-linked roles, then AWS created the `AWSServiceRoleForSupport` role in your account. For more information, see [A new role appeared in my IAM account](#).

Editing and deleting a service-linked role for Support

You can use IAM to edit the description for the `AWSServiceRoleForSupport` service-linked role. For more information, see [Editing a service-linked role](#) in the *IAM User Guide*.

The `AWSServiceRoleForSupport` role is necessary for Support to provide administrative, operational, and technical support for your account. As a result, this role can't be deleted through the IAM console, API, or AWS Command Line Interface (AWS CLI). This protects your AWS account, because you can't inadvertently remove necessary permissions for administering support services.

Customers onboarded to AWS Organizations and who have an Enterprise Support plan can delete the `AWSServiceRoleForSupport` service-linked role. Deleting this role restricts access to your resources by AWS Support engineers, limiting their ability to perform actions on your behalf. For more information, or to request to delete the `AWSServiceRoleForSupport` service-linked role, contact your Technical Account Manager (TAM).

For more information about the `AWSServiceRoleForSupport` role or its uses, contact [Support](#).

Using service-linked roles for AWS Support Plans

AWS Support Plans uses an AWS Identity and Access Management (IAM) [service-linked role](#) to read and update AWS resources used to manage your account's support plan on your behalf.

The `AWSServiceRoleForSupportPlans` service-linked role is a unique IAM role that is linked directly to Support Plans. This service-linked role is predefined, and it includes the permissions that Support Plans requires to call other AWS services on your behalf.

The `AWSServiceRoleForSupportPlans` service-linked role trusts the `supportplans.amazonaws.com` service to assume the role.

Service-linked role permissions for AWS Support Plans

This role uses the `AWSSupportPlansServiceRolePolicy` AWS managed policy. This managed policy is attached to the role and allows the role permission to complete actions on your behalf.

The role's permissions policy allows Support Plans to perform the following actions on your behalf:

- `organizations:ListAccounts`
- `supportplans:AcceptSupportAgreement`
- `supportplans:CancelSupportAgreement`
- `supportplans:CreateSupportAgreement`
- `supportplans:GetSupportAgreement`
- `supportplans:ListSupportAgreementRevisions`
- `supportplans:ListSupportAgreements`
- `supportplans:RejectSupportAgreement`
- `supportplans:StartSupportPlanUpdate`
- `supportplans:UpdateSupportAgreement`

For more information about the allowed actions, see the [AWSSupportPlansServiceRolePolicy](#) policy in the IAM console.

Creating a service-linked role for AWS Support Plans

You don't need to manually create the `AWSServiceRoleForSupportPlans` role. When you create a support plan through the Support Plan Management page in the AWS Management Console, Support Plans automatically creates this role for you.

Note

If you need to create this role manually, you can use the IAM console, AWS CLI, or the IAM API. For more information, see [Creating a service-linked role](#) in the *IAM User Guide*.

Editing a service-linked role for AWS Support Plans

You can use IAM to edit the description for the `AWSServiceRoleForSupportPlans` service-linked role. You can't edit the trust policy or permissions policy, and you can't attach additional policies to this role. For more information, see [Editing a service-linked role](#) in the *IAM User Guide*.

Deleting a service-linked role for AWS Support Plans

Support Plans doesn't delete the `AWSServiceRoleForSupportPlans` service-linked role on your behalf. If you want to delete this role, you must first cancel your active support plan. To do this, contact Support. After Support confirms the cancellation, you can use the IAM console, AWS CLI, or the IAM API to delete the role. For more information, see [Deleting a service-linked role](#) in the *IAM User Guide*.

Using service-linked roles for Trusted Advisor

AWS Trusted Advisor uses the AWS Identity and Access Management (IAM) [service-linked role](#). A service-linked role is a unique IAM role that is linked directly to AWS Trusted Advisor. Service-linked roles are predefined by Trusted Advisor, and they include all the permissions that the service requires to call other AWS services on your behalf. Trusted Advisor uses this role to check your usage across AWS and to provide recommendations to improve your AWS environment. For example, Trusted Advisor analyzes your Amazon Elastic Compute Cloud (Amazon EC2) instance use to help you reduce costs, increase performance, tolerate failures, and improve security.

Note

AWS Support uses a separate IAM service-linked role for accessing your account's resources to provide billing, administrative, and support services. For more information, see [Using service-linked roles for AWS Support](#).

For information about other services that support service-linked roles, see [AWS services that work with IAM](#). Look for the services that have **Yes** in the **Service-linked role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Topics

- [Service-linked role permissions for Trusted Advisor](#)
- [Manage permissions for service-linked roles](#)
- [Creating a service-linked role for Trusted Advisor](#)
- [Editing a service-linked role for Trusted Advisor](#)
- [Deleting a service-linked role for Trusted Advisor](#)

Service-linked role permissions for Trusted Advisor

Trusted Advisor uses two service-linked roles:

- [AWSServiceRoleForTrustedAdvisor](#) – This role trusts the Trusted Advisor service to assume the role to access AWS services on your behalf. The role permissions policy allows Trusted Advisor read-only access for all AWS resources. This role simplifies getting started with your AWS account, because you don't have to add the necessary permissions for Trusted Advisor. When you open an AWS account, Trusted Advisor creates this role for you. The defined permissions include the trust policy and the permissions policy. You can't attach the permissions policy to any other IAM entity.

For more information about the attached policy, see [AWSTrustedAdvisorServiceRolePolicy](#).

- [AWSServiceRoleForTrustedAdvisorReporting](#) – This role trusts the Trusted Advisor service to assume the role for the organizational view feature. This role enables Trusted Advisor as a trusted service in your AWS Organizations organization. Trusted Advisor creates this role for you when you enable organizational view.

For more information about the attached policy, see [AWSTrustedAdvisorReportingServiceRolePolicy](#).

You can use the organizational view to create reports for Trusted Advisor check results for all accounts in your organization. For more information about this feature, see [Organizational view for AWS Trusted Advisor](#).

Manage permissions for service-linked roles

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. The following examples use the `AWSServiceRoleForTrustedAdvisor` service-linked role.

Example: Allow an IAM entity to create the `AWSServiceRoleForTrustedAdvisor` service-linked role

This step is necessary only if the Trusted Advisor account is disabled, the service-linked role is deleted, and the user must recreate the role to reenable Trusted Advisor.

You can add the following statement to the permissions policy for the IAM entity to create the service-linked role.

```
{
  "Effect": "Allow",
  "Action": [
    "iam:CreateServiceLinkedRole",
    "iam:PutRolePolicy"
  ],
  "Resource": "arn:aws:iam::*:role/aws-service-role/trustedadvisor.amazonaws.com/AWSServiceRoleForTrustedAdvisor*",
  "Condition": {"StringLike": {"iam:AWSServiceName": "trustedadvisor.amazonaws.com"}}
}
```

Example: Allow an IAM entity to edit the description of the `AWSServiceRoleForTrustedAdvisor` service-linked role

You can only edit the description for the `AWSServiceRoleForTrustedAdvisor` role. You can add the following statement to the permissions policy for the IAM entity to edit the description of a service-linked role.

```
{
  "Effect": "Allow",
  "Action": [
    "iam:UpdateRoleDescription"
  ],
  "Resource": "arn:aws:iam::*:role/aws-service-role/trustedadvisor.amazonaws.com/AWSServiceRoleForTrustedAdvisor*",
  "Condition": {"StringLike": {"iam:AWSServiceName": "trustedadvisor.amazonaws.com"}}
}
```

```
}
```

Example: Allow an IAM entity to delete the `AWSServiceRoleForTrustedAdvisor` service-linked role

You can add the following statement to the permissions policy for the IAM entity to delete a service-linked role.

```
{
  "Effect": "Allow",
  "Action": [
    "iam:DeleteServiceLinkedRole",
    "iam:GetServiceLinkedRoleDeletionStatus"
  ],
  "Resource": "arn:aws:iam::*:role/aws-service-role/trustedadvisor.amazonaws.com/AWSServiceRoleForTrustedAdvisor*",
  "Condition": {"StringLike": {"iam:AWSServiceName": "trustedadvisor.amazonaws.com"}}
}
```

You can also use an AWS managed policy, such as [AdministratorAccess](#), to provide full access to Trusted Advisor.

Creating a service-linked role for Trusted Advisor

You don't need to manually create the `AWSServiceRoleForTrustedAdvisor` service-linked role. When you open an AWS account, Trusted Advisor creates the service-linked role for you.

Important

If you were using the Trusted Advisor service before it began supporting service-linked roles, then Trusted Advisor already created the `AWSServiceRoleForTrustedAdvisor` role in your account. To learn more, see [A new role appeared in my IAM account](#) in the *IAM User Guide*.

If your account doesn't have the `AWSServiceRoleForTrustedAdvisor` service-linked role, then Trusted Advisor won't work as expected. This can happen if someone in your account disabled Trusted Advisor and then deleted the service-linked role. In this case, you can use IAM to create the `AWSServiceRoleForTrustedAdvisor` service-linked role, and then reenable Trusted Advisor.

To enable Trusted Advisor (console)

1. Use the IAM console, AWS CLI, or the IAM API to create a service-linked role for Trusted Advisor. For more information, see [Creating a service-linked role](#).
2. Sign in to the AWS Management Console, and then navigate to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor>.

The **Disabled Trusted Advisor** status banner appears in the console.

3. Choose **Enable Trusted Advisor Role** from the status banner. If the required `AWSServiceRoleForTrustedAdvisor` isn't detected, the disabled status banner remains.

Editing a service-linked role for Trusted Advisor

You can't change the name of a service-linked role because various entities might reference the role. However, you can use the IAM console, AWS CLI, or the IAM API to edit the description of the role. For more information, see [Editing a service-linked role](#) in the *IAM User Guide*.

Deleting a service-linked role for Trusted Advisor

If you don't need to use the features or services of Trusted Advisor, you can delete the `AWSServiceRoleForTrustedAdvisor` role. You must disable Trusted Advisor before you can delete this service-linked role. This prevents you from removing permissions required by Trusted Advisor operations. When you disable Trusted Advisor, you disable all service features, including offline processing and notifications. Also, if you disable Trusted Advisor for a member account, then the separate payer account is also affected, which means you won't receive Trusted Advisor checks that identify ways to save costs. You can't access the Trusted Advisor console. API calls to Trusted Advisor return an access denied error.

You must recreate the `AWSServiceRoleForTrustedAdvisor` service-linked role in the account before you can reenable Trusted Advisor.

You must first disable Trusted Advisor in the console before you can delete the `AWSServiceRoleForTrustedAdvisor` service-linked role.

To disable Trusted Advisor

1. Sign in to the AWS Management Console and navigate to the Trusted Advisor console at <https://console.aws.amazon.com/trustedadvisor>.
2. In the navigation pane, choose **Preferences**.

3. In the **Service Linked Role Permissions** section, choose **Disable Trusted Advisor**.
4. In the confirmation dialog box, choose **OK** to confirm that you want to disable Trusted Advisor.

After you disable Trusted Advisor, all Trusted Advisor functionality is disabled, and the Trusted Advisor console displays only the disabled status banner.

You can then use the IAM console, the AWS CLI, or the IAM API to delete the Trusted Advisor service-linked role named `AWSServiceRoleForTrustedAdvisor`. For more information, see [Deleting a service-linked role](#) in the *IAM User Guide*.

AWS managed policies for AWS Support

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS customers to use. We recommend that you reduce permissions further by defining [customer managed policies](#) that are specific to your use cases.

You cannot change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see [AWS managed policies](#) in the *IAM User Guide*.

Topics

- [AWS managed policies for AWS Support](#)
- [AWS managed policies for AWS Support App in Slack](#)
- [AWS managed policies for AWS Trusted Advisor](#)
- [AWS managed policies for AWS Support Plans](#)
- [AWS managed policies for AWS Partner-Led Support](#)

AWS managed policies for AWS Support

AWS Support has the following managed policies.

Contents

- [AWS managed policy: AWSSupportAccess](#)
- [AWS managed policy: AWSSupportServiceRolePolicy](#)
- [AWS Support updates to AWS managed policies](#)
- [Permission changes for AWSSupportServiceRolePolicy](#)

AWS managed policy: AWSSupportAccess

AWS Support uses the [AWSSupportAccess](#) AWS managed policy. This policy manages your support case lifecycle through the Support API. Enhanced functionality in the AWS Support Center Console is provided through the support-console API service. You can attach this policy to your IAM entities. For more information, see [Service-linked role permissions for Support](#).

To view the permissions for this policy, see [AWSSupportAccess](#) in the *AWS Managed Policy Reference*.

AWS managed policy: AWSSupportServiceRolePolicy

AWS Support uses the [AWSSupportServiceRolePolicy](#) AWS managed policy. This managed policy is attached to the `AWSServiceRoleForSupport` service-linked role. The policy allows the service-linked role to complete actions on your behalf. You can't attach this policy to your IAM entities. For more information, see [Service-linked role permissions for Support](#).

To view the permissions for this policy, see [AWSSupportServiceRolePolicy](#) in the *AWS Managed Policy Reference*.

For a list of changes to the policy, see [AWS Support updates to AWS managed policies](#) and [Permission changes for AWSSupportServiceRolePolicy](#).

AWS Support updates to AWS managed policies

View details about updates to AWS managed policies for AWS Support since these services began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the [Document history](#) page.

The following table describes important updates to the AWS Support managed policies since February 17, 2022.

AWS Support

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Amazon API Gateway – To debug issues related to Amazon API Gateway. • AWS AppConfig – To debug issues related to AWS AppConfig. • AWS Backup – To debug issues related to AWS Backup. • Amazon Bedrock – To debug issues related to Amazon Bedrock. • Amazon Bedrock AgentCore – To debug issues related to Amazon Bedrock AgentCore. • AWS Certificate Manager – To debug issues related to AWS Certificate Manager.	August 19, 2026

Change	Description	Date
	• Amazon Q Developer in chat applications – To debug issues related to Amazon Q Developer in chat applications. • AWS CloudTrail – To debug issues related to . • Amazon CloudWatch – To debug issues related to Amazon CloudWatch. • Amazon Connect – To debug issues related to Amazon Connect. • AWS Database Migration Service – To debug issues related to AWS Database Migration Service. • Amazon DataZone – To debug issues related to Amazon DataZone. • Amazon EC2 – To debug issues related to Amazon EC2. • Amazon EKS – To debug issues related to Amazon EKS. • Elastic Load Balancing – To debug issues related to Elastic Load Balancing. • AWS Lambda – To debug issues related to AWS Lambda.	

Change	Description	Date
	• Amazon Managed Streaming for Apache Kafka – To debug issues related to Amazon Managed Streaming for Apache Kafka. • Amazon OpenSearch Service – To debug issues related to Amazon OpenSearch Service. • Oracle Database@AWS – To debug issues related to Oracle Database@AWS. • Amazon Quick – To debug issues related to Amazon Quick. • Amazon Redshift Serverless – To debug issues related to Amazon Redshift Serverless. • Amazon S3 Files – To debug issues related to Amazon S3 Files. • Savings Plans – To debug issues related to Savings Plans. • AWS Security Agent – To debug issues related to AWS Security Agent. • AWS Shield – To debug issues related to AWS Shield.	

Change	Description	Date
	AWS Sustainability – To debug issues related to AWS Sustainability.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 217 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • AWS DevOps Agent – To debug issues related to AWS DevOps Agent. • Amazon API Gateway – To debug issues related to Amazon API Gateway. • AWS Batch – To debug issues related to AWS Batch. • Amazon Bedrock – To debug issues related to Amazon Bedrock. • AWS Certificate Manager – To debug issues related to AWS Certificate Manager. • AWS CloudFormation – To debug issues related to AWS CloudFormation. • Amazon CloudFront – To debug issues related to Amazon CloudFront. • AWS CloudTrail – To debug issues related to AWS CloudTrail. • Amazon Cognito Identity Pools – To debug issues	June 22, 2026

Change	Description	Date
	related to Amazon Cognito Identity Pools. • AWS Control Catalog – To debug issues related to AWS Control Catalog. • AWS Elastic Beanstalk – To debug issues related to AWS Elastic Beanstalk. • Amazon EMR Serverless – To debug issues related to Amazon EMR Serverless. • AWS Entity Resolution – To debug issues related to AWS Entity Resolution. • Amazon EventBridge Streaming for Apache Kafka – To debug issues related to Amazon EventBridge Streaming for Apache Kafka. • AWS Identity Store – To debug issues related to AWS IAM Identity Center. • Amazon Kendra – To debug issues related to Amazon Kendra. • AWS Key Management Service – To debug issues related to AWS Key Management Service. • Amazon Managed Streaming for Apache Kafka – To debug issues	

Change	Description	Date
	related to Amazon Managed Streaming for Apache Kafka. • AWS Marketplace – To debug issues related to AWS Marketplace. • Amazon OpenSearch Service – To debug issues related to Amazon OpenSearch Service. • Amazon OpenSearch Serverless – To debug issues related to Amazon OpenSearch Serverless. • Amazon Personalize – To debug issues related to Amazon Personalize. • AWS Private Certificate Authority – To debug issues related to AWS Private Certificate Authority. • Amazon S3 Tables – To debug issues related to Amazon S3 Tables. • Amazon SageMaker – To debug issues related to Amazon SageMaker. • AWS Security Incident Response – To debug issues related to AWS Security Incident Response.	

Change	Description	Date
	• AWS Sustainability – To debug issues related to AWS Sustainability. • Amazon VPC Lattice – To debug issues related to Amazon VPC Lattice. • AWS WAF – To debug issues related to AWS WAF.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 146 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Amazon Managed Workflows for Apache Airflow Serverless – To debug issues related to Amazon Managed Workflows for Apache Airflow Serverless. • – To debug issues related to . • Amazon Bedrock Agent Core – To debug issues related to Amazon Bedrock Agent Core. • Campaigns – To debug issues related to Campaigns . • AWS Control Tower – To debug issues related to AWS Control Tower. • AWS Device Farm – To debug issues related to AWS Device Farm. • Amazon EC2 – To debug issues related to Amazon EC2.	April 21, 2026

Change	Description	Date
	• Amazon Elastic Container Service – To debug issues related to Amazon Elastic Container Service. • Amazon EKS – To debug issues related to Amazon EKS. • AWS Lambda – To debug issues related to AWS Lambda. • Amazon Macie – To debug issues related to Amazon Macie. • Amazon Route 53 Recovery Control Config – To debug issues related to Amazon Route 53 Recovery Control Config. • Amazon Route 53 Recovery Readiness – To debug issues related to Amazon Route 53 Recovery Readiness. • Amazon Route 53 Resolver – To debug issues related to Amazon Route 53 Resolver. • Amazon Route 53 Global Resolver – To debug issues related to Amazon Route 53 Global Resolver. • Amazon S3 on Outposts – To debug issues related to Amazon S3 on Outposts.	

Change	Description	Date
	• AWS Security Hub CSPM – To debug issues related to AWS Security Hub CSPM. • Amazon Simple Email Service – To debug issues related to Amazon Simple Email Service. • Amazon Pinpoint SMS Voice – To debug issues related to Amazon Pinpoint SMS Voice. • Amazon EC2 Systems Manager for SAP – To debug issues related to Amazon EC2 Systems Manager for SAP. • Amazon WorkSpaces – To debug issues related to Amazon WorkSpaces.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 105 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Amazon Bedrock – To troubleshoot issues related to Amazon Bedrock. • Amazon Braket – To debug issues related to Amazon Braket. • Amazon Connect – To debug issues related to Amazon Connect. • Amazon Aurora DSQL – To troubleshoot issues related to Amazon Aurora DSQL. • AWS Compute Optimizer – To debug issues related to AWS Compute Optimizer. • Amazon Connect – To debug issues related to Amazon Connect. • Amazon EBS – To debug issues related to Amazon EBS. • Amazon EC2 – To troubleshoot issues related to Amazon EC2.	January 29, 2026

Change	Description	Date
	• Amazon EKS – To debug issues related to Amazon EKS. • AWS Elastic Beanstalk – To debug issues related to AWS Elastic Beanstalk. • AWS Glue – To troubleshoot issues related to AWS Glue. • AWS IoT Greengrass – To debug issues related to AWS IoT Greengrass. • AWS Invoicing – To debug issues related to AWS Invoicing. • Amazon Kinesis Video Streams – To debug issues related to Amazon Kinesis Video Streams. • AWS License Manager – To debug issues related to AWS License Manager. • Amazon MemoryDB – To troubleshoot issues related to Amazon MemoryDB. • Multi-party approval – To debug issues related to Multi-party approval. • Amazon MQ – To troubleshoot issues related to Amazon MQ. • Oracle Database@AWS – To troubleshoot issues related to Oracle Database@AWS.	

Change	Description	Date
	• AWS Payment Cryptography – To debug issues related to AWS Payment Cryptography. • Amazon RDS – To debug issues related to Amazon RDS. • AWS Tax Settings – To troubleshoot issues related to AWS Tax Settings.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 145 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • AWS Amplify – To troubleshoot issues related to AWS Amplify. • AWS AppSync – To debug issues related to AWS AppSync. • AWS Outposts – To debug issues related to AWS Outposts. • AWS Clean Rooms – To troubleshoot issues related to AWS Clean Rooms. • AWS Compute Optimizer – To debug issues related to AWS Compute Optimizer. • Amazon Connect – To debug issues related to Amazon Connect. • Amazon DynamoDB – To debug issues related to Amazon DynamoDB. • Amazon EMR – To troubleshoot issues related to Amazon EMR. • Amazon Location Service – To troubleshoot issues	December 08, 2025

Change	Description	Date
	related to Amazon Location Service. • Amazon GuardDuty – To debug issues related to Amazon GuardDuty. • AWS Network Firewall – To debug issues related to AWS Network Firewall. • AWS HealthOmics – To troubleshoot issues related to AWS HealthOmics. • AWS Organizations – To debug issues related to AWS Organizations. • Amazon S3 – To debug issues related to Amazon S3. • Amazon S3 Tables – To debug issues related to Amazon S3 Tables. • Amazon S3 Vectors – To debug issues related to Amazon S3 Vectors. • Amazon SageMaker AI – To troubleshoot issues related to Amazon SageMaker AI. • AWS Security Hub CSPM – To troubleshoot issues related to AWS Security Hub CSPM. • Amazon SES – To debug issues related to Amazon SES.	

Change	Description	Date
	• AWS Signer – To debug issues related to AWS Signer. • AWS STS – To troubleshoot issues related to AWS STS.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 125 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Artificial intelligence for IT operations(AIOps) – To debug issues related to Artificial intelligence for IT operations(AIOps). • AWS Backup gateway – To troubleshoot issues related to AWS Backup gateway. • CloudFormation – To debug issues related to CloudFormation. • Amazon Cognito Identity – To debug issues related to Amazon Cognito Identity. • Amazon Cognito – To troubleshoot issues related to Amazon Cognito. • AWS Backup – To debug issues related to AWS Backup. • Directory Service – To debug issues related to Directory Service. • Amazon EC2 – To debug issues related to Amazon EC2.	Sep 30, 2025

Change	Description	Date
	• AWS FIS – To troubleshoot issues related to AWS FIS. • Amazon FSx – To troubleshoot issues related to Amazon FSx. • AWS Global Accelerator – To debug issues related to AWS Global Accelerator. • AWS Identity Store – To debug issues related to AWS Identity Store. • AWS Invoicing – To debug issues related to AWS Invoicing. • AWS Lake Formation – To troubleshoot issues related to AWS Lake Formation. • AWS Network Firewall – To debug issues related to AWS Network Firewall. • Oracle Database@AWS – To troubleshoot issues related to Oracle Database@AWS. • Amazon S3 – To debug issues related to Amazon S3. • Amazon SES – To troubleshoot issues related to Amazon SES. • AWS IAM Identity Center – To troubleshoot issues	

Change	Description	Date
	related to AWS IAM Identity Center.	
AWSSupportAccess – Update to an existing policy	Added permissions for the support-console API to the AWSSupportAccess managed policy.	July 18, 2025

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 25 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • CloudFormation – To debug issues related to CloudFormation. • AWS Config – To troubleshoot issues related to AWS Config. • Amazon OpenSearch Service – To debug issues related to Amazon OpenSearch Service. • AWS Glue – To debug issues related to AWS Glue. • AWS IAM – To troubleshoot issues related to AWS IAM. • Amazon Pinpoint – To troubleshoot issues related to Amazon Pinpoint. • AWS Outposts – To debug issues related to AWS Outposts. • AWS STS – To debug issues related to AWS STS. For more information, see Permission changes for	Jul 15, 2025

Change	Description	Date
	AWSSupportServiceRolePolicy :	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 257 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • AWS App Runner – To debug issues related to AWS App Runner. • AWS AppSync – To troubleshoot issues related to AWS AppSync. • AWS Batch – To debug issues related to AWS Batch. • Amazon Bedrock – To troubleshoot issues related to Amazon Bedrock. • Amazon CloudFront – To debug issues related to Amazon CloudFront. • AWS CodePipeline – To troubleshoot issues related to AWS CodePipeline. • AWS Config – To troubleshoot issues related to AWS Config. • Amazon Connect – To debug issues related to Amazon Connect.	Jun 17, 2025

Change	Description	Date
	• AWS DataSync – To debug issues related to AWS DataSync. • Direct Connect – To troubleshoot issues related to the Direct Connect. • Amazon EC2 – To troubleshoot issues related to the Amazon EC2. • AWS Fault Injection Service – To debug issues related to the AWS Fault Injection Service. • AWS Firewall Manager – To troubleshoot issues related to the AWS Firewall Manager. • AWS Glue – To debug issues related to the AWS Glue. • Amazon GuardDuty – To debug issues related to the Amazon GuardDuty. • EC2 Image Builder – To troubleshoot issues related to the EC2 Image Builder. • AWS IoT – To troubleshoot issues related to the AWS IoT. • AWS IoT FleetWise – To debug issues related to the AWS IoT FleetWise. • Amazon CloudWatch Logs – To debug issues related	

Change	Description	Date
	to the Amazon CloudWatch Logs. • AWS Elemental MediaLive – To debug issues related to the AWS Elemental MediaLive. • Network Flow Monitor – To troubleshoot issues related to the Network Flow Monitor. • AWS Network Manager – To troubleshoot issues related to AWS Network Manager. • Amazon CloudWatch Observability Admin Service – To troubleshoot issues related to the Amazon CloudWatch Observability Admin Service. • AWS Parallel Computing Service – To debug issues related to the AWS Parallel Computing Service. • Amazon Redshift Serverless – To troubleshoot issues related to the Amazon Redshift Serverless. • Amazon Redshift – To troubleshoot issues related to the Amazon Redshift.	

Change	Description	Date
	• AWS Resilience Hub – To debug issues related to the AWS Resilience Hub. • AWS Identity and Access Management Roles Anywhere– To debug issues related to the AWS Identity and Access Management Roles Anywhere. • Amazon S3 on Outposts – To troubleshoot issues related to the Amazon S3 on Outposts. • Amazon S3 – To troubleshoot issues related to the Amazon S3. • Amazon S3 Tables– To troubleshoot issues related to the Amazon S3 Tables. • Amazon SageMaker AI – To debug issues related to the Amazon SageMaker AI. • AWS Security Hub CSPM – To debug issues related to the AWS Security Hub CSPM. • Amazon SQS – To debug issues related to the Amazon SQS. • AWS Systems Manager Incident Manager – To troubleshoot issues related	

Change	Description	Date
	to the AWS Systems Manager Incident Manager. • AWS Systems Manager Quick Setup – To debug issues related to the AWS Systems Manager Quick Setup. • AWS Systems Manager – To debug issues related to the AWS Systems Manager. • Amazon WorkSpaces Thin Client – To troubleshoot issues related to the Amazon WorkSpaces Thin Client. • Amazon Timestream – To debug issues related to the Amazon Timestream. • AWS Telco Network Builder – To troubleshoot issues related to the AWS Telco Network Builder. • AWS Transfer Family – To debug issues related to the AWS Transfer Family. • Amazon VPC Lattice – To troubleshoot issues related to the Amazon VPC Lattice.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 88 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Amazon Bedrock – To troubleshoot issues related to Amazon Bedrock. • Amazon Connect – To debug issues related to Amazon Connect. • Amazon DataZone – To debug issues related to Amazon DataZone. • Amazon EC2 – To troubleshoot issues related to the Amazon EC2. • Amazon EKS – To debug issues related to the Amazon EKS. • AWS Glue – To troubleshoot issues related to AWS Glue. • Amazon Managed Service for Apache Flink – To troubleshoot issues related to the Amazon Managed Service for Apache Flink. • AWS Lambda – To debug issues related to the AWS Lambda.	Nov 25, 2024

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 79 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Amazon OpenSearch Serverless – To troubleshoot issues related to Amazon OpenSearch Serverless. • AWS AppConfig – To debug issues related to AWS AppConfig. • Application Signals– To debug issues related to Application Signals. • Amazon Athena – To troubleshoot issues related to the Amazon Athena. • Amazon CloudWatch – To debug issues related to the Amazon CloudWatch. • Amazon DynamoDB – To troubleshoot issues related to Amazon DynamoDB. • Amazon EC2 – To troubleshoot issues related to the Amazon EC2. • AWS IoT – To debug issues related to the AWS IoT.	Oct 8, 2024

Change	Description	Date
	• AWS Lambda – To troubleshoot issues related to the AWS Lambda. • AWS Launch Wizard – To troubleshoot issues related to the AWS Launch Wizard. • AWS Security Hub CSPM – To debug issues related to AWS Security Hub CSPM. • Amazon WorkSpaces – To debug issues related to the Amazon WorkSpaces.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 79 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • AWS account – To troubleshoot issues related to the AWS account. • AWS Auto Scaling – To debug issues related to AWS Auto Scaling. • Amazon Bedrock – To debug issues related to Amazon Bedrock. • AWS CodeConnections – To troubleshoot issues related to the AWS CodeConnections. • AWS Deadline Cloud – To debug issues related to the AWS Deadline Cloud. • Amazon Elastic Kubernetes Service – To troubleshoot issues related to Amazon Elastic Kubernetes Service. • Elastic Load Balancing – To troubleshoot issues related to the Elastic Load Balancing.	Aug 5, 2024

Change	Description	Date
	• AWS Free Tier – To debug issues related to the AWS Free Tier. • Amazon Inspector – To troubleshoot issues related to the Amazon Inspector. • Amazon OpenSearch Ingestion – To troubleshoot issues related to the Amazon OpenSearch Ingestion. • Amazon WorkSpaces – To debug issues related to Amazon WorkSpaces. • AWS X-Ray – To debug issues related to the AWS X-Ray.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 17 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Amazon CloudWatch Network Monitor – To troubleshoot issues related to the Network Monitor service. • Amazon CloudWatch Logs – To debug issues related to Amazon CloudWatch Logs. • Amazon Managed Streaming for Apache Kafka – To debug issues related to Amazon Managed Streaming for Apache Kafka. • Amazon Managed Service for Prometheus – To troubleshoot issues related to the Amazon Managed Service for Prometheus.	Mar 22, 2024

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 63 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • AWS Clean Rooms – To troubleshoot issues related to the AWS Clean Rooms. • CodeConnections – To troubleshoot issues related to CodeConnections. • Amazon EKS – To debug issues related to Amazon EKS. • Image Builder – To debug issues related to the Image Builder. • Amazon Inspector2 – To troubleshoot issues related to Amazon Inspector2. • Amazon Inspector Scan – To debug issues related to the Amazon Inspector Scan. • Amazon CloudWatch Logs – To troubleshoot issues related to Amazon CloudWatch Logs. • AWS Outposts – To troubleshoot issues related to the AWS Outposts.	Jan 17, 2024

Change	Description	Date
	• Amazon RDS – To debug issues related to Amazon RDS. • AWS IAM Identity Center – To troubleshoot issues related to AWS IAM Identity Center. • Amazon S3 Express – To debug issues related to Amazon S3 Express. • AWS Trusted Advisor – To troubleshoot issues related to AWS Trusted Advisor.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 126 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • AWS Direct Connect – To troubleshoot issues related to the AWS Direct Connect service. • Amazon SageMaker AI – To troubleshoot issues related to Amazon SageMaker AI service. • Amazon AppStream – To debug issues related to Amazon AppStream. • AWS Resource Explorer – To debug issues related to the AWS Resource Explorer. • Amazon Redshift serverless – To troubleshoot issues related to Amazon Redshift serverless. • Amazon ElastiCache – To debug issues related to the Amazon ElastiCache. • Amazon Comprehend – To troubleshoot issues related to Amazon Comprehend.	Dec 6, 2023

Change	Description	Date
	• Amazon EC2 – To troubleshoot issues related to the Amazon EC2. • Amazon Elastic Kubernetes Service – To debug issues related to Amazon Elastic Kubernetes Service. • AWS Elastic Disaster Recovery – To troubleshoot issues related to AWS Elastic Disaster Recovery. • AWS AppSync – To debug issues related to AWS AppSync. • Amazon CloudWatch Logs – To troubleshoot issues related to Amazon CloudWatch Logs. • AWS Health – To debug issues related to the AWS Health Service. • Connect Customer – To debug issues related to the Connect Customer. • AWS Snowball Edge – To troubleshoot issues related to AWS Snowball Edge. • AWS HealthImaging – To troubleshoot issues related to AWS HealthImaging.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 163 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Amazon CloudFront – To troubleshoot issues related to the CloudFront service. • Amazon EC2 – To troubleshoot issues related to Amazon EC2 service. • Amazon AppStream – To debug issues related to Amazon AppStream. • AWS WAF – To debug issues related to the AWS Web Application Firewall. • Connect Customer – To troubleshoot issues related to Connect Customer. • AWS IoT – To debug issues related to the AWS IoT. • Amazon Route 53 – To troubleshoot issues related to Amazon Route 53. • AWS Verified Access – To troubleshoot issues related to the AWS Verified Access service. • Amazon Simple Email Service – To debug issues	Oct 27, 2023

Change	Description	Date
	related to Amazon Simple Email Service. • AWS Elastic Beanstalk – To troubleshoot issues related to AWS Elastic Beanstalk. • Amazon DynamoDB – To debug issues related to Amazon DynamoDB. • AWS EC2 Image Builder – To troubleshoot issues related to AWS EC2 Image Builder. • AWS Outposts – To debug issues related to the AWS Outposts Service. • AWS Glue – To debug issues related to the AWS Glue. • AWS Directory Service – To troubleshoot issues related to AWS Directory Service. • AWS Elastic Disaster Recovery – To troubleshoot issues related to AWS Elastic Disaster Recovery. • AWS Step Functions – To debug issues related to AWS Step Functions. • Amazon EMR – To troubleshoot issues related to Amazon EMR. • Amazon Relational Database Service – To troubleshoot issues related	

Change	Description	Date
	to Amazon Relational Database Service. Amazon EC2 Systems Manager – To debug issues related to Amazon EC2 Systems Manager.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 176 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • AWS Glue – To troubleshoot issues related to the AWS Glue service • Amazon EMR – To troubleshoot issues related to Amazon EMR service. • Amazon Security Lake – To debug issues related to Amazon Security Lake. • AWS Systems Manager – To debug issues related to the Systems Manager service. • Amazon Verified Permissions – To troubleshoot issues related to Amazon Verified Permissions. • AWS IAM Access Analyzer – To debug issues related to the IAM Access Analyzer service. • AWS Backup – To troubleshoot issues related to AWS Backup. • AWS Database Migration Service – To troubleshoot	Aug 28, 2023

Change	Description	Date
	issues related to the DMS service. • Amazon DynamoDB – To debug issues related to Dynamo DB. • Amazon Elastic Container Registry (Amazon ECR) – To troubleshoot issues related to Amazon Elastic Container Registry (Amazon ECR). • Amazon Elastic Container Service – To debug issues related to Amazon Elastic Container Service. • Amazon Elastic Kubernetes Service – To troubleshoot issues related to Amazon Elastic Kubernetes Service. • Amazon EMR Serverless – To debug issues related to the Amazon EMR Serverless Service. • AWS Identity and Access Management – To troubleshoot issues related to AWS Identity and Access Management. • AWS Network Firewall – To troubleshoot issues related to AWS Network Firewall.	

Change	Description	Date
	• AWS HealthOmics – To debug issues related to AWS HealthOmics. • Amazon Quick – To debug issues related to Amazon Quick. • Amazon Relational Database Service – To troubleshoot issues related to Amazon Relational Database Service. • Amazon Redshift – To troubleshoot issues related to Amazon Redshift. • Amazon Redshift Serverless – To debug issues related to Amazon Redshift Serverless. • Amazon SageMaker AI – To debug issues related to Amazon SageMaker AI.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 141 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Lambda – To troubleshoot issues related to Lambda service. • Amazon Lex – To troubleshoot issues related to Amazon Lex service. • AWS Transfer – To debug issues related to Transfer service. • AWS Amplify – To debug issues related to Amplify service. • Amazon EventBridge Pipes – To troubleshoot permissions and billing issues related to Pipes. • Amazon EventBridge – To debug issues related to Amazon EventBridge • Amazon CloudWatch Logs – To troubleshoot issues related to Amazon CloudWatch Logs. • AWS Systems Manager – To troubleshoot issues related to Systems Manager.	June 26, 2023

Change	Description	Date
	• Amazon CloudWatch – To debug issues related to CloudWatch. • Amazon ElastiCache – To troubleshoot issues related to Amazon ElastiCache. • Amazon Athena – To debug issues related to Athena. • AWS Elastic Disaster Recovery – To troubleshoot issues related to Elastic Disaster Recovery. • Amazon CloudWatch – To troubleshoot configurations of Amazon CloudWatch. • Amazon EC2 – To debug issues related to the EC2 service. • AWS Certificate Manager – To troubleshoot issues related to Certificate Manager. • Amazon EventBridge Scheduler – To troubleshoot issues related to EventBridge Scheduler. • Amazon OpenSearch Service – To troubleshoot issues related to OpenSearch. • Amazon EventBridge Schemas – To debug issues	

Change	Description	Date
	related to EventBridge Schemas. • AWS User Notifications – To troubleshoot issues related to User Notifications. • Amazon CloudWatch Application Insights – To troubleshoot issues related to CloudWatch Application Insights. • Amazon DynamoDB – To troubleshoot issues related to DynamoDB. • Amazon DocumentDB Elastic Clusters – To troubleshoot issues related to DocumentDB Elastic Clusters.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 53 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Auto Scaling – To troubleshoot issues related to Auto Scaling service. • Amazon CloudWatch – To troubleshoot issues related to Amazon CloudWatch. • AWS Compute Optimizer – To troubleshoot issues related to Compute Optimizer. • Amazon CloudWatch Evidently – To troubleshoot issues related to Evidently. • EC2 Image Builder – To troubleshoot issues related to Image Builder service. • AWS IoT TwinMaker – To troubleshoot issues related to AWS IoT TwinMaker. • Amazon CloudWatch Logs – To troubleshoot issues related to Amazon CloudWatch Logs. • Amazon Pinpoint – To troubleshoot issues related to Amazon Pinpoint.	May 02, 2023

Change	Description	Date
	• AWS OAM Link – To debug issues related to OAM resources. • AWS Outposts – To troubleshoot issues related to AWS Outposts. • Amazon RDS – To debug issues related to Amazon RDS. • AWS Resource Explorer – To troubleshoot issues related to Resource Explorer. • Amazon CloudWatch RUM – To troubleshoot configurations of RUM service resources. • Amazon SNS – To troubleshoot issues related to Amazon SNS. • Amazon CloudWatch Synthetics – To troubleshoot issues related to CloudWatch Synthetics.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 52 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • AWS Backup gateway – To troubleshoot issues related to Backup gateway. • Amazon S3 – To debug issues related to Amazon S3. • AWS Transform MGN – To troubleshoot issues related to MGN. • AWS Clean Rooms – To debug issues related to AWS Clean Rooms; • AWS Systems Manager for SAP – To troubleshoot issues related to AWS Systems Manager for SAP. • Amazon VPC Lattice – To debug issues related to Amazon VPC Lattice.	March 16, 2023

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 220 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Amazon Athena – To enable AWS Support to develop tools that can be used to help customers with their queries related to Athena. • Amazon Chime – To troubleshoot issues related to Amazon Chime. • Amazon CloudWatch Internet Monitor – To debug issues related to Internet Monitor. • Amazon Comprehend – To troubleshoot issues related to Amazon Comprehend. • Amazon Elastic Compute Cloud – To debug issues related to Transit Gateway Connect and multicast features. • Amazon EventBridge Pipes – To troubleshoot issues related to EventBridge Pipes. • Amazon Interactive Video Service – To enable AWS Support to query Amazon	January 10, 2023

Change	Description	Date
	IVS resources to troubleshoot customer issues. • Amazon FSx – To enable AWS Support to develop tools to support importing and exporting for an Amazon FSx data repository. • Amazon GameLift Servers – To troubleshoot issues related to Amazon GameLift Servers. • AWS Glue– To troubleshoot issues related to AWS Glue Data Quality. • Amazon Kinesis Video Streams– To troubleshoot issues related to Kinesis Video Streams. • Amazon Managed Service for Prometheus – To troubleshoot issues related to Amazon Managed Service for Prometheus. • Amazon Managed Streaming for Apache Kafka – To troubleshoot issues related to Amazon MSK Connect. • AWS Network Manager – To troubleshoot issues related to Network Manager.	

Change	Description	Date
	• Amazon Nimble Studio – To debug issues related to Nimble Studio. • Amazon Personalize – To debug issues related to Amazon Personalize. • Amazon Pinpoint – To troubleshoot issues related to Amazon Pinpoint. • AWS HealthOmics – To troubleshoot issues related to HealthOmics. • Amazon Transcribe – To debug issues related to Amazon Transcribe.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 47 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • AWS Transform MGN – To troubleshoot replication and launch issues. • AWS CloudFormation hooks – To enable AWS Support to develop automation tools that can help resolve issues. • Amazon Elastic Kubernetes Service – To troubleshoot issues related to Amazon EKS. • AWS IoT FleetWise – To troubleshoot issues related to AWS IoT FleetWise. • AWS Mainframe Modernization – To debug issues related to AWS Mainframe Modernization. • AWS Outposts – To help AWS Support get a list of dedicated hosts and assets. • AWS Private 5G – To troubleshoot issues related to Private 5G.	October 4, 2022

Change	Description	Date
	AWS Tiros – To debug issues related to Tiros.	

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 46 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Amazon Managed Streaming for Apache Kafka – To troubleshoot issues related to Amazon MSK. • AWS DataSync – To troubleshoot issues related to DataSync. • AWS Elastic Disaster Recovery – To troubleshoot replication and launch issues. • Amazon GameSparks – To troubleshoot issues related to GameSparks. • AWS IoT TwinMaker – To debug issues related to AWS IoT TwinMaker. • AWS Lambda – To view the configuration of a function URL to troubleshooting issues. • Amazon Lookout for Equipment – To troubleshoot issues related to Lookout for Equipment.	August 17, 2022

Change	Description	Date
	• Amazon Route 53 and Amazon Route 53 Resolver – To get resolver configurations so that AWS Support can check the DNS resolution behavior of a VPC.	
AWSSupportServiceRolePolicy – Update to an existing policy	Added new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Amazon CloudWatch Logs – To help troubleshoot CloudWatch Logs related issues. • Amazon Interactive Video Service – To help Support check existing Amazon IVS resources for support cases regarding fraud or compromised accounts. • Amazon Inspector – To troubleshoot Amazon Inspector related issues. Removed permissions for services, such as Amazon WorkLink. Amazon WorkLink was deprecated on April 19, 2022.	June 23, 2022

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 25 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • AWS Amplify UI Builder – To troubleshoot issues related to component and theme generation. • Amazon AppStream – To troubleshoot issues by retrieving resources for features that launched recently. • AWS Backup – To troubleshoot issues related to backup jobs. • AWS CloudFormation – To perform diagnostics on issues related to IAM, extension, and versioning. • Amazon Kinesis – To troubleshoot issues related to Kinesis. • AWS Transfer Family – To troubleshoot issues related to Transfer Family.	April 27, 2022

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 54 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • Amazon Elastic Compute Cloud • To troubleshoot issues related to customer and AWS-managed prefixed lists. • To troubleshoot issues related to Amazon VPC IP Address Manager (IPAM). • AWS Network Manager – To troubleshoot issues related to Network Manager. • Savings Plans – To get metadata about outstanding Savings Plans commitments. • AWS Serverless Application Repository – To improve and support response actions as part of researching and resolving support cases. • Amazon WorkSpaces Web – To debug and troubleshoot issues with WorkSpaces Web services.	March 14, 2022

Change	Description	Date
AWSSupportServiceRolePolicy – Update to an existing policy	Added 74 new permissions to the following services to perform actions that help troubleshoot customer issues related to billing, administrative, and technical support: • AWS Transform MGN – To support agentless replication in the MGN. • AWS CloudFormation – To perform diagnostics on IAM, extension, and versioning related issues. • Amazon CloudWatch Logs – To validate resource policies. • Amazon EC2 Recycle Bin – To get metadata about Recycle Bin retention rules. • AWS Elastic Disaster Recovery – To troubleshoot replication and launch problems in customer accounts. • Amazon FSx – To view the description of Amazon FSx snapshots. • Amazon Lightsail – To view metadata and configurations details for Lightsail buckets.	February 17, 2022

Change	Description	Date
	• Amazon Macie – To view Macie configurations, such as classification jobs, custom data identifiers, regular expressions and findings. • Amazon S3 – To gather metadata and configurations for Amazon S3 buckets. • AWS Storage Gateway – To view metadata about customers' automatic tape creation policies. • Elastic Load Balancing – To view the description of resource limits when using the Service Quotas console. For more information, see Permission changes for AWSSupportServiceRolePolicy.	
Change log published	Change log for the AWS Support managed policies.	February 17, 2022

Permission changes for AWSSupportServiceRolePolicy

Most permissions added to AWSSupportServiceRolePolicy allow AWS Support to call an API operation with the same name. However, some API operations require permissions that have a different name.

The following table only lists the API operations that require permissions with a different name. This table describes these differences beginning on February 17, 2022.

Date	API operation name	Required policy permission
Added permissions on February 17, 2022	s3.GetBucketAnalyticsConfiguration	s3:GetAnalyticsConfiguration
	s3.ListBucketAnalyticsConfiguration	
	s3.GetBucketNotificationConfiguration	s3:GetBucketNotification
	s3.GetBucketEncryption	s3:GetEncryptionConfiguration
	s3.GetBucketIntelligentTieringConfiguration	s3:GetIntelligentTieringConfiguration
	s3.ListBucketIntelligentTieringConfiguration	
	s3.GetBucketInventoryConfiguration	s3:GetInventoryConfiguration
	s3.ListBucketInventoryConfiguration	
	s3.GetBucketLifecycleConfiguration	s3:GetLifecycleConfiguration
	s3.GetBucketMetricsConfiguration	s3:GetMetricsConfiguration
	s3.ListBucketMetricsConfiguration	

Date	API operation name	Required policy permission
	s3.GetBucketReplication	s3:GetReplicationConfiguration
	s3.HeadBucket	s3:ListBucket
	s3.ListObjects	
	s3.ListBuckets	s3:ListAllMyBuckets
	s3.ListMultipartUploads	s3:ListBucketMultipartUploads
	s3.ListObjectVersions	s3:ListBucketVersions
	s3.ListParts	s3:ListMultipartUploadParts
Added permissions on July 15, 2025	cloudcontrolapi:GetResource	cloudformation:GetResource
	cloudcontrolapi:ListResources	cloudformation:ListResources

AWS managed policies for AWS Support App in Slack

 **Note**

To access and view support cases in the AWS Support Center Console, see [Manage access to AWS Support Center](#).

AWS Support App has the following managed policies.

Contents

- [AWS managed policy: AWSSupportAppFullAccess](#)

- [AWS managed policy: AWSSupportAppReadOnlyAccess](#)
- [AWS Support App updates to AWS managed policies](#)

AWS managed policy: AWSSupportAppFullAccess

You can use the [AWSSupportAppFullAccess](#) managed policy to grant the IAM role the permissions to your Slack channel configurations. You can also attach the AWSSupportAppFullAccess policy to your IAM entities.

For more information, see [AWS Support App in Slack](#).

To view the permissions for this policy, see [AWSSupportAppFullAccess](#) in the *AWS Managed Policy Reference*.

Permissions details

This policy includes the following permissions:

- `servicequotas` – Describes your existing service quotas and requests, and creates service quota increases for your account.
- `support` – Creates, updates, and resolves your support cases. Updates and describes information about your cases, such as file attachments, correspondences, and severity levels. Initiates live chat sessions with a support agent.
- `iam` – Creates a service-linked role for Service Quotas.

For more information, see [Managing access to the AWS Support App](#).

AWS managed policy: AWSSupportAppReadOnlyAccess

The [AWSSupportAppReadOnlyAccess](#) policy grants permissions that allow the entity to perform read-only AWS Support App actions. For more information, see [AWS Support App in Slack](#).

To view the permissions for this policy, see [AWSSupportAppReadOnlyAccess](#) in the *AWS Managed Policy Reference*.

Permissions details

This policy includes the following permissions:

- `support` – Describes support case details and communications added to the support cases.

AWS Support App updates to AWS managed policies

View details about updates to AWS managed policies for AWS Support App since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the [Document history](#) page.

The following table describes important updates to the AWS Support App managed policies since August 17, 2022.

AWS Support App

Change	Description	Date
AWSSupportAppFullAccess and AWSSupportAppReadOnlyAccess New AWS managed policies for the AWS Support App	You can use these policies for the IAM role that you configure for your Slack channel configuration. For more information, see Managing access to the AWS Support App .	August 19, 2022
Change log published	Change log for the AWS Support App managed policies.	August 19, 2022

AWS managed policies for AWS Trusted Advisor

Trusted Advisor has the following AWS managed policies.

Contents

- [AWS managed policy: AWSTrustedAdvisorPriorityFullAccess](#)
- [AWS managed policy: AWSTrustedAdvisorPriorityReadOnlyAccess](#)
- [AWS managed policy: AWSTrustedAdvisorServiceRolePolicy](#)

- [AWS managed policy: AWSTrustedAdvisorReportingServiceRolePolicy](#)
- [Trusted Advisor updates to AWS managed policies](#)

AWS managed policy: AWSTrustedAdvisorPriorityFullAccess

The [AWSTrustedAdvisorPriorityFullAccess](#) policy grants full access to Trusted Advisor Priority. This policy also allows the user to add Trusted Advisor as a trusted service with AWS Organizations and to specify the delegated administrator accounts for Trusted Advisor Priority.

Permissions details

In the first statement, the policy includes the following permissions for `trustedadvisor`:

- Describes your account and organization.
- Describes identified risks from Trusted Advisor Priority. The permissions allow you to download and update the risk status.
- Describes your configurations for Trusted Advisor Priority email notifications. The permissions allow you to configure the email notifications and disable them for your delegated administrators.
- Sets up Trusted Advisor so that your account can enable AWS Organizations.

In the second statement, the policy includes the following permissions for `organizations`:

- Describes your Trusted Advisor account and organization.
- Lists the AWS services that you enabled to use Organizations.

In the third statement, the policy includes the following permissions for `organizations`:

- Lists the delegated administrators for Trusted Advisor Priority.
- Enables and disables trusted access with Organizations.

In the fourth statement, the policy includes the following permissions for `iam`:

- Creates the `AWSServiceRoleForTrustedAdvisorReporting` service-linked role.

In the fifth statement, the policy includes the following permissions for `organizations`:

- Allows you to register and deregister delegated administrators for Trusted Advisor Priority.

To view the full JSON policy document, see [AWSTrustedAdvisorPriorityFullAccess](#) in the *AWS Managed Policy Reference*.

AWS managed policy: AWSTrustedAdvisorPriorityReadOnlyAccess

The [AWSTrustedAdvisorPriorityReadOnlyAccess](#) policy grants read-only permissions to Trusted Advisor Priority, including permission to view the delegated administrator accounts.

Permissions details

In the first statement, the policy includes the following permissions for `trustedadvisor:`

- Describes your Trusted Advisor account and organization.
- Describes the identified risks from Trusted Advisor Priority and allows you to download them.
- Describes the configurations for Trusted Advisor Priority email notifications.

In the second and third statement, the policy includes the following permissions for `organizations:`

- Describes your organization with Organizations.
- Lists the AWS services that you enabled to use Organizations.
- Lists the delegated administrators for Trusted Advisor Priority

To view the full JSON policy document, see [AWSTrustedAdvisorPriorityReadOnlyAccess](#) in the *AWS Managed Policy Reference*.

AWS managed policy: AWSTrustedAdvisorServiceRolePolicy

This policy is attached to the `AWSServiceRoleForTrustedAdvisor` service-linked role. It allows the service-linked role to perform actions for you. You can't attach the [AWSTrustedAdvisorServiceRolePolicy](#) to your AWS Identity and Access Management (IAM) entities. For more information, see [Using service-linked roles for Trusted Advisor](#).

This policy grants administrative permissions that allow the service-linked role to access AWS services. These permissions allow the checks for Trusted Advisor to evaluate your account.

Permissions details

This policy includes the following permissions.

- `accessanalyzer` – Describes AWS Identity and Access Management Access Analyzer resources
- `Auto Scaling` – Describes Amazon EC2 Auto Scaling account quotas and resources
- `cloudformation` – Describes AWS CloudFormation (CloudFormation) account quotas and stacks
- `cloudfront` – Describes Amazon CloudFront distributions
- `cloudtrail` – Describes AWS CloudTrail (CloudTrail) trails
- `dynamodb` – Describes Amazon DynamoDB account quotas and resources
- `dynamodbaccelerator` – Describes DynamoDB Accelerator resources
- `ec2` – Describes Amazon Elastic Compute Cloud (Amazon EC2) account quotas and resources
- `elasticloadbalancing` – Describes Elastic Load Balancing (ELB) account quotas and resources
- `iam` – Gets IAM resources, such as credentials, password policy, and certificates
- `networkfirewall` – Describes AWS Network Firewall resources
- `kinesis` – Describes Amazon Kinesis (Kinesis) account quotas
- `rds` – Describes Amazon Relational Database Service (Amazon RDS) resources
- `redshift` – Describes Amazon Redshift resources
- `route53` – Describes Amazon Route 53 account quotas and resources
- `s3` – Describes Amazon Simple Storage Service (Amazon S3) resources
- `ses` – Gets Amazon Simple Email Service (Amazon SES) send quotas
- `sqs` – Lists Amazon Simple Queue Service (Amazon SQS) queues
- `cloudwatch` – Gets Amazon CloudWatch Events (CloudWatch Events) metric statistics
- `ce` – Gets Cost Explorer Service (Cost Explorer) recommendations
- `route53resolver` – Gets Amazon Route 53 Resolver Resolver Endpoints and resources
- `kafka` – Gets Amazon Managed Streaming for Apache Kafka resources
- `ecs` – Gets Amazon ECS resources

- `outposts` – Gets AWS Outposts resources

To view the full JSON policy document, see [AWSTrustedAdvisorServiceRolePolicy](#) in the *AWS Managed Policy Reference*.

AWS managed policy: `AWSTrustedAdvisorReportingServiceRolePolicy`

This policy is attached to the `AWSServiceRoleForTrustedAdvisorReporting` service-linked role that allows Trusted Advisor to perform actions for the organizational view feature. You can't attach the [AWSTrustedAdvisorReportingServiceRolePolicy](#) to your IAM entities. For more information, see [Using service-linked roles for Trusted Advisor](#).

This policy grants administrative permissions that allow the service-linked role to perform AWS Organizations actions.

Permissions details

This policy includes the following permissions.

- `organizations` – Describes your organization and lists the service access, accounts, parents, children, and organizational units

To view the full JSON policy document, see [AWSTrustedAdvisorReportingServiceRolePolicy](#) in the *AWS Managed Policy Reference*.

Trusted Advisor updates to AWS managed policies

View details about updates to AWS managed policies for AWS Support and Trusted Advisor since these services began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the [Document history](#) page.

The following table describes important updates to the Trusted Advisor managed policies since August 10, 2021.

Trusted Advisor

Change	Description	Date
AWSTrustedAdvisorServiceRolePolicy Update to an existing policy.	Trusted Advisor added new actions to grant the <code>ecs:ListClusters</code> , <code>ecs:ListTasks</code> , <code>ecs:DescribeTasks</code> , and <code>ecs:ListTaskDefinitionFamilies</code> permissions.	May 14, 2026
AWSTrustedAdvisorServiceRolePolicy Update to an existing policy.	Trusted Advisor added new actions to grant the <code>elasticloadbalancing:DescribeListeners</code> , and <code>elasticloadbalancing:DescribeRules</code> permissions.	October 30, 2024
AWSTrustedAdvisorServiceRolePolicy Update to an existing policy.	Trusted Advisor added new actions to grant the <code>access-analyzer:ListAnalyzers</code> , <code>cloudwatch:ListMetrics</code> , <code>dax:DescribeClusters</code> , <code>ec2:DescribeNatGateways</code> , <code>ec2:DescribeRouteTables</code> , <code>ec2:DescribeVpcEndpoints</code> , <code>ec2:GetManagedPrefixListEntries</code> , <code>elasticloadbalancing:Describe</code>	June 11, 2024

Change	Description	Date
	beTargetHealth , iam:ListSAMLProvid ers , kafka:Des cribeClusterV2 network-firewall:L istFirewalls network- firewall:DescribeFi rewall and sqs:GetQu eueAttributes permissio ns.	
AWSTrustedAdvisorServiceRolePolicy Update to an existing policy.	Trusted Advisor added new actions to grant the cloudtrail:GetTrai l cloudtrail:ListTra ils cloudtrai l:GetEventSelectors outposts:GetOutpost , outposts:ListAssets and outposts:ListOutpo sts permissions.	January 18, 2024
AWSTrustedAdvisorPriorityFullAccess Update to an existing policy.	Trusted Advisor updated the AWSTrustedAdvisorP riorityFullAccess AWS managed policy to include statement IDs.	December 6, 2023
AWSTrustedAdvisorPriorityReadOnlyAccess Update to an existing policy.	Trusted Advisor updated the AWSTrustedAdvisorP riorityReadOnlyAcc ess AWS managed policy to include statement IDs.	December 6, 2023

Change	Description	Date
AWSTrustedAdvisorServiceRolePolicy – Update to an existing policy	Trusted Advisor added new actions to grant the <code>ec2:DescribeRegions</code> , <code>s3:GetLifecycleConfiguration</code> , <code>ecs:DescribeTaskDefinition</code> and <code>ecs:ListTaskDefinitions</code> permissions.	November 9, 2023
AWSTrustedAdvisorServiceRolePolicy – Update to an existing policy	Trusted Advisor added new IAM actions <code>route53resolver:ListResolverEndpoints</code> , <code>route53resolver:ListResolverEndpointIpAddresses</code> , <code>ec2:DescribeSubnets</code> , <code>kafka:ListClustersV2</code> and <code>kafka:ListNodes</code> to onboard new resilience checks.	September 14, 2023
AWSTrustedAdvisorReportingServiceRolePolicy V2 of managed policy attached on Trusted Advisor <code>AWSServiceRoleForTrustedAdvisorReporting</code> service-linked role	Upgrade AWS managed policy to V2 for the Trusted Advisor <code>AWSServiceRoleForTrustedAdvisorReporting</code> service-linked role. The V2 will add one more IAM action <code>organizations:ListDelegatedAdministrators</code>	Feb 28, 2023

Change	Description	Date
AWSTrustedAdvisorPriorityFullAccess and AWSTrustedAdvisorPriorityReadOnlyAccess New AWS managed policies for the Trusted Advisor	Trusted Advisor added two new managed policies that you can use to control access to Trusted Advisor Priority.	August 17, 2022
AWSTrustedAdvisorServiceRolePolicy – Update to an existing policy	Trusted Advisor added new actions to grant the <code>DescribeTargetGroups</code> and <code>GetAccountPublicAccessBlock</code> permissions. The <code>DescribeTargetGroup</code> permission is required for the Auto Scaling Group Health Check to retrieve non-Classic Load Balancers that are attached to an Auto Scaling group. The <code>GetAccountPublicAccessBlock</code> permission is required for the Amazon S3 Bucket Permissions check to retrieve the block public access settings for an AWS account.	August 10, 2021
Change log published	Trusted Advisor started tracking changes for its AWS managed policies.	August 10, 2021

AWS managed policies for AWS Support Plans

AWS Support Plans has the following managed policies.

Contents

- [AWS managed policy: AWSSupportPlansFullAccess](#)
- [AWS managed policy: AWSSupportPlansReadOnlyAccess](#)
- [AWS managed policy: AWSSupportPlansServiceRolePolicy](#)
- [AWS Support Plans updates to AWS managed policies](#)

AWS managed policy: AWSSupportPlansFullAccess

AWS Support Plans uses the [AWSSupportPlansFullAccess](#) AWS managed policy. The IAM entity uses this policy to complete the following Support Plans actions for you:

- View your support plan for your AWS account
- View details about the status for a request to change your support plan
- Change the support plan for your AWS account
- Create support plan schedules for your AWS account
- View a list of all support plan modifiers for your AWS account
- Accept a support agreement for your AWS account
- Cancel a support agreement for your AWS account
- Create a support agreement for your AWS account
- View a support agreement for your AWS account
- View a list of support agreements for your AWS account
- View a list of support agreement revisions for your AWS account
- Reject a support agreement for your AWS account
- Update a support agreement for your AWS account
- Create the AWS Support Plans service-linked role for your AWS account
- View the AWS Support Plans service-linked role for your AWS account

To view the permissions for this policy, see [AWSSupportPlansFullAccess](#) in the *AWS Managed Policy Reference*.

For a list of changes to the policies, see [AWS Support Plans updates to AWS managed policies](#).

AWS managed policy: `AWSSupportPlansReadOnlyAccess`

AWS Support Plans uses the [AWSSupportPlansReadOnlyAccess](#) AWS managed policy. The IAM entity uses this policy to complete the following read-only Support Plans actions for you:

- View your support plan for your AWS account
- View details about the status for a request to change your support plan
- View a list of all support plan modifiers for your AWS account
- View a support agreement for your AWS account
- View a list of support agreements for your AWS account
- View a list of support agreement revisions for your AWS account

To view the permissions for this policy, see [AWSSupportPlansReadOnlyAccess](#) in the *AWS Managed Policy Reference*.

For a list of changes to the policies, see [AWS Support Plans updates to AWS managed policies](#).

AWS managed policy: `AWSSupportPlansServiceRolePolicy`

This policy is attached to the `AWSServiceRoleForSupportPlans` service-linked role and allows Support Plans to read and update AWS resources used to manage your account's support plan on your behalf. You can't attach this policy to your IAM entities. For more information, see [Using service-linked roles for AWS Support Plans](#).

To view the permissions for this policy, see [AWSSupportPlansServiceRolePolicy](#) in the *AWS Managed Policy Reference*.

AWS Support Plans updates to AWS managed policies

View details about updates to AWS managed policies for Support Plans since these services began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the [Document history](#) page.

The following table describes important updates to the Support Plans managed policies since September 29, 2022.

AWS Support

Change	Description	Date
AWSSupportPlansReadOnlyAccess - Update to an existing policy AWSSupportPlansFullAccess - Update to an existing policy	Added support agreement actions to the <code>AWSSupportPlansFullAccess</code> and <code>AWSSupportPlansReadOnlyAccess</code> managed policies. Also added service-linked role actions to the <code>AWSSupportPlansFullAccess</code> managed policy.	August 27, 2026
AWSSupportPlansReadOnlyAccess - Update to an existing policy AWSSupportPlansFullAccess - Update to an existing policy	Add <code>ListSupportPlanModifiers</code> action to <code>AWSSupportPlansFullAccess</code> and <code>AWSSupportPlansReadOnlyAccess</code> managed policies.	September 9, 2024
AWSSupportPlansFullAccess - Update to an existing policy	Add <code>CreateSupportPlanSchedule</code> action to <code>AWSSupportPlansFullAccess</code> managed policy.	May 8, 2023
Change log published	Change log for the Support Plans managed policies.	September 29, 2022

AWS managed policies for AWS Partner-Led Support

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS customers to use. We recommend that you

reduce permissions further by defining [customer managed policies](#) that are specific to your use cases.

You cannot change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see [AWS managed policies](#) in the *IAM User Guide*.

AWS managed policy: `AWSPartnerLedSupportReadOnlyAccess`

You can attach `AWSPartnerLedSupportReadOnlyAccess` to your users, groups, and roles.

This policy can be used to grant read-only access to APIs that can read service metadata for services in your AWS account. You can use this policy to provide your partners in the AWS Partner-Led Support Program with access to the services specified in the following permissions details section.

Important

Although `AWSPartnerLedSupportReadOnlyAccess` is a managed policy provided by AWS, you're responsible for reviewing the services and permissions that are included in the policy to verify that they meet your specific support requirements. Don't assume that this managed policy automatically includes all existing or new AWS services. You might need to create and maintain additional custom policies to cover services outside the scope of this managed policy.

Permissions details

This policy includes the following permissions.

- `acm` – Allow principals to troubleshoot technical support cases related to AWS Certificate Manager.
- `acm-pca` – Allow principals to troubleshoot technical support cases related to AWS Private Certificate Authority.
- `apigateway` – Allow principals to troubleshoot technical support cases related to Amazon API Gateway.
- `athena` – Allow principals to troubleshoot technical support cases related to Amazon Athena.
- `backup` – Allow principals to troubleshoot technical support cases related to AWS Backup.
- `backup-gateway` – Allow principals to troubleshoot technical support cases related to AWS Backup Gateway.
- `cloudformation` – Allow principals to troubleshoot technical support cases related to AWS CloudFormation.
- `cloudfront` – Allow principals to troubleshoot technical support cases related to Amazon CloudFront.
- `cloudtrail` – Allow principals to troubleshoot technical support cases related to AWS CloudTrail.
- `cloudwatch` – Allow principals to troubleshoot technical support cases related to Amazon CloudWatch.
- `codepipeline` – Allow principals to troubleshoot technical support cases related to AWS CodePipeline.
- `cognito-identity` – Allow principals to troubleshoot technical support cases related to Amazon Cognito Identity.
- `cognito-idp` – Allow principals to troubleshoot technical support cases related to Amazon Cognito user pools.
- `cognito-sync` – Allow principals to troubleshoot technical support cases related to Amazon Cognito Sync.
- `connect` – Allow principals to troubleshoot technical support cases related to Amazon Connect Customer.
- `directconnect` – Allow principals to troubleshoot technical support cases related to AWS Direct Connect.
- `dms` – Allow principals to troubleshoot technical support cases related to AWS Database Migration Service.

- `ds` – Allow principals to troubleshoot technical support cases related to AWS Directory Service.
- `ec2` – Allow principals to troubleshoot technical support cases related to Amazon Elastic Compute Cloud. This include technical support categories in EC2 (Windows and Linux), Virtual Private Cloud (VPC) and VPC.
- `ecs` – Allow principals to troubleshoot technical support cases related to Amazon Elastic Container Service.
- `eks` – Allow principals to troubleshoot technical support cases related to Amazon Elastic Kubernetes Service.
- `elasticache` – Allow principals to troubleshoot technical support cases related to Amazon ElastiCache.
- `elasticbeanstalk` – Allow principals to troubleshoot technical support cases related to AWS Elastic Beanstalk.
- `elasticfilesystem` – Allow principals to troubleshoot technical support cases related to Amazon Elastic File System.
- `elasticloadbalancing` – Allow principals to troubleshoot technical support cases related to Elastic Load Balancing.
- `emr-containers` – Allow principals to troubleshoot technical support cases related to Amazon EMR on EKS.
- `emr-serverless` – Allow principals to troubleshoot technical support cases related to Amazon EMR Serverless.
- `es` – Allow principals to troubleshoot technical support cases related to Amazon OpenSearch Service. This includes technical support categories such as OpenSearch Service Managed Cluster.
- `events` – Allow principals to troubleshoot technical support cases related to Amazon EventBridge.
- `fsx` – Allow principals to troubleshoot technical support cases related to Amazon FSx. This includes technical support categories such as FSX for Windows File Server.
- `glue` – Allow principals to troubleshoot technical support cases related to AWS Glue.
- `guardduty` – Allow principals to troubleshoot technical support cases related to Amazon GuardDuty.
- `iam` – Allow principals to troubleshoot technical support cases related to AWS Identity and Access Management.
- `kafka` – Allow principals to troubleshoot technical support cases related to Amazon Managed Streaming for Apache Kafka.

- `kafkaconnect` – Allow principals to troubleshoot technical support cases related to Amazon Managed Streaming for Apache Kafka Connect.
- `lambda` – Allow principals to troubleshoot technical support cases related to AWS Lambda.
- `logs` – Allow principals to troubleshoot technical support cases related to Amazon CloudWatch Logs.
- `medialive` – Allow principals to troubleshoot technical support cases related to AWS Elemental MediaLive.
- `mobiletargeting` – Allow principals to troubleshoot technical support cases related to Amazon Pinpoint.
- `pipes` – Allow principals to troubleshoot technical support cases related to Amazon EventBridge Pipes.
- `polly` – Allow principals to troubleshoot technical support cases related to Amazon Polly.
- `quicksight` – Allow principals to troubleshoot technical support cases related to Amazon Quick.
- `rds` – Allow principals to troubleshoot technical support cases related to Amazon Relational Database Service. This includes technical support categories such as: Relational Database Service (Aurora - MySQL-Compatible), Relational Database Service (Aurora - PostgreSQL-compatible), Relational Database Service (PostgreSQL), Relational Database Service (SQL Server), Relational Database Service (MySQL) and Relational Database Service (Oracle).
- `redshift` – Allow principals to troubleshoot technical support cases related to Amazon Redshift.
- `redshift-data` – Allow principals to troubleshoot technical support cases related to Amazon Redshift Data API.
- `redshift-serverless` – Allow principals to troubleshoot technical support cases related to Amazon Redshift Serverless.
- `route53` – Allow principals to troubleshoot technical support cases related to Amazon Route 53.
- `route53domains` – Allow principals to troubleshoot technical support cases related to Amazon Route 53 Domains.
- `route53-recovery-cluster` – Allow principals to troubleshoot technical support cases related to Amazon Route 53 Recovery Cluster.
- `route53-recovery-control-config` – Allow principals to troubleshoot technical support cases related to Amazon Route 53 Recovery Controls.

- `route53-recovery-readiness` – Allow principals to troubleshoot technical support cases related to Amazon Route 53 Recovery Readiness.
- `route53resolver` – Allow principals to troubleshoot technical support cases related to Amazon Route 53 Resolver.
- `s3` – Allow principals to troubleshoot technical support cases related to Amazon Simple Storage Service.
- `s3express` – Allow principals to troubleshoot technical support cases related to Amazon S3 Express.
- `sagemaker` – Allow principals to troubleshoot technical support cases related to Amazon SageMaker AI.
- `scheduler` – Allow principals to troubleshoot technical support cases related to Amazon EventBridge Scheduler.
- `servicequotas` – Allow principals to troubleshoot technical support cases related to Service Quotas.
- `ses` – Allow principals to troubleshoot technical support cases related to Amazon Simple Email Service.
- `sns` – Allow principals to troubleshoot technical support cases related to Amazon Simple Notification Service.
- `ssm` – Allow principals to troubleshoot technical support cases related to AWS Systems Manager.
- `ssm-contacts` – Allow principals to troubleshoot technical support cases related to AWS Systems Manager Incident Manager Contacts.
- `ssm-incidents` – Allow principals to troubleshoot technical support cases related to AWS Systems Manager Incident Manager.
- `ssm-sap` – Allow principals to troubleshoot technical support cases related to AWS Systems Manager for SAP.
- `swf` – Allow principals to troubleshoot technical support cases related to Amazon Simple Workflow Service.
- `vpc-lattice` – Allow principals to troubleshoot technical support cases related to Amazon VPC Lattice. This includes technical support categories such as VPC - Transit Gateway.
- `waf` – Allow principals to troubleshoot technical support cases related to AWS WAF.
- `waf-regional` – Allow principals to troubleshoot technical support cases related to AWS WAF Regional.

- `wafv2` – Allow principals to troubleshoot technical support cases related to AWS WAF V2.
- `workspaces` – Allow principals to troubleshoot technical support cases related to Amazon WorkSpaces. This includes technical support categories such as Workspaces (Windows).
- `workspaces-web` – Allow principals to troubleshoot technical support cases related to Amazon WorkSpaces Secure Browser. This includes technical support categories such as Workspaces (Windows).

To view the permissions for this policy, see [AWSPartnerLedSupportReadOnlyAccess](#) in the *AWS Managed Policy Reference*.

AWS Partner-Led Support updates to AWS managed policies

View details about updates to AWS managed policies for AWS Partner-Led Support since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the AWS Partner-Led Support Document history page.

Change	Description	Date
AWSPartnerLedSupportReadOnlyAccess – New policy	Added a new AWS managed policy that contains permissions that can read service metadata for services in your AWS account.	November 22, 2024
AWS Partner-Led Support started tracking changes	AWS Partner-Led Support started tracking changes for its AWS managed policies.	November 22, 2024

Manage access to AWS Support Center

You must have permissions to access Support Center and to [create a support case](#).

You can use one of the following options to access Support Center:

- Use AWS Identity and Access Management (IAM).
- Use the email address and password associated with your AWS account. This identity is called the AWS account *root user* (not recommended).

If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, you can also use the [Support API](#) to access Support and Trusted Advisor operations programmatically. For more information, see the [AWS Support API Reference](#).

Note

If you can't sign in to Support Center, you can use the [Contact Us](#) page instead. You can use this page to get help with billing and account issues.

For information about the Support Center Console API operations and how to add them to your IAM policies, see [Adding IAM policies for the Support Center Console API operations](#).

AWS account (not recommended)

You can sign in to the AWS Management Console and access the Support Center by using your AWS account email address and password. This identity is called the AWS account *root user*. However, we strongly recommend that you don't use the root user for your everyday tasks, even the administrative ones. Instead, we recommend that you use IAM, which lets you control who can perform certain tasks in your account.

AWS support actions

You can perform the following Support actions in the console. You can also specify these Support actions in an IAM policy to allow or deny specific actions.

Note

Denying any of the following actions in your IAM policies, might result in unintended behavior in Support Center when creating or interacting with a support case.

Action	Description
AddAttachmentsToSet	Grants permission to add one or more attachments to an attachment set. An attachment set is a temporary container for attachments that you add to a case or case communication. The set is available for 1 hour after it's created. The expiryTime returned in the response is when the set expires.
AddCommunicationToCase	Grants permission to add additional customer communication to an Support case, including a set of email addresses to copy on the communication.
CreateCase	Grants permission to create a case.
DescribeAttachment	Grants permission to retrieve an attachment on a case.
DescribeCaseAttributes	Grants permission to allow secondary services to read Support case attributes. *This is used internally by Support Center to get attributes tagged on your case.
DescribeCases	Grants permission to return a list of Support cases that matches a case ID or case IDs.
DescribeCommunication	Grants permission to get a single communication and attachments for a single AWS Support case.
DescribeCommunications	Grants permission to return communications and attachments for one or more Support cases.

Action	Description
DescribeCreateCaseOptions	Grants permission to return a list of CreateCaseOption types along with the corresponding supported hours and language availability.
DescribeIssueTypes	Grants permission to return issue types for Support cases. This is used internally by Support Center to get available issue types for your account.
DescribeServices	Grants permission to return the current list of AWS services and a list of service categories for each service. You then use service names and categories to create a case. Each AWS service has its own set of categories.
DescribeSeverityLevels	Grants permission to return the list of severity levels that you can assign to a Support case.
DescribeSupportedLanguages	Grants permission to return a list of supported languages for a specified categoryCode, issueType and serviceCode.
DescribeSupportLevel	Grants permission to return the support level for an AWS account identifier. This is used internally by Support Center to identify your support level.
DescribeTrustedAdvisorCheckRefreshStatuses	Grants permission to return the refresh status of the AWS Trusted Advisor checks that have the specified check IDs.
DescribeTrustedAdvisorCheckResult	Grants permission to return the results of the AWS Trusted Advisor check that has the specified check ID.

Action	Description
<code>DescribeTrustedAdvisorChecks</code>	Grants permission to return information about all available AWS Trusted Advisor checks, including the name, ID, category, description, and metadata.
<code>DescribeTrustedAdvisorCheckSummaries</code>	Grants permission to return the results for the AWS Trusted Advisor check summaries for the check IDs that you specified.
<code>GetInteraction</code>	Grants permission to retrieve details about a specific interaction by its unique identifier. This is used internally by Support Center to retrieve personalized recommendations.
<code>InitiateCallForCase</code>	Grants permission to initiate a call on Support Center. This is used internally by Support Center to start a call on your behalf.
<code>ListInteractionEntries</code>	Grants permission to retrieve a list of entries within a specific interaction, including messages, status updates, or other relevant data points. This is used internally by Support Center to track the detailed trail of an interaction.
<code>ListInteractions</code>	Grants permission to retrieve a list of interactions, potentially with filters or pagination. This is used internally by Support Center to manage and overview multiple interactions.
<code>InitiateChatForCase</code>	Grants permission to initiate a chat on Support Center. This is used internally by Support Center to start a chat on your behalf.

Action	Description
PutCaseAttributes	Grants permission to allow secondary services to attach attributes to Support cases. This is used internally by Support Center to add operational tags to your Support cases.
RateCaseCommunication	Grants permission to rate a Support case communication.
RefreshTrustedAdvisorCheck	Grants permission to refresh the AWS Trusted Advisor check that you specify using the check ID.
ResolveCase	Grants permission to resolve a Support case.
ResolveInteraction	Grants permission to mark an interaction as resolved using its unique identifier, indicating that the issue has been fully addressed and requires no further action. Once resolved, the interaction's status is set to CLOSED and becomes accessible to all users within the same account.
SearchForCases	Grants permission to return a list of Support cases that matches the given inputs. This is used internally by Support Center to find searched cases.
StartInteraction	Grants permission to initiate a new interaction to receive personalized troubleshooting assistance for account and technical issues. This is used internally by Support Center to initiate the troubleshooting process.

Action	Description
UpdateInteraction	Grants permission to update a specific interaction by its unique identifier with another message. This is used internally by Support Center to update the troubleshooting process.

IAM

By default, IAM users can't access the Support Center. You can use IAM to create individual users or groups. Then, you attach IAM policies to these entities, so that they have permission to perform actions and access resources, such as to open Support Center cases and use the Support API.

After you create IAM users, you can give those users individual passwords and an account-specific sign-in page. They can then sign in to your AWS account and work in the Support Center. IAM users who have AWS Support access can see all cases that are created for the account.

For more information, see [Sign in to the AWS Management Console as an IAM user](#) in the *IAM User Guide*.

The easiest way to grant permissions is to attach the AWS managed policy [AWSSupportAccess](#) to the user, group, or role. AWS Support allows action-level permissions to control access to specific AWS Support operations. AWS Support doesn't provide resource-level access, so the Resource element is always set to *. You can't allow or deny access to specific support cases.

Example: Allow access to all Support actions

The AWS managed policy [AWSSupportAccess](#) grants an IAM user access to Support. An IAM user with this policy can access all AWS Support operations and resources.

JSON

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



```

        "Action": ["support:*"],
        "Resource": "*"
      }
    ]
  }
}

```

For more information about how to attach the `AWSSupportAccess` policy to your entities, see [Adding IAM identity permissions \(console\)](#) in the *IAM User Guide*.

Example: Allow access to all actions except the `ResolveCase` action

You can also create *customer managed policies* in IAM to specify what actions to allow or deny. The following policy statement allows an IAM user to perform all actions in Support except resolve a case.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "support:*",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": "support:ResolveCase",
      "Resource": "*"
    }
  ]
}

```

For more information about how to create a customer managed IAM policy, see [Creating IAM policies \(console\)](#) in the *IAM User Guide*.

If the user or group already has a policy, you can add the AWS Support-specific policy statement to that policy.

⚠ Important

- If you can't view cases in the Support Center, make sure that you have the required permissions. You might need to contact your IAM administrator. For more information, see [Identity and access management for AWS Support](#).

Access to AWS Trusted Advisor

In the AWS Management Console, a separate `trustedadvisor` IAM namespace controls access to Trusted Advisor. In the Support API, the `support` IAM namespace controls access to Trusted Advisor. For more information, see [Manage access to AWS Trusted Advisor](#).

Manage access to AWS Support Plans

Topics

- [Permissions for the Support Plans console](#)
- [Support Plans actions](#)
- [Example IAM policies for Support Plans](#)
- [Troubleshooting](#)

Permissions for the Support Plans console

To access the Support Plans console, a user must have a minimum set of permissions. These permissions must allow the user to list and view details about the Support Plans resources in your AWS account.

You can create an AWS Identity and Access Management (IAM) policy with the `supportplans` namespace. You can use this policy to specify permissions for actions and resources.

When you create a policy, you can specify the namespace of the service to allow or deny an action. The namespace for Support Plans is `supportplans`.

You can use AWS managed policies and attach them to your IAM entities. For more information, see [AWS managed policies for AWS Support Plans](#).

Support Plans actions

You can perform the following Support Plans actions in the console. You can also specify these Support Plans actions in an IAM policy to allow or deny specific actions.

Action	Description
GetSupportPlan	Grants permission to view details about the current support plan for this AWS account.
GetSupportPlanUpdateStatus	Grants permission to view details about the status for a request to update a support plan.
StartSupportPlanUpdate	Grants permission to start the request to update the support plan for this AWS account.
CreateSupportPlanSchedule	Grants permission to create support plan schedules for this AWS account.
ListSupportPlanModifiers	Grants permission to view a list of all support plan modifiers for this AWS account.

Example IAM policies for Support Plans

You can use the following example policies to manage access to Support Plans.

Full access to Support Plans

The following policy allows users full access to Support Plans.

JSON

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

```
}  
]  
}
```

Read-only access to Support Plans

The following policy allows read-only access to Support Plans.

JSON

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": "supportplans:Get*",  
      "Resource": "*"   
    },  
    {  
      "Effect": "Allow",  
      "Action": "supportplans:List*",  
      "Resource": "*"   
    }  
  ]  
}
```

Deny access to Support Plans

The following policy doesn't allow users access to Support Plans.

JSON

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Deny",  
      "Action": "supportplans:*",  
      "Resource": "*"   
    }  
  ]  
}
```

```
} ]
```

Troubleshooting

See the following topics to manage access to Support Plans.

When I try to view or change my support plan, the Support Plans console says that I'm missing the `GetSupportPlan` permission

IAM users must have the required permissions to access the Support Plans console. You can update your IAM policy to include the missing permission or use an AWS managed policy, such as `AWSSupportPlansFullAccess` or `AWSSupportPlansReadOnlyAccess`. For more information, see [AWS managed policies for AWS Support Plans](#).

If you don't have access to update your IAM policies, contact your AWS account administrator.

Related information

For more information, see the following topics in the *IAM User Guide*:

- [Testing IAM policies with the IAM policy simulator](#)
- [Troubleshooting access denied error messages](#)

I have the correct Support Plans permissions, but I still get the same error

If your AWS account is a member account that's part of AWS Organizations, the service control policy (SCP) might need to be updated. SCPs are a type of policy that manages permissions in an organization.

Because Support Plans is a *global* service, policies that restrict AWS Regions might prevent member accounts from viewing or changing their support plan. To allow global services for your organization, such as IAM and Support Plans, you must add the service to the exclusion list in any applicable SCP. This means that accounts in the organization can access these services, even if the SCP denies a specified AWS Region.

To add Support Plans as an exception, enter `"supportplans: *"` to the `"NotAction"` list in the SCP.

```
"supportplans:*",
```

Your SCP might appear as the following policy snippet.

Example: SCP that allows Support Plans access in an organization

```
{ "Version": "2012-10-17",  
  "Statement": [  
    { "Sid": "GRREGIONDENY",  
      "Effect": "Deny",  
      "NotAction": [  
        "aws-portal:*",  
        "budgets:*",  
        "chime:*",  
        "iam:*",  
        "supportplans:*",  
        ....
```

If you have a member account and can't update the SCP, contact your AWS account administrator. The management account might need to update the SCP so that all member accounts can access Support Plans.

Notes for AWS Control Tower

- If your organization uses an SCP with AWS Control Tower, you can update the **Deny access to AWS based on the requested AWS Region** control (commonly referred to as the Region deny control).
- If you update the SCP for AWS Control Tower to allow supportplans, repairing the drift will remove your update to the SCP. For more information, see [Detect and resolve drift in AWS Control Tower](#).

Related information

For more information, see the following topics:

- [Service control policies \(SCPs\)](#) in the *AWS Organizations User Guide*.
- [Configure the Region deny control](#) in the *AWS Control Tower User Guide*
- [Deny access to AWS based on the requested AWS Region](#) in the *AWS Control Tower User Guide*

Manage access to AWS Trusted Advisor

You can access AWS Trusted Advisor from the AWS Management Console. All AWS accounts have access to a select core [Trusted Advisor checks](#). If you have a AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan, you can access all checks. for more information, see [AWS Trusted Advisor check reference](#).

You can use AWS Identity and Access Management (IAM) to control access to Trusted Advisor.

Topics

- [Permissions for the Trusted Advisor console](#)
- [Trusted Advisor actions](#)
- [IAM policy examples](#)
- [See also](#)

Permissions for the Trusted Advisor console

To access the Trusted Advisor console, a user must have a minimum set of permissions. These permissions must allow the user to list and view details about the Trusted Advisor resources in your AWS account.

You can use the following options to control access to Trusted Advisor:

- Use the tag filter feature of the Trusted Advisor console. The user or role must have permissions associated with the tags.

You can use AWS managed policies or custom policies to assign permissions by tags. For more information, see [Controlling access to and for IAM users and roles using tags](#).

- Create an IAM policy with the `trustedadvisor` namespace. You can use this policy to specify permissions for actions and resources.

When you create a policy, you can specify the namespace of the service to allow or deny an action. The namespace for Trusted Advisor is `trustedadvisor`. However, you can't use the `trustedadvisor` namespace to allow or deny Trusted Advisor API operations in the Support API. You must use the `support` namespace for Support instead.

Note

If you have permissions to the [AWS Support](#) API, the Trusted Advisor widget in the AWS Management Console shows a summary view of your Trusted Advisor results. To view your results in the Trusted Advisor console, you must have permission to the `trustedadvisor` namespace.

Trusted Advisor actions

You can perform the following Trusted Advisor actions in the console. You can also specify these Trusted Advisor actions in an IAM policy to allow or deny specific actions.

Action	Description
<code>DescribeAccount</code>	Grants permission to view the Support plan and various Trusted Advisor preferences.
<code>DescribeAccountAccess</code>	Grants permission to view if the AWS account has enabled or disabled Trusted Advisor.
<code>DescribeCheckItems</code>	Grants permission to view details for the check items.
<code>DescribeCheckRefreshStatuses</code>	Grants permission to view the refresh statuses for Trusted Advisor checks.
<code>DescribeCheckSummaries</code>	Grants permission to view Trusted Advisor check summaries.
<code>DescribeChecks</code>	Grants permission to view details for Trusted Advisor checks.
<code>DescribeNotificationPreferences</code>	Grants permission to view the notification preferences for the AWS account.
<code>ExcludeCheckItems</code>	Grants permission to exclude recommendations for Trusted Advisor checks.

Action	Description
IncludeCheckItems	Grants permission to include recommendations for Trusted Advisor checks.
RefreshCheck	Grants permission to refresh a Trusted Advisor check.
SetAccountAccess	Grants permission to enable or disable Trusted Advisor for the account.
UpdateNotificationPreferences	Grants permission to update notification preferences for Trusted Advisor.
DescribeCheckStatusHistoryChanges	Grants permission to view the results and changed statuses for checks in the last 30 days.

Trusted Advisor actions for organizational view

The following Trusted Advisor actions are for the organizational view feature. For more information, see [Organizational view for AWS Trusted Advisor](#).

Action	Description
DescribeOrganization	Grants permission to view if the AWS account meets the requirements to enable the organizational view feature.
DescribeOrganizationAccounts	Grants permission to view the linked AWS accounts that are in the organization.
DescribeReports	Grants permission to view details for organizational view reports, such as the report name, runtime, date created, status, and format.
DescribeServiceMetadata	Grants permission to view information about organizational view reports, such as the AWS

Action	Description
	Regions, check categories, check names, and resource statuses.
GenerateReport	Grants permission to create a report for Trusted Advisor checks in your organization.
ListAccountsForParent	Grants permission to view, in the Trusted Advisor console, all of the accounts in an AWS organization that are contained by a root or organizational unit (OU).
ListOrganizationalUnitsForParent	Grants permission to view, in the Trusted Advisor console, all of the organizational units (OUs) in a parent organizational unit or root.
ListRoots	Grants permission to view, in the Trusted Advisor console, all of the roots that are defined in an AWS organization.
SetOrganizationAccess	Grants permission to enable the organizational view feature for Trusted Advisor.

Trusted Advisor Priority actions

If you have Trusted Advisor Priority enabled for your account, you can perform the following Trusted Advisor actions in the console. You can also add these Trusted Advisor actions in an IAM policy to allow or deny specific actions. For more information, see [Example IAM policies for Trusted Advisor Priority](#).

Note

The risks that appear in Trusted Advisor Priority are recommendations that your technical account manager (TAM) has identified for your account. Recommendations from a service, such as a Trusted Advisor check, are created for you automatically. Recommendations from your TAM are created for you manually. Next, your TAM sends these recommendations so that they appear in Trusted Advisor Priority for your account.

For more information, see [Get started with AWS Trusted Advisor Priority](#).

Action	Description
DescribeRisks	Grants permission to view risks in Trusted Advisor Priority.
DescribeRisk	Grants permission to view risk details in Trusted Advisor Priority.
DescribeRiskResources	Grants permission to view affected resources for a risk in Trusted Advisor Priority.
DownloadRisk	Grants permission to download a file that contains details about the risk in Trusted Advisor Priority.
UpdateRiskStatus	Grants permission to update the risk status in Trusted Advisor Priority.
DescribeNotificationConfigurations	Grants permission to get your email notification preferences for Trusted Advisor Priority.
UpdateNotificationConfigurations	Grants permission to create or update your email notification preferences for Trusted Advisor Priority.
DeleteNotificationConfigurationForDelegatedAdmin	Grants permission to the organization management account to delete email notification preferences from a delegated administrator account for Trusted Advisor Priority.

IAM policy examples

The following policies show you how to allow and deny access to Trusted Advisor. You can use one of the following policies to create a *customer managed policy* in the IAM console. For example, you can copy an example policy, and then paste it into the [JSON tab](#) of the IAM console. Then, you attach the policy to your IAM user, group, or role.

For more information about how to create an IAM policy, see [Creating IAM policies \(console\)](#) in the *IAM User Guide*.

Examples

- [Full access to Trusted Advisor](#)
- [Read-only access to Trusted Advisor](#)
- [Deny access to Trusted Advisor](#)
- [Allow and deny specific actions](#)
- [Control access to the Support API operations for Trusted Advisor](#)
- [Example IAM policies for Trusted Advisor Priority](#)

Full access to Trusted Advisor

The following policy allows users to view and take all actions on all Trusted Advisor checks in the Trusted Advisor console.

JSON

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

Read-only access to Trusted Advisor

The following policy allows users read-only access to the Trusted Advisor console. Users can't make changes, such as refresh checks or change notification preferences.

JSON

```
{
```

```
"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Allow",
    "Action": [
      "trustedadvisor:Describe*",
      "trustedadvisor:Get*",
      "trustedadvisor:List*"
    ],
    "Resource": "*"
  }
]
```

Deny access to Trusted Advisor

The following policy doesn't allow users to view or take actions for Trusted Advisor checks in the Trusted Advisor console.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "trustedadvisor:*",
      "Resource": "*"
    }
  ]
}
```

Allow and deny specific actions

The following policy allows users to view all Trusted Advisor checks in the Trusted Advisor console, but doesn't allow them to refresh any checks.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "trustedadvisor:*",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": "trustedadvisor:RefreshCheck",
      "Resource": "*"
    }
  ]
}
```

Control access to the Support API operations for Trusted Advisor

In the AWS Management Console, a separate `trustedadvisor` IAM namespace controls access to Trusted Advisor. You can't use the `trustedadvisor` namespace to allow or deny Trusted Advisor API operations in the Support API. Instead, you use the `support` namespace. You must have permissions to the Support API to call Trusted Advisor programmatically.

For example, if you want to call the [RefreshTrustedAdvisorCheck](#) operation, you must have permissions to this action in the policy.

Example: Allow Trusted Advisor API operations only

The following policy allows users access to the Support API operations for Trusted Advisor, but not the rest of the Support API operations. For example, users can use the API to view and refresh checks. They can't create, view, update, or resolve AWS Support cases.

You can use this policy to call the Trusted Advisor API operations programmatically, but you can't use this policy to view or refresh checks in the Trusted Advisor console.

JSON

```
{
```

```

"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Allow",
    "Action": [
      "support:DescribeTrustedAdvisorCheckRefreshStatuses",
      "support:DescribeTrustedAdvisorCheckResult",
      "support:DescribeTrustedAdvisorChecks",
      "support:DescribeTrustedAdvisorCheckSummaries",
      "support:RefreshTrustedAdvisorCheck",
      "trustedadvisor:Describe*"
    ],
    "Resource": "*"
  },
  {
    "Effect": "Deny",
    "Action": [
      "support:AddAttachmentsToSet",
      "support:AddCommunicationToCase",
      "support:CreateCase",
      "support:DescribeAttachment",
      "support:DescribeCases",
      "support:DescribeCommunications",
      "support:DescribeServices",
      "support:DescribeSeverityLevels",
      "support:ResolveCase"
    ],
    "Resource": "*"
  }
]
}

```

For more information about how IAM works with Support and Trusted Advisor, see [Actions](#).

Example IAM policies for Trusted Advisor Priority

You can use the following AWS managed policies to control access to Trusted Advisor Priority. For more information, see [AWS managed policies for AWS Trusted Advisor](#) and [Get started with AWS Trusted Advisor Priority](#).

See also

For more information about Trusted Advisor permissions, see the following resources:

- [Actions defined by AWS Trusted Advisor](#) in the *IAM User Guide*.
- [Controlling Access to the Trusted Advisor Console](#)

Troubleshooting AWS Support identity and access

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

Topics

- [I'm not authorized to perform iam:PassRole](#)
- [I want to view my access keys](#)
- [I'm an administrator and want to allow others to access Support](#)
- [I want to allow people outside of my AWS account to access my Support resources](#)

I'm 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 Support.

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 Support. 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 view my access keys

After you create your IAM user access keys, you can view your access key ID at any time. However, you can't view your secret access key again. If you lose your secret key, you must create a new access key pair.

Access keys consist of two parts: an access key ID (for example, AKIAIOSFODNN7EXAMPLE) and a secret access key (for example, wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY). Like a user name and password, you must use both the access key ID and secret access key together to authenticate your requests. Manage your access keys as securely as you do your user name and password.

Important

Do not provide your access keys to a third party, even to help [find your canonical user ID](#). By doing this, you might give someone permanent access to your AWS account.

When you create an access key pair, you are prompted to save the access key ID and secret access key in a secure location. The secret access key is available only at the time you create it. If you lose your secret access key, you must add new access keys to your IAM user. You can have a maximum of two access keys. If you already have two, you must delete one key pair before creating a new one. To view instructions, see [Managing access keys](#) in the *IAM User Guide*.

I'm an administrator and want to allow others to access Support

To allow others to access Support, you must grant permission to the people or applications that need access. If you are using AWS IAM Identity Center to manage people and applications, you assign permission sets to users or groups to define their level of access. Permission sets automatically create and assign IAM policies to IAM roles that are associated with the person or application. For more information, see [Permission sets](#) in the *AWS IAM Identity Center User Guide*.

If you are not using IAM Identity Center, you must create IAM entities (users or roles) for the people or applications that need access. You must then attach a policy to the entity that grants them the correct permissions in Support. After the permissions are granted, provide the credentials to the user or application developer. They will use those credentials to access AWS. To learn more about creating IAM users, groups, policies, and permissions, see [IAM Identities](#) and [Policies and permissions in IAM](#) in the *IAM User Guide*.

I want to allow people outside of my AWS account to access my Support 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 Support supports these features, see [How AWS Support works 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*.

Incident response

Incident response for Support is an AWS responsibility. AWS has a formal, documented policy and program that governs incident response. For more information, see the [AWS Security Incident Response Technical Guide](#).

Use the following options to inform yourself about operational issues:

- View AWS operational issues with broad impact on the [AWS Service Health Dashboard](#). For example, events that affect a service or Region that isn't specific to your account.
- View operational issues for individual accounts in the [Health Dashboard](#). For example, events that affect services or resources in your account. For more information, see [Getting started with the Health Dashboard](#) in the *AWS Health User Guide*.

Logging and monitoring in AWS Support and AWS Trusted Advisor

Monitoring is an important part of maintaining the reliability, availability, and performance of AWS Support and AWS Trusted Advisor and your other AWS solutions. AWS provides the following monitoring tools to watch AWS Support and AWS Trusted Advisor, report when something is wrong, and take actions when appropriate:

- *Amazon CloudWatch* monitors your AWS resources and the applications that you run on AWS in real time. You can collect and track metrics, create customized dashboards, and set alarms that notify you or take actions when a specified metric reaches a threshold that you specify. For example, you can have CloudWatch track CPU usage or other metrics of your Amazon Elastic Compute Cloud (Amazon EC2) instances and automatically launch new instances when needed. For more information, see the [Amazon CloudWatch User Guide](#).
- *Amazon EventBridge* delivers a near real-time stream of system events that describe changes in AWS resources. EventBridge enables automated event-driven computing, as you can write rules that watch for certain events and trigger automated actions in other AWS services when these events happen. For more information, see the [Amazon EventBridge User Guide](#).
- *AWS CloudTrail* captures API calls and related events made by or on behalf of your AWS account and delivers the log files to an Amazon Simple Storage Service (Amazon S3) bucket that you specify. You can identify which users and accounts called AWS, the source IP address from which the calls were made, and when the calls occurred. For more information, see the [AWS CloudTrail User Guide](#).

For more information, see [Monitoring and logging for AWS Support](#) and [Monitoring and logging for AWS Trusted Advisor](#).

Compliance validation for AWS Support

Our new AWS sign-up experience is not designed for regulated workloads. If you're using our new AWS sign-up experience, but you want to use AWS for regulated workloads, you can [sign up for AWS \(advanced\)](#) or [activate advanced features](#) for your AWS environment.

To learn whether an AWS service is within the scope of specific compliance programs, see [AWS services in Scope by Compliance Program](#) and choose the compliance program that you are interested in. For general information, see [AWS Compliance Programs](#).

You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#).

Your compliance responsibility when using AWS services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. For more information about your compliance responsibility when using AWS services, see [AWS Security Documentation](#).

Resilience in AWS Support

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 zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see [AWS global infrastructure](#).

Infrastructure security in AWS Support

As a managed service, AWS Support is protected by the AWS global network security procedures that are described in the [Amazon Web Services: Overview of security processes](#) whitepaper.

You use AWS published API calls to access Support through the network. Clients must support Transport Layer Security (TLS) 1.0 or later. We recommend 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 IAM principal. Or you can use the [AWS Security Token Service](#) (AWS STS) to generate temporary security credentials to sign requests.

Configuration and vulnerability analysis in Support

For AWS Trusted Advisor, AWS handles basic security tasks such as guest operating system (OS) and database patching, firewall configuration, and disaster recovery.

Configuration and IT controls are a shared responsibility between AWS and you, our customer. For more information, see the AWS [shared responsibility model](#).

Code examples for Support using AWS SDKs

The following code examples show how to use Support with an AWS software development kit (SDK).

Basics are code examples that show you how to perform the essential operations within a service.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Scenarios are code examples that show you how to accomplish specific tasks by calling multiple functions within a service or combined with other AWS services.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Code examples

- [Basic examples for Support using AWS SDKs](#)
 - [Hello Support](#)
 - [Learn the basics of Support with an AWS SDK](#)
 - [Actions for Support using AWS SDKs](#)
 - [Use AddAttachmentsToSet with an AWS SDK or CLI](#)
 - [Use AddCommunicationToCase with an AWS SDK or CLI](#)
 - [Use CreateCase with an AWS SDK or CLI](#)
 - [Use DescribeAttachment with an AWS SDK or CLI](#)
 - [Use DescribeCases with an AWS SDK or CLI](#)
 - [Use DescribeCommunications with an AWS SDK or CLI](#)
 - [Use DescribeServices with an AWS SDK or CLI](#)
 - [Use DescribeSeverityLevels with an AWS SDK or CLI](#)
 - [Use DescribeTrustedAdvisorCheckRefreshStatuses with a CLI](#)
 - [Use DescribeTrustedAdvisorCheckResult with a CLI](#)
 - [Use DescribeTrustedAdvisorCheckSummaries with a CLI](#)
 - [Use DescribeTrustedAdvisorChecks with a CLI](#)

- [Use RefreshTrustedAdvisorCheck with a CLI](#)
- [Use ResolveCase with an AWS SDK or CLI](#)
- [Scenarios for Support using AWS SDKs](#)
 - [Getting started with technical support](#)

Basic examples for Support using AWS SDKs

The following code examples show how to use the basics of AWS Support with AWS SDKs.

Examples

- [Hello Support](#)
- [Learn the basics of Support with an AWS SDK](#)
- [Actions for Support using AWS SDKs](#)
 - [Use AddAttachmentsToSet with an AWS SDK or CLI](#)
 - [Use AddCommunicationToCase with an AWS SDK or CLI](#)
 - [Use CreateCase with an AWS SDK or CLI](#)
 - [Use DescribeAttachment with an AWS SDK or CLI](#)
 - [Use DescribeCases with an AWS SDK or CLI](#)
 - [Use DescribeCommunications with an AWS SDK or CLI](#)
 - [Use DescribeServices with an AWS SDK or CLI](#)
 - [Use DescribeSeverityLevels with an AWS SDK or CLI](#)
 - [Use DescribeTrustedAdvisorCheckRefreshStatuses with a CLI](#)
 - [Use DescribeTrustedAdvisorCheckResult with a CLI](#)
 - [Use DescribeTrustedAdvisorCheckSummaries with a CLI](#)
 - [Use DescribeTrustedAdvisorChecks with a CLI](#)
 - [Use RefreshTrustedAdvisorCheck with a CLI](#)
 - [Use ResolveCase with an AWS SDK or CLI](#)

Hello Support

The following code examples show how to get started using Support.

.NET

SDK for .NET

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
using Amazon.AWSSupport;
using Microsoft.Extensions.DependencyInjection;
using Microsoft.Extensions.Hosting;

public static class HelloSupport
{
    static async Task Main(string[] args)
    {
        // Use the AWS .NET Core Setup package to set up dependency injection for
        // the AWS Support service.
        // Use your AWS profile name, or leave it blank to use the default
        // profile.
        // You must have one of the following AWS Support plans: Business,
        // Enterprise On-Ramp, or Enterprise. Otherwise, an exception will be thrown.
        using var host = Host.CreateDefaultBuilder(args)
            .ConfigureServices((_, services) =>
                services.AddAWSService<IAmazonAWSSupport>()
            ).Build();

        // Now the client is available for injection.
        var supportClient =
            host.Services.GetRequiredService<IAmazonAWSSupport>();

        // You can use await and any of the async methods to get a response.
        var response = await supportClient.DescribeServicesAsync();
        Console.WriteLine($"{response.Services.Count} services available.");
    }
}
```


- For API details, see [DescribeServices](#) in *AWS SDK for .NET API Reference*.

Java

SDK for Java 2.x

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
import software.amazon.awssdk.regions.Region;
import software.amazon.awssdk.services.support.SupportClient;
import software.amazon.awssdk.services.support.model.Category;
import software.amazon.awssdk.services.support.model.DescribeServicesRequest;
import software.amazon.awssdk.services.support.model.DescribeServicesResponse;
import software.amazon.awssdk.services.support.model.Service;
import software.amazon.awssdk.services.support.model.SupportException;
import java.util.ArrayList;
import java.util.List;

/**
 * Before running this Java (v2) code example, set up your development
 * environment, including your credentials.
 *
 * For more information, see the following documentation topic:
 *
 * https://docs.aws.amazon.com/sdk-for-java/latest/developer-guide/get-
started.html
 *
 * In addition, you must have the AWS Business Support Plan to use the AWS
 * Support Java API. For more information, see:
 *
 * https://aws.amazon.com/premiumsupport/plans/
 *
 * This Java example performs the following task:
 *
 * 1. Gets and displays available services.
 *
 * NOTE: To see multiple operations, see SupportScenario.
```

```
*/

public class HelloSupport {
    public static void main(String[] args) {
        Region region = Region.US_WEST_2;
        SupportClient supportClient = SupportClient.builder()
            .region(region)
            .build();

        System.out.println("***** Step 1. Get and display available services.");
        displayServices(supportClient);
    }

    // Return a List that contains a Service name and Category name.
    public static void displayServices(SupportClient supportClient) {
        try {
            DescribeServicesRequest servicesRequest =
                DescribeServicesRequest.builder()
                    .language("en")
                    .build();

            DescribeServicesResponse response =
                supportClient.describeServices(servicesRequest);
            List<Service> services = response.services();

            System.out.println("Get the first 10 services");
            int index = 1;
            for (Service service : services) {
                if (index == 11)
                    break;

                System.out.println("The Service name is: " + service.name());

                // Display the Categories for this service.
                List<Category> categories = service.categories();
                for (Category cat : categories) {
                    System.out.println("The category name is: " + cat.name());
                }
                index++;
            }
        } catch (SupportException e) {
            System.out.println(e.getLocalizedMessage());
            System.exit(1);
        }
    }
}
```

```
    }  
  }  
}
```

- For API details, see [DescribeServices](#) in *AWS SDK for Java 2.x API Reference*.

JavaScript

SDK for JavaScript (v3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

Invoke `main()` to run the example.

```
import {  
  DescribeServicesCommand,  
  SupportClient,  
} from "@aws-sdk/client-support";  
  
// Change the value of 'region' to your preferred AWS Region.  
const client = new SupportClient({ region: "us-east-1" });  
  
const getServiceCount = async () => {  
  try {  
    const { services } = await client.send(new DescribeServicesCommand({}));  
    return services.length;  
  } catch (err) {  
    if (err.name === "SubscriptionRequiredException") {  
      throw new Error(  
        "You must be subscribed to the AWS Support plan to use this feature.",  
      );  
    }  
    throw err;  
  }  
};  
  
export const main = async () => {
```

```
try {
    const count = await getServiceCount();
    console.log(`Hello, AWS Support! There are ${count} services available.`);
} catch (err) {
    console.error("Failed to get service count: ", err.message);
}
};
```

- For API details, see [DescribeServices](#) in *AWS SDK for JavaScript API Reference*.

Kotlin

SDK for Kotlin

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
/**
```

Before running this Kotlin code example, set up your development environment, including your credentials.

For more information, see the following documentation topic:

<https://docs.aws.amazon.com/sdk-for-kotlin/latest/developer-guide/setup.html>

In addition, you must have the AWS Business Support Plan to use the AWS Support Java API. For more information, see:

<https://aws.amazon.com/premiumsupport/plans/>

This Kotlin example performs the following task:

1. Gets and displays available services.

```
*/
```

```
suspend fun main() {
    displaySomeServices()
}
```

```
// Return a List that contains a Service name and Category name.
suspend fun displaySomeServices() {
    val servicesRequest =
        DescribeServicesRequest {
            language = "en"
        }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response = supportClient.describeServices(servicesRequest)
        println("Get the first 10 services")
        var index = 1

        response.services?.forEach { service ->
            if (index == 11) {
                return@forEach
            }

            println("The Service name is: " + service.name)

            // Get the categories for this service.
            service.categories?.forEach { cat ->
                println("The category name is ${cat.name}")
                index++
            }
        }
    }
}
```

- For API details, see [DescribeServices](#) in *AWS SDK for Kotlin API reference*.

Python

SDK for Python (Boto3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```

import logging
import boto3
from botocore.exceptions import ClientError

logger = logging.getLogger(__name__)

def hello_support(support_client):
    """
    Use the AWS SDK for Python (Boto3) to create an AWS Support client and count
    the available services in your account.
    This example uses the default settings specified in your shared credentials
    and config files.

    :param support_client: A Boto3 Support Client object.
    """
    try:
        print("Hello, AWS Support! Let's count the available Support services:")
        response = support_client.describe_services()
        print(f"There are {len(response['services'])} services available.")
    except ClientError as err:
        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                "examples."
            )
        else:
            logger.error(
                "Couldn't count services. Here's why: %s: %s",
                err.response["Error"]["Code"],
                err.response["Error"]["Message"],
            )
            raise

if __name__ == "__main__":
    hello_support(boto3.client("support"))

```

- For API details, see [DescribeServices](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Learn the basics of Support with an AWS SDK

The following code examples show how to:

- Get and display available services and severity levels for cases.
- Create a support case using a selected service, category, and severity level.
- Get and display a list of open cases for the current day.
- Add an attachment set and a communication to the new case.
- Describe the new attachment and communication for the case.
- Resolve the case.
- Get and display a list of resolved cases for the current day.

.NET

SDK for .NET

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

Run an interactive scenario at a command prompt.

```
/// <summary>
/// Hello AWS Support example.
/// </summary>
public static class SupportCaseScenario
{
    /*
    Before running this .NET code example, set up your development environment,
    including your credentials.

    To use the AWS Support API, you must have one of the following AWS Support
    plans: Business, Enterprise On-Ramp, or Enterprise.
```

```

This .NET example performs the following tasks:
1. Get and display services. Select a service from the list.
2. Select a category from the selected service.
3. Get and display severity levels and select a severity level from the
list.
4. Create a support case using the selected service, category, and severity
level.
5. Get and display a list of open support cases for the current day.
6. Create an attachment set with a sample text file to add to the case.
7. Add a communication with the attachment to the support case.
8. List the communications of the support case.
9. Describe the attachment set.
10. Resolve the support case.
11. Get a list of resolved cases for the current day.
*/

private static SupportWrapper _supportWrapper = null!;

static async Task Main(string[] args)
{
    // Set up dependency injection for the AWS Support service.
    // Use your AWS profile name, or leave it blank to use the default
profile.
    using var host = Host.CreateDefaultBuilder(args)
        .ConfigureLogging(logging =>
            logging.AddFilter("System", LogLevel.Debug)
                .AddFilter<DebugLoggerProvider>("Microsoft",
LogLevel.Information)
                .AddFilter<ConsoleLoggerProvider>("Microsoft",
LogLevel.Trace))
        .ConfigureServices((_, services) =>
            services.AddAWSService<IAmazonAWSSupport>(new AWSOptions()
{ Profile = "default" }))
        .AddTransient<SupportWrapper>()
        )
        .Build();

    var logger = LoggerFactory.Create(builder =>
    {
        builder.AddConsole();
    }).CreateLogger(typeof(SupportCaseScenario));

    _supportWrapper = host.Services.GetRequiredService<SupportWrapper>();

```



```
Console.WriteLine(new string('-', 80));
Console.WriteLine("Welcome to the AWS Support case example scenario.");
Console.WriteLine(new string('-', 80));

try
{
    var apiSupported = await _supportWrapper.VerifySubscription();
    if (!apiSupported)
    {
        logger.LogError("You must have a Business, Enterprise On-Ramp, or
Enterprise Support " +
                        "plan to use the AWS Support API. \n\tPlease
upgrade your subscription to run these examples.");
        return;
    }

    var service = await DisplayAndSelectServices();

    var category = DisplayAndSelectCategories(service);

    var severityLevel = await DisplayAndSelectSeverity();

    var caseId = await CreateSupportCase(service, category,
severityLevel);

    await DescribeTodayOpenCases();

    var attachmentSetId = await CreateAttachmentSet();

    await AddCommunicationToCase(attachmentSetId, caseId);

    var attachmentId = await ListCommunicationsForCase(caseId);

    await DescribeCaseAttachment(attachmentId);

    await ResolveCase(caseId);

    await DescribeTodayResolvedCases();

    Console.WriteLine(new string('-', 80));
    Console.WriteLine("AWS Support case example scenario complete.");
    Console.WriteLine(new string('-', 80));
}
```

```

        catch (Exception ex)
        {
            logger.LogError(ex, "There was a problem executing the scenario.");
        }
    }

    /// <summary>
    /// List some available services from AWS Support, and select a service for
the example.
    /// </summary>
    /// <returns>The selected service.</returns>
    private static async Task<Service> DisplayAndSelectServices()
    {
        Console.WriteLine(new string('-', 80));
        var services = await _supportWrapper.DescribeServices();
        Console.WriteLine($"AWS Support client returned {services.Count}
services.");

        Console.WriteLine($"1. Displaying first 10 services:");
        for (int i = 0; i < 10 && i < services.Count; i++)
        {
            Console.WriteLine($"  {i + 1}. {services[i].Name}");
        }

        var choiceNumber = 0;
        while (choiceNumber < 1 || choiceNumber > services.Count)
        {
            Console.WriteLine(
                "Select an example support service by entering a number from the
preceding list:");
            var choice = Console.ReadLine();
            Int32.TryParse(choice, out choiceNumber);
        }
        Console.WriteLine(new string('-', 80));

        return services[choiceNumber - 1];
    }

    /// <summary>
    /// List the available categories for a service and select a category for the
example.
    /// </summary>
    /// <param name="service">Service to use for displaying categories.</param>
    /// <returns>The selected category.</returns>

```

```

private static Category DisplayAndSelectCategories(Service service)
{
    Console.WriteLine(new string('-', 80));

    Console.WriteLine($"2. Available support categories for Service
\"{service.Name}\"");
    for (int i = 0; i < service.Categories.Count; i++)
    {
        Console.WriteLine($"\\t{i + 1}. {service.Categories[i].Name}");
    }

    var choiceNumber = 0;
    while (choiceNumber < 1 || choiceNumber > service.Categories.Count)
    {
        Console.WriteLine(
            "Select an example support category by entering a number from the
preceding list:");
        var choice = Console.ReadLine();
        Int32.TryParse(choice, out choiceNumber);
    }

    Console.WriteLine(new string('-', 80));

    return service.Categories[choiceNumber - 1];
}

/// <summary>
/// List available severity levels from AWS Support, and select a level for
the example.
/// </summary>
/// <returns>The selected severity level.</returns>
private static async Task<SeverityLevel> DisplayAndSelectSeverity()
{
    Console.WriteLine(new string('-', 80));
    var severityLevels = await _supportWrapper.DescribeSeverityLevels();

    Console.WriteLine($"3. Get and display available severity levels:");
    for (int i = 0; i < 10 && i < severityLevels.Count; i++)
    {
        Console.WriteLine($"\\t{i + 1}. {severityLevels[i].Name}");
    }

    var choiceNumber = 0;
    while (choiceNumber < 1 || choiceNumber > severityLevels.Count)

```

```

        {
            Console.WriteLine(
                "Select an example severity level by entering a number from the
preceding list:");
            var choice = Console.ReadLine();
            Int32.TryParse(choice, out choiceNumber);
        }
        Console.WriteLine(new string('-', 80));

        return severityLevels[choiceNumber - 1];
    }

    /// <summary>
    /// Create an example support case.
    /// </summary>
    /// <param name="service">Service to use for the new case.</param>
    /// <param name="category">Category to use for the new case.</param>
    /// <param name="severity">Severity to use for the new case.</param>
    /// <returns>The caseId of the new support case.</returns>
    private static async Task<string> CreateSupportCase(Service service,
        Category category, SeverityLevel severity)
    {
        Console.WriteLine(new string('-', 80));
        Console.WriteLine($"4. Create an example support case" +
            $" with the following settings:" +
            $" \n\tService: {service.Name}, Category:
{category.Name} " +
            $"and Severity Level: {severity.Name}.");
        var caseId = await _supportWrapper.CreateCase(service.Code,
            category.Code, severity.Code,
            "Example case for testing, ignore.", "This is my example support
case.");

        Console.WriteLine($" \tNew case created with ID {caseId}");

        Console.WriteLine(new string('-', 80));

        return caseId;
    }

    /// <summary>
    /// List open cases for the current day.
    /// </summary>
    /// <returns>Async task.</returns>

```

```

private static async Task DescribeTodayOpenCases()
{
    Console.WriteLine($"5. List the open support cases for the current
day.");
    // Describe the cases. If it is empty, try again and allow time for the
new case to appear.
    List<CaseDetails> currentOpenCases = null!;
    while (currentOpenCases == null || currentOpenCases.Count == 0)
    {
        Thread.Sleep(1000);
        currentOpenCases = await _supportWrapper.DescribeCases(
            new List<string>(),
            null,
            false,
            false,
            DateTime.UtcNow.Date,
            DateTime.UtcNow);
    }

    foreach (var openCase in currentOpenCases)
    {
        Console.WriteLine($"\\tCase: {openCase.CaseId} created
{openCase.TimeCreated}");
    }

    Console.WriteLine(new string('-', 80));
}

/// <summary>
/// Create an attachment set for a support case.
/// </summary>
/// <returns>The attachment set id.</returns>
private static async Task<string> CreateAttachmentSet()
{
    Console.WriteLine(new string('-', 80));
    Console.WriteLine($"6. Create an attachment set for a support case.");
    var fileName = "example_attachment.txt";

    // Create the file if it does not already exist.
    if (!File.Exists(fileName))
    {
        await using StreamWriter sw = File.CreateText(fileName);
        await sw.WriteLineAsync(
            "This is a sample file for attachment to a support case.");
    }
}

```

```

    }

    await using var ms = new MemoryStream(await
File.ReadAllBytesAsync(fileName));

    var attachmentSetId = await _supportWrapper.AddAttachmentToSet(
        ms,
        fileName);

    Console.WriteLine($"\\tNew attachment set created with id: \\n
\\t{attachmentSetId.Substring(0, 65)}...");

    Console.WriteLine(new string('-', 80));

    return attachmentSetId;
}

/// <summary>
/// Add an attachment set and communication to a case.
/// </summary>
/// <param name="attachmentSetId">Id of the attachment set.</param>
/// <param name="caseId">Id of the case to receive the attachment set.</
param>
/// <returns>Async task.</returns>
private static async Task AddCommunicationToCase(string attachmentSetId,
string caseId)
{
    Console.WriteLine(new string('-', 80));
    Console.WriteLine($"7. Add attachment set and communication to
{caseId}.");

    await _supportWrapper.AddCommunicationToCase(
        caseId,
        "This is an example communication added to a support case.",
        attachmentSetId);

    Console.WriteLine($"\\tNew attachment set and communication added to
{caseId}");

    Console.WriteLine(new string('-', 80));
}

/// <summary>
/// List the communications for a case.

```

```

    /// </summary>
    /// <param name="caseId">Id of the case to describe.</param>
    /// <returns>An attachment id.</returns>
    private static async Task<string> ListCommunicationsForCase(string caseId)
    {
        Console.WriteLine(new string('-', 80));
        Console.WriteLine($"8. List communications for case {caseId}.");

        var communications = await
        _supportWrapper.DescribeCommunications(caseId);
        var attachmentId = "";
        foreach (var communication in communications)
        {
            Console.WriteLine(
                $"\\tCommunication created on: {communication.TimeCreated} has
{communication.AttachmentSet.Count} attachments.");
            if (communication.AttachmentSet.Any())
            {
                attachmentId = communication.AttachmentSet.First().AttachmentId;
            }
        }

        Console.WriteLine(new string('-', 80));
        return attachmentId;
    }

    /// <summary>
    /// Describe an attachment by id.
    /// </summary>
    /// <param name="attachmentId">Id of the attachment to describe.</param>
    /// <returns>Async task.</returns>
    private static async Task DescribeCaseAttachment(string attachmentId)
    {
        Console.WriteLine(new string('-', 80));
        Console.WriteLine($"9. Describe the attachment set.");

        var attachment = await _supportWrapper.DescribeAttachment(attachmentId);
        var data = Encoding.ASCII.GetString(attachment.Data.ToArray());
        Console.WriteLine($"\\tAttachment includes {attachment.FileName} with
data: \\n\\t{data}");

        Console.WriteLine(new string('-', 80));
    }

```

```
/// <summary>
/// Resolve the support case.
/// </summary>
/// <param name="caseId">Id of the case to resolve.</param>
/// <returns>Async task.</returns>
private static async Task ResolveCase(string caseId)
{
    Console.WriteLine(new string('-', 80));
    Console.WriteLine($"10. Resolve case {caseId}.");

    var status = await _supportWrapper.ResolveCase(caseId);
    Console.WriteLine($"\\tCase {caseId} has final status {status}");

    Console.WriteLine(new string('-', 80));
}

/// <summary>
/// List resolved cases for the current day.
/// </summary>
/// <returns>Async Task.</returns>
private static async Task DescribeTodayResolvedCases()
{
    Console.WriteLine(new string('-', 80));
    Console.WriteLine($"11. List the resolved support cases for the current
day.");
    var currentCases = await _supportWrapper.DescribeCases(
        new List<string>(),
        null,
        false,
        true,
        DateTime.UtcNow.Date,
        DateTime.UtcNow);

    foreach (var currentCase in currentCases)
    {
        if (currentCase.Status == "resolved")
        {
            Console.WriteLine(
                $"\\tCase: {currentCase.CaseId}: status
{currentCase.Status}");
        }
    }

    Console.WriteLine(new string('-', 80));
}
```



```
}  
}
```

Wrapper methods used by the scenario for Support actions.

```
/// <summary>  
/// Wrapper methods to use AWS Support for working with support cases.  
/// </summary>  
public class SupportWrapper  
{  
    private readonly IAmazonAWSSupport _amazonSupport;  
    public SupportWrapper(IAmazonAWSSupport amazonSupport)  
    {  
        _amazonSupport = amazonSupport;  
    }  
  
    /// <summary>  
    /// Get the descriptions of AWS services.  
    /// </summary>  
    /// <param name="name">Optional language for services.  
    /// Currently Chinese ("zh"), English ("en"), Japanese ("ja") and Korean  
    /// ("ko") are supported.</param>  
    /// <returns>The list of AWS service descriptions.</returns>  
    public async Task<List<Service>> DescribeServices(string language = "en")  
    {  
        var response = await _amazonSupport.DescribeServicesAsync(  
            new DescribeServicesRequest()  
            {  
                Language = language  
            });  
        return response.Services;  
    }  
  
    /// <summary>  
    /// Get the descriptions of support severity levels.  
    /// </summary>  
    /// <param name="name">Optional language for severity levels.
```

```

    /// Currently Chinese ("zh"), English ("en"), Japanese ("ja") and Korean
    ("ko") are supported.</param>
    /// <returns>The list of support severity levels.</returns>
    public async Task<List<SeverityLevel>> DescribeSeverityLevels(string language
    = "en")
    {
        var response = await _amazonSupport.DescribeSeverityLevelsAsync(
            new DescribeSeverityLevelsRequest()
            {
                Language = language
            });
        return response.SeverityLevels;
    }

    /// <summary>
    /// Create a new support case.
    /// </summary>
    /// <param name="serviceCode">Service code for the new case.</param>
    /// <param name="categoryCode">Category for the new case.</param>
    /// <param name="severityCode">Severity code for the new case.</param>
    /// <param name="subject">Subject of the new case.</param>
    /// <param name="body">Body text of the new case.</param>
    /// <param name="language">Optional language support for your case.
    /// Currently Chinese ("zh"), English ("en"), Japanese ("ja") and Korean
    ("ko") are supported.</param>
    /// <param name="attachmentSetId">Optional Id for an attachment set for the
    new case.</param>
    /// <param name="issueType">Optional issue type for the new case. Options are
    "customer-service" or "technical".</param>
    /// <returns>The caseId of the new support case.</returns>
    public async Task<string> CreateCase(string serviceCode, string categoryCode,
    string severityCode, string subject,
        string body, string language = "en", string? attachmentSetId = null,
    string issueType = "customer-service")
    {
        var response = await _amazonSupport.CreateCaseAsync(
            new CreateCaseRequest()
            {
                ServiceCode = serviceCode,
                CategoryCode = categoryCode,
                SeverityCode = severityCode,
                Subject = subject,

```

```

        Language = language,
        AttachmentSetId = attachmentSetId,
        IssueType = issueType,
        CommunicationBody = body
    });
    return response.CaseId;
}

/// <summary>
/// Add an attachment to a set, or create a new attachment set if one does
not exist.
/// </summary>
/// <param name="data">The data for the attachment.</param>
/// <param name="fileName">The file name for the attachment.</param>
/// <param name="attachmentSetId">Optional setId for the attachment. Creates
a new attachment set if empty.</param>
/// <returns>The setId of the attachment.</returns>
public async Task<string> AddAttachmentToSet(MemoryStream data, string
fileName, string? attachmentSetId = null)
{
    var response = await _amazonSupport.AddAttachmentsToSetAsync(
        new AddAttachmentsToSetRequest
        {
            AttachmentSetId = attachmentSetId,
            Attachments = new List<Attachment>
            {
                new Attachment
                {
                    Data = data,
                    FileName = fileName
                }
            }
        });
    return response.AttachmentSetId;
}

/// <summary>
/// Get description of a specific attachment.
/// </summary>

```

```

    /// <param name="attachmentId">Id of the attachment, usually fetched by
    describing the communications of a case.</param>
    /// <returns>The attachment object.</returns>
    public async Task<Attachment> DescribeAttachment(string attachmentId)
    {
        var response = await _amazonSupport.DescribeAttachmentAsync(
            new DescribeAttachmentRequest()
            {
                AttachmentId = attachmentId
            });
        return response.Attachment;
    }

    /// <summary>
    /// Add communication to a case, including optional attachment set ID and CC
    email addresses.
    /// </summary>
    /// <param name="caseId">Id for the support case.</param>
    /// <param name="body">Body text of the communication.</param>
    /// <param name="attachmentSetId">Optional Id for an attachment set.</param>
    /// <param name="ccEmailAddresses">Optional list of CC email addresses.</
param>
    /// <returns>True if successful.</returns>
    public async Task<bool> AddCommunicationToCase(string caseId, string body,
        string? attachmentSetId = null, List<string>? ccEmailAddresses = null)
    {
        var response = await _amazonSupport.AddCommunicationToCaseAsync(
            new AddCommunicationToCaseRequest()
            {
                CaseId = caseId,
                CommunicationBody = body,
                AttachmentSetId = attachmentSetId,
                CcEmailAddresses = ccEmailAddresses
            });
        return response.Result;
    }

    /// <summary>
    /// Describe the communications for a case, optionally with a date filter.
    /// </summary>

```

```

    /// <param name="caseId">The ID of the support case.</param>
    /// <param name="afterTime">The optional start date for a filtered search.</
param>
    /// <param name="beforeTime">The optional end date for a filtered search.</
param>
    /// <returns>The list of communications for the case.</returns>
    public async Task<List<Communication>> DescribeCommunications(string caseId,
DateTime? afterTime = null, DateTime? beforeTime = null)
    {
        var results = new List<Communication>();
        var paginateCommunications =
        _amazonSupport.Paginators.DescribeCommunications(
            new DescribeCommunicationsRequest()
            {
                CaseId = caseId,
                AfterTime = afterTime?.ToString("s"),
                BeforeTime = beforeTime?.ToString("s")
            });
        // Get the entire list using the paginator.
        await foreach (var communications in
paginateCommunications.Communications)
        {
            results.Add(communications);
        }
        return results;
    }

    /// <summary>
    /// Get case details for a list of case ids, optionally with date filters.
    /// </summary>
    /// <param name="caseIds">The list of case IDs.</param>
    /// <param name="displayId">Optional display ID.</param>
    /// <param name="includeCommunication">True to include communication.
Defaults to true.</param>
    /// <param name="includeResolvedCases">True to include resolved cases.
Defaults to false.</param>
    /// <param name="afterTime">The optional start date for a filtered search.</
param>
    /// <param name="beforeTime">The optional end date for a filtered search.</
param>
    /// <param name="language">Optional language support for your case.

```

```

    /// Currently Chinese ("zh"), English ("en"), Japanese ("ja") and Korean
    /// ("ko") are supported.</param>
    /// <returns>A list of CaseDetails.</returns>
    public async Task<List<CaseDetails>> DescribeCases(List<string> caseIds,
string? displayId = null, bool includeCommunication = true,
    bool includeResolvedCases = false, DateTime? afterTime = null, DateTime?
beforeTime = null,
    string language = "en")
    {
        var results = new List<CaseDetails>();
        var paginateCases = _amazonSupport.Paginators.DescribeCases(
            new DescribeCasesRequest()
            {
                CaseIdList = caseIds,
                DisplayId = displayId,
                IncludeCommunications = includeCommunication,
                IncludeResolvedCases = includeResolvedCases,
                AfterTime = afterTime?.ToString("s"),
                BeforeTime = beforeTime?.ToString("s"),
                Language = language
            });
        // Get the entire list using the paginator.
        await foreach (var cases in paginateCases.Cases)
        {
            results.Add(cases);
        }
        return results;
    }

    /// <summary>
    /// Resolve a support case by caseId.
    /// </summary>
    /// <param name="caseId">Id for the support case.</param>
    /// <returns>The final status of the case after resolving.</returns>
    public async Task<string> ResolveCase(string caseId)
    {
        var response = await _amazonSupport.ResolveCaseAsync(
            new ResolveCaseRequest()
            {
                CaseId = caseId
            });
        return response.FinalCaseStatus;
    }

```

```
}

/// <summary>
/// Verify the support level for AWS Support API access.
/// </summary>
/// <returns>True if the subscription level supports API access.</returns>
public async Task<bool> VerifySubscription()
{
    try
    {
        var response = await _amazonSupport.DescribeServicesAsync(
            new DescribeServicesRequest()
            {
                Language = "en"
            });
        return response.HttpStatusCode == HttpStatusCode.OK;
    }
    catch (Amazon.AWSSupport.AmazonAWSSupportException ex)
    {
        if (ex.ErrorCode == "SubscriptionRequiredException")
        {
            return false;
        }
        else throw;
    }
}
}
```

- For API details, see the following topics in *AWS SDK for .NET API Reference*.
 - [AddAttachmentsToSet](#)
 - [AddCommunicationToCase](#)
 - [CreateCase](#)
 - [DescribeAttachment](#)
 - [DescribeCases](#)
 - [DescribeCommunications](#)
 - [DescribeServices](#)
 - [DescribeSeverityLevels](#)
 - [ResolveCase](#)

Java

SDK for Java 2.x

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

Run various Support operations.

```
import software.amazon.awssdk.core.SdkBytes;
import software.amazon.awssdk.regions.Region;
import software.amazon.awssdk.services.support.SupportClient;
import software.amazon.awssdk.services.support.model.AddAttachmentsToSetResponse;
import
    software.amazon.awssdk.services.support.model.AddCommunicationToCaseRequest;
import
    software.amazon.awssdk.services.support.model.AddCommunicationToCaseResponse;
import software.amazon.awssdk.services.support.model.Attachment;
import software.amazon.awssdk.services.support.model.AttachmentDetails;
import software.amazon.awssdk.services.support.model.CaseDetails;
import software.amazon.awssdk.services.support.model.Category;
import software.amazon.awssdk.services.support.model.Communication;
import software.amazon.awssdk.services.support.model.CreateCaseRequest;
import software.amazon.awssdk.services.support.model.CreateCaseResponse;
import software.amazon.awssdk.services.support.model.DescribeAttachmentRequest;
import software.amazon.awssdk.services.support.model.DescribeAttachmentResponse;
import software.amazon.awssdk.services.support.model.DescribeCasesRequest;
import software.amazon.awssdk.services.support.model.DescribeCasesResponse;
import
    software.amazon.awssdk.services.support.model.DescribeCommunicationsRequest;
import
    software.amazon.awssdk.services.support.model.DescribeCommunicationsResponse;
import software.amazon.awssdk.services.support.model.DescribeServicesRequest;
import software.amazon.awssdk.services.support.model.DescribeServicesResponse;
import
    software.amazon.awssdk.services.support.model.DescribeSeverityLevelsRequest;
import
    software.amazon.awssdk.services.support.model.DescribeSeverityLevelsResponse;
import software.amazon.awssdk.services.support.model.ResolveCaseRequest;
import software.amazon.awssdk.services.support.model.ResolveCaseResponse;
```



```
import software.amazon.awssdk.services.support.model.Service;
import software.amazon.awssdk.services.support.model.SeverityLevel;
import software.amazon.awssdk.services.support.model.SupportException;
import software.amazon.awssdk.services.support.model.AddAttachmentsToSetRequest;
import java.io.File;
import java.io.FileInputStream;
import java.io.FileNotFoundException;
import java.io.InputStream;
import java.time.Instant;
import java.time.temporal.ChronoUnit;
import java.util.ArrayList;
import java.util.List;

/**
 * Before running this Java (v2) code example, set up your development
 * environment, including your credentials.
 *
 * For more information, see the following documentation topic:
 *
 * https://docs.aws.amazon.com/sdk-for-java/latest/developer-guide/get-started.html
 *
 * In addition, you must have the AWS Business Support Plan to use the AWS
 * Support Java API. For more information, see:
 *
 * https://aws.amazon.com/premiumsupport/plans/
 *
 * This Java example performs the following tasks:
 *
 * 1. Gets and displays available services.
 * 2. Gets and displays severity levels.
 * 3. Creates a support case by using the selected service, category, and
 * severity level.
 * 4. Gets a list of open cases for the current day.
 * 5. Creates an attachment set with a generated file.
 * 6. Adds a communication with the attachment to the support case.
 * 7. Lists the communications of the support case.
 * 8. Describes the attachment set included with the communication.
 * 9. Resolves the support case.
 * 10. Gets a list of resolved cases for the current day.
 */
public class SupportScenario {
```

```

    public static final String DASHES = new String(new char[80]).replace("\0",
    "-");

    public static void main(String[] args) {
        final String usage = ""

            Usage:
            <fileAttachment>Where:
            fileAttachment - The file can be a simple saved .txt file to
use as an email attachment.\s
            "";

        // if (args.length != 1) {
        //     System.out.println(usage);
        //     System.exit(1);
        // }

        String fileAttachment = "C:\\\\AWS\\test.txt" ; //args[0];
        Region region = Region.US_WEST_2;
        SupportClient supportClient = SupportClient.builder()
            .region(region)
            .build();

        System.out.println(DASHES);
        System.out.println("***** Welcome to the AWS Support case example
scenario.");
        System.out.println(DASHES);

        System.out.println(DASHES);
        System.out.println("1. Get and display available services.");
        List<String> sevCatList = displayServices(supportClient);
        System.out.println(DASHES);

        System.out.println(DASHES);
        System.out.println("2. Get and display Support severity levels.");
        String sevLevel = displaySevLevels(supportClient);
        System.out.println(DASHES);

        System.out.println(DASHES);
        System.out.println("3. Create a support case using the selected service,
category, and severity level.");
        String caseId = createSupportCase(supportClient, sevCatList, sevLevel);
        if (caseId.compareTo("") == 0) {
            System.out.println("A support case was not successfully created!");

```

```
        System.exit(1);
    } else
        System.out.println("Support case " + caseId + " was successfully
created!");
    System.out.println(DASHES);

    System.out.println(DASHES);
    System.out.println("4. Get open support cases.");
    getOpenCase(supportClient);
    System.out.println(DASHES);

    System.out.println(DASHES);
    System.out.println("5. Create an attachment set with a generated file to
add to the case.");
    String attachmentSetId = addAttachment(supportClient, fileAttachment);
    System.out.println("The Attachment Set id value is" + attachmentSetId);
    System.out.println(DASHES);

    System.out.println(DASHES);
    System.out.println("6. Add communication with the attachment to the
support case.");
    addAttachSupportCase(supportClient, caseId, attachmentSetId);
    System.out.println(DASHES);

    System.out.println(DASHES);
    System.out.println("7. List the communications of the support case.");
    String attachId = listCommunications(supportClient, caseId);
    System.out.println("The Attachment id value is" + attachId);
    System.out.println(DASHES);

    System.out.println(DASHES);
    System.out.println("8. Describe the attachment set included with the
communication.");
    describeAttachment(supportClient, attachId);
    System.out.println(DASHES);

    System.out.println(DASHES);
    System.out.println("9. Resolve the support case.");
    resolveSupportCase(supportClient, caseId);
    System.out.println(DASHES);

    System.out.println(DASHES);
    System.out.println("10. Get a list of resolved cases for the current
day.");
```

```

        getResolvedCase(supportClient);
        System.out.println(DASHES);

        System.out.println(DASHES);
        System.out.println("***** This Scenario has successfully completed");
        System.out.println(DASHES);
    }

    public static void getResolvedCase(SupportClient supportClient) {
        try {
            // Specify the start and end time.
            Instant now = Instant.now();
            java.time.LocalDate.now();
            Instant yesterday = now.minus(1, ChronoUnit.DAYS);

            DescribeCasesRequest describeCasesRequest =
DescribeCasesRequest.builder()
                .maxResults(30)
                .afterTime(yesterday.toString())
                .beforeTime(now.toString())
                .includeResolvedCases(true)
                .build();

            DescribeCasesResponse response =
supportClient.describeCases(describeCasesRequest);
            List<CaseDetails> cases = response.cases();
            for (CaseDetails sinCase : cases) {
                if (sinCase.status().compareTo("resolved") == 0)
                    System.out.println("The case status is " + sinCase.status());
            }

        } catch (SupportException e) {
            System.out.println(e.getLocalizedMessage());
            System.exit(1);
        }
    }

    public static void resolveSupportCase(SupportClient supportClient, String
caseId) {
        try {
            ResolveCaseRequest caseRequest = ResolveCaseRequest.builder()
                .caseId(caseId)
                .build();

```

```
        ResolveCaseResponse response =
supportClient.resolveCase(caseRequest);
        System.out.println("The status of case " + caseId + " is " +
response.finalCaseStatus());

    } catch (SupportException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
}

public static void describeAttachment(SupportClient supportClient, String
attachId) {
    try {
        DescribeAttachmentRequest attachmentRequest =
DescribeAttachmentRequest.builder()
            .attachmentId(attachId)
            .build();

        DescribeAttachmentResponse response =
supportClient.describeAttachment(attachmentRequest);
        System.out.println("The name of the file is " +
response.attachment().fileName());

    } catch (SupportException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
}

public static String listCommunications(SupportClient supportClient, String
caseId) {
    try {
        String attachId = null;
        DescribeCommunicationsRequest communicationsRequest =
DescribeCommunicationsRequest.builder()
            .caseId(caseId)
            .maxResults(10)
            .build();

        DescribeCommunicationsResponse response =
supportClient.describeCommunications(communicationsRequest);
        List<Communication> communications = response.communications();
        for (Communication comm : communications) {
```

```
        System.out.println("the body is: " + comm.body());

        // Get the attachment id value.
        List<AttachmentDetails> attachments = comm.attachmentSet();
        for (AttachmentDetails detail : attachments) {
            attachId = detail.attachmentId();
        }
    }
    return attachId;

} catch (SupportException e) {
    System.out.println(e.getLocalizedMessage());
    System.exit(1);
}
return "";
}

public static void addAttachSupportCase(SupportClient supportClient, String
caseId, String attachmentSetId) {
    try {
        AddCommunicationToCaseRequest caseRequest =
AddCommunicationToCaseRequest.builder()
            .caseId(caseId)
            .attachmentSetId(attachmentSetId)
            .communicationBody("Please refer to attachment for details.")
            .build();

        AddCommunicationToCaseResponse response =
supportClient.addCommunicationToCase(caseRequest);
        if (response.result())
            System.out.println("You have successfully added a communication
to an AWS Support case");
        else
            System.out.println("There was an error adding the communication
to an AWS Support case");

    } catch (SupportException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
}

public static String addAttachment(SupportClient supportClient, String
fileAttachment) {
```

```
try {
    File myFile = new File(fileAttachment);
    InputStream sourceStream = new FileInputStream(myFile);
    SdkBytes sourceBytes = SdkBytes.fromInputStream(sourceStream);

    Attachment attachment = Attachment.builder()
        .fileName(myFile.getName())
        .data(sourceBytes)
        .build();

    AddAttachmentsToSetRequest setRequest =
AddAttachmentsToSetRequest.builder()
        .attachments(attachment)
        .build();

    AddAttachmentsToSetResponse response =
supportClient.addAttachmentsToSet(setRequest);
    return response.attachmentSetId();

} catch (SupportException | FileNotFoundException e) {
    System.out.println(e.getLocalizedMessage());
    System.exit(1);
}
return "";
}

public static void getOpenCase(SupportClient supportClient) {
    try {
        // Specify the start and end time.
        Instant now = Instant.now();
        java.time.LocalDate.now();
        Instant yesterday = now.minus(1, ChronoUnit.DAYS);

        DescribeCasesRequest describeCasesRequest =
DescribeCasesRequest.builder()
            .maxResults(20)
            .afterTime(yesterday.toString())
            .beforeTime(now.toString())
            .build();

        DescribeCasesResponse response =
supportClient.describeCases(describeCasesRequest);
        List<CaseDetails> cases = response.cases();
        for (CaseDetails sinCase : cases) {
```

```

        System.out.println("The case status is " + sinCase.status());
        System.out.println("The case Id is " + sinCase.caseId());
        System.out.println("The case subject is " + sinCase.subject());
    }

    } catch (SupportException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
}

public static String createSupportCase(SupportClient supportClient,
List<String> sevCatList, String sevLevel) {
    try {
        String serviceCode = sevCatList.get(0);
        String caseCat = sevCatList.get(1);
        CreateCaseRequest caseRequest = CreateCaseRequest.builder()
            .categoryCode(caseCat.toLowerCase())
            .serviceCode(serviceCode.toLowerCase())
            .severityCode(sevLevel.toLowerCase())
            .communicationBody("Test issue with " +
serviceCode.toLowerCase())
            .subject("Test case, please ignore")
            .language("en")
            .issueType("technical")
            .build();

        CreateCaseResponse response = supportClient.createCase(caseRequest);
        return response.caseId();

    } catch (SupportException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
    return "";
}

public static String displaySevLevels(SupportClient supportClient) {
    try {
        DescribeSeverityLevelsRequest severityLevelsRequest =
DescribeSeverityLevelsRequest.builder()
            .language("en")
            .build();

```



```

        DescribeSeverityLevelsResponse response =
supportClient.describeSeverityLevels(severityLevelsRequest);
        List<SeverityLevel> severityLevels = response.severityLevels();
        String levelName = null;
        for (SeverityLevel sevLevel : severityLevels) {
            System.out.println("The severity level name is: " +
sevLevel.name());
            if (sevLevel.name().compareTo("High") == 0)
                levelName = sevLevel.name();
        }
        return levelName;

    } catch (SupportException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
    return "";
}

// Return a List that contains a Service name and Category name.
public static List<String> displayServices(SupportClient supportClient) {
    try {
        DescribeServicesRequest servicesRequest =
DescribeServicesRequest.builder()
            .language("en")
            .build();

        DescribeServicesResponse response =
supportClient.describeServices(servicesRequest);
        String serviceCode = null;
        String catName = null;
        List<String> sevCatList = new ArrayList<>();
        List<Service> services = response.services();

        System.out.println("Get the first 10 services");
        int index = 1;
        for (Service service : services) {
            if (index == 11)
                break;

            System.out.println("The Service name is: " + service.name());
            if (service.name().compareTo("Account") == 0)
                serviceCode = service.code();
        }
    }
}

```

```
        // Get the Categories for this service.
        List<Category> categories = service.categories();
        for (Category cat : categories) {
            System.out.println("The category name is: " + cat.name());
            if (cat.name().compareTo("Security") == 0)
                catName = cat.name();
        }
        index++;
    }

    // Push the two values to the list.
    sevCatList.add(serviceCode);
    sevCatList.add(catName);
    return sevCatList;

} catch (SupportException e) {
    System.out.println(e.getLocalizedMessage());
    System.exit(1);
}
return null;
}
}
```

- For API details, see the following topics in *AWS SDK for Java 2.x API Reference*.
 - [AddAttachmentsToSet](#)
 - [AddCommunicationToCase](#)
 - [CreateCase](#)
 - [DescribeAttachment](#)
 - [DescribeCases](#)
 - [DescribeCommunications](#)
 - [DescribeServices](#)
 - [DescribeSeverityLevels](#)
 - [ResolveCase](#)

JavaScript

SDK for JavaScript (v3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

Run an interactive scenario in the terminal.

```
import {
  AddAttachmentsToSetCommand,
  AddCommunicationToCaseCommand,
  CreateCaseCommand,
  DescribeAttachmentCommand,
  DescribeCasesCommand,
  DescribeCommunicationsCommand,
  DescribeServicesCommand,
  DescribeSeverityLevelsCommand,
  ResolveCaseCommand,
  SupportClient,
} from "@aws-sdk/client-support";
import * as inquirer from "@inquirer/prompts";
import { retry } from "@aws-doc-sdk-examples/lib/utils/util-timers.js";

const wrapText = (text, char = "=") => {
  const rule = char.repeat(80);
  return `${rule}\n    ${text}\n${rule}\n`;
};

const client = new SupportClient({ region: "us-east-1" });

// Verify that the account has a Support plan.
export const verifyAccount = async () => {
  const command = new DescribeServicesCommand({});

  try {
    await client.send(command);
  } catch (err) {
    if (err.name === "SubscriptionRequiredException") {
      throw new Error(
```

```

        "You must be subscribed to the AWS Support plan to use this feature.",
    );
    }
    throw err;
  }
};

/**
 * Select a service from the list returned from DescribeServices.
 */
export const getService = async () => {
  const { services } = await client.send(new DescribeServicesCommand({}));
  const selectedService = await inquirer.select({
    message:
      "Select a service. Your support case will be created for this service. The
      list of services is truncated for readability.",
    choices: services.slice(0, 10).map((s) => ({ name: s.name, value: s })),
  });
  return selectedService;
};

/**
 * @param {{ categories: import('@aws-sdk/client-support').Category[] }} service
 */
export const getCategory = async (service) => {
  const selectedCategory = await inquirer.select({
    message: "Select a category.",
    choices: service.categories.map((c) => ({ name: c.name, value: c })),
  });
  return selectedCategory;
};

// Get the available severity levels for the account.
export const getSeverityLevel = async () => {
  const command = new DescribeSeverityLevelsCommand({});
  const { severityLevels } = await client.send(command);
  const selectedSeverityLevel = await inquirer.select({
    message: "Select a severity level.",
    choices: severityLevels.map((s) => ({ name: s.name, value: s })),
  });
  return selectedSeverityLevel;
};

/**

```

```
* Create a new support case
* @param {{
*   selectedService: import('@aws-sdk/client-support').Service
*   selectedCategory: import('@aws-sdk/client-support').Category
*   selectedSeverityLevel: import('@aws-sdk/client-support').SeverityLevel
* }} selections
* @returns
*/
export const createCase = async ({
  selectedService,
  selectedCategory,
  selectedSeverityLevel,
}) => {
  const command = new CreateCaseCommand({
    subject: "IGNORE: Test case",
    communicationBody: "This is a test. Please ignore.",
    serviceCode: selectedService.code,
    categoryCode: selectedCategory.code,
    severityCode: selectedSeverityLevel.code,
  });
  const { caseId } = await client.send(command);
  return caseId;
};

// Get a list of open support cases created today.
export const getTodaysOpenCases = async () => {
  const d = new Date();
  const startOfToday = new Date(d.getFullYear(), d.getMonth(), d.getDate());
  const command = new DescribeCasesCommand({
    includeCommunications: false,
    afterTime: startOfToday.toISOString(),
  });

  const { cases } = await client.send(command);

  if (cases.length === 0) {
    throw new Error(
      "Unexpected number of cases. Expected more than 0 open cases.",
    );
  }
  return cases;
};

// Create an attachment set.
```

```

export const createAttachmentSet = async () => {
  const command = new AddAttachmentsToSetCommand({
    attachments: [
      {
        fileName: "example.txt",
        data: new TextEncoder().encode("some example text"),
      },
    ],
  });
  const { attachmentSetId } = await client.send(command);
  return attachmentSetId;
};

export const linkAttachmentSetToCase = async (attachmentSetId, caseId) => {
  const command = new AddCommunicationToCaseCommand({
    attachmentSetId,
    caseId,
    communicationBody: "Adding attachment set to case.",
  });
  await client.send(command);
};

// Get all communications for a support case.
export const getCommunications = async (caseId) => {
  const command = new DescribeCommunicationsCommand({
    caseId,
  });
  const { communications } = await client.send(command);
  return communications;
};

/**
 * @param {import('@aws-sdk/client-support').Communication[]} communications
 */
export const getFirstAttachment = (communications) => {
  const firstCommWithAttachment = communications.find(
    (c) => c.attachmentSet.length > 0,
  );
  return firstCommWithAttachment?.attachmentSet[0].attachmentId;
};

// Get an attachment.
export const getAttachment = async (attachmentId) => {
  const command = new DescribeAttachmentCommand({

```

```

        attachmentId,
    });
    const { attachment } = await client.send(command);
    return attachment;
};

// Resolve the case matching the given case ID.
export const resolveCase = async (caseId) => {
    const shouldResolve = await inquirer.confirm({
        message: `Do you want to resolve ${caseId}?`,
    });

    if (shouldResolve) {
        const command = new ResolveCaseCommand({
            caseId: caseId,
        });

        await client.send(command);
        return true;
    }
    return false;
};

/**
 * Find a specific case in the list of provided cases by case ID.
 * If the case is not found, and the results are paginated, continue
 * paging through the results.
 * @param {{
 *   caseId: string,
 *   cases: import('@aws-sdk/client-support').CaseDetails[]
 *   nextToken: string
 * }} options
 * @returns
 */
export const findCase = async ({ caseId, cases, nextToken }) => {
    const foundCase = cases.find((c) => c.caseId === caseId);

    if (foundCase) {
        return foundCase;
    }

    if (nextToken) {
        const response = await client.send(
            new DescribeCasesCommand({

```

```
        nextToken,
        includeResolvedCases: true,
      })),
    );
    return findCase({
      caseId,
      cases: response.cases,
      nextToken: response.nextToken,
    });
  }

  throw new Error(`${caseId} not found.`);
};

// Get all cases created today.
export const getTodayResolvedCases = async (caseIdToWaitFor) => {
  const d = new Date("2023-01-18");
  const startOfToday = new Date(d.getFullYear(), d.getMonth(), d.getDate());
  const command = new DescribeCasesCommand({
    includeCommunications: false,
    afterTime: startOfToday.toISOString(),
    includeResolvedCases: true,
  });
  const { cases, nextToken } = await client.send(command);
  await findCase({ cases, caseId: caseIdToWaitFor, nextToken });
  return cases.filter((c) => c.status === "resolved");
};

const main = async () => {
  let caseId;
  try {
    console.log(wrapText("Welcome to the AWS Support basic usage scenario."));

    // Verify that the account is subscribed to support.
    await verifyAccount();

    // Provided a truncated list of services and prompt the user to select one.
    const selectedService = await getService();

    // Provided the categories for the selected service and prompt the user to
    select one.
    const selectedCategory = await getCategory(selectedService);
```



```
// Provide the severity available severity levels for the account and prompt
the user to select one.
const selectedSeverityLevel = await getSeverityLevel();

// Create a support case.
console.log("\nCreating a support case.");
caseId = await createCase({
  selectedService,
  selectedCategory,
  selectedSeverityLevel,
});
console.log(`Support case created: ${caseId}`);

// Display a list of open support cases created today.
const todaysOpenCases = await retry(
  { intervalInMs: 1000, maxRetries: 15 },
  getTodaysOpenCases,
);
console.log(
  `\nOpen support cases created today: ${todaysOpenCases.length}`,
);
console.log(todaysOpenCases.map((c) => `${c.caseId}`).join("\n"));

// Create an attachment set.
console.log("\nCreating an attachment set.");
const attachmentSetId = await createAttachmentSet();
console.log(`Attachment set created: ${attachmentSetId}`);

// Add the attachment set to the support case.
console.log(`\nAdding attachment set to ${caseId}`);
await linkAttachmentSetToCase(attachmentSetId, caseId);
console.log(`Attachment set added to ${caseId}`);

// List the communications for a support case.
console.log(`\nListing communications for ${caseId}`);
const communications = await getCommunications(caseId);
console.log(
  communications
    .map(
      (c) =>
        `Communication created on ${c.timeCreated}. Has
        ${c.attachmentSet.length} attachments.`
    )
    .join("\n"),
);
```

```

    );

    // Describe the first attachment.
    console.log(`\nDescribing attachment ${attachmentSetId}`);
    const attachmentId = getFirstAttachment(communications);
    const attachment = await getAttachment(attachmentId);
    console.log(
      `Attachment is the file '${
        attachment.fileName
      }' with data: \n${new TextDecoder().decode(attachment.data)}`,
    );

    // Confirm that the support case should be resolved.
    const isResolved = await resolveCase(caseId);
    if (isResolved) {
      // List the resolved cases and include the one previously created.
      // Resolved cases can take a while to appear.
      console.log(
        "\nWaiting for case status to be marked as resolved. This can take some
time.",
      );
      const resolvedCases = await retry(
        { intervalInMs: 20000, maxRetries: 15 },
        () => getTodayResolvedCases(caseId),
      );
      console.log("Resolved cases:");
      console.log(resolvedCases.map((c) => c.caseId).join("\n"));
    }
  } catch (err) {
    console.error(err);
  }
};

```

- For API details, see the following topics in *AWS SDK for JavaScript API Reference*.
 - [AddAttachmentsToSet](#)
 - [AddCommunicationToCase](#)
 - [CreateCase](#)
 - [DescribeAttachment](#)
 - [DescribeCases](#)
 - [DescribeCommunications](#)

- [DescribeServices](#)
- [DescribeSeverityLevels](#)
- [ResolveCase](#)

Kotlin

SDK for Kotlin

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
/**
```

```
Before running this Kotlin code example, set up your development environment, including your credentials.
```

```
For more information, see the following documentation topic:
```

```
https://docs.aws.amazon.com/sdk-for-kotlin/latest/developer-guide/setup.html
```

```
In addition, you must have the AWS Business Support Plan to use the AWS Support Java API. For more information, see:
```

```
https://aws.amazon.com/premiumsupport/plans/
```

```
This Kotlin example performs the following tasks:
```

1. Gets and displays available services.
2. Gets and displays severity levels.
3. Creates a support case by using the selected service, category, and severity level.
4. Gets a list of open cases for the current day.
5. Creates an attachment set with a generated file.
6. Adds a communication with the attachment to the support case.
7. Lists the communications of the support case.
8. Describes the attachment set included with the communication.
9. Resolves the support case.
10. Gets a list of resolved cases for the current day.

```
*/
```

```
suspend fun main(args: Array<String>) {
    val usage = """
    Usage:
        <fileAttachment>
    Where:
        fileAttachment - The file can be a simple saved .txt file to use as an
    email attachment.
    """

    if (args.size != 1) {
        println(usage)
        exitProcess(0)
    }

    val fileAttachment = args[0]
    println("***** Welcome to the AWS Support case example scenario.")
    println("***** Step 1. Get and display available services.")
    val sevCatList = displayServices()

    println("***** Step 2. Get and display Support severity levels.")
    val sevLevel = displaySevLevels()

    println("***** Step 3. Create a support case using the selected service,
category, and severity level.")
    val caseIdVal = createSupportCase(sevCatList, sevLevel)
    if (caseIdVal != null) {
        println("Support case $caseIdVal was successfully created!")
    } else {
        println("A support case was not successfully created!")
        exitProcess(1)
    }

    println("***** Step 4. Get open support cases.")
    getOpenCase()

    println("***** Step 5. Create an attachment set with a generated file to add
to the case.")
    val attachmentSetId = addAttachment(fileAttachment)
    println("The Attachment Set id value is $attachmentSetId")

    println("***** Step 6. Add communication with the attachment to the support
case.")
    addAttachSupportCase(caseIdVal, attachmentSetId)
```

```

println("***** Step 7. List the communications of the support case.")
val attachId = listCommunications(caseIdVal)
println("The Attachment id value is $attachId")

println("***** Step 8. Describe the attachment set included with the
communication.")
describeAttachment(attachId)

println("***** Step 9. Resolve the support case.")
resolveSupportCase(caseIdVal)

println("***** Step 10. Get a list of resolved cases for the current day.")
getResolvedCase()
println("***** This Scenario has successfully completed")
}

suspend fun getResolvedCase() {
    // Specify the start and end time.
    val now = Instant.now()
    LocalDate.now()
    val yesterday = now.minus(1, ChronoUnit.DAYS)
    val describeCasesRequest =
        DescribeCasesRequest {
            maxResults = 30
            afterTime = yesterday.toString()
            beforeTime = now.toString()
            includeResolvedCases = true
        }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response = supportClient.describeCases(describeCasesRequest)
        response.cases?.forEach { sinCase ->
            println("The case status is ${sinCase.status}")
            println("The case Id is ${sinCase.caseId}")
            println("The case subject is ${sinCase.subject}")
        }
    }
}

suspend fun resolveSupportCase(caseIdVal: String) {
    val caseRequest =
        ResolveCaseRequest {
            caseId = caseIdVal
        }
}

```

```

        SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
            val response = supportClient.resolveCase(caseRequest)
            println("The status of case $caseIdVal is ${response.finalCaseStatus}")
        }
    }

suspend fun describeAttachment(attachId: String?) {
    val attachmentRequest =
        DescribeAttachmentRequest {
            attachmentId = attachId
        }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response = supportClient.describeAttachment(attachmentRequest)
        println("The name of the file is ${response.attachment?.fileName}")
    }
}

suspend fun listCommunications(caseIdVal: String?): String? {
    val communicationsRequest =
        DescribeCommunicationsRequest {
            caseId = caseIdVal
            maxResults = 10
        }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response =
            supportClient.describeCommunications(communicationsRequest)
        response.communications?.forEach { comm ->
            println("the body is: " + comm.body)
            comm.attachmentSet?.forEach { detail ->
                return detail.attachmentId
            }
        }
    }
    return ""
}

suspend fun addAttachSupportCase(
    caseIdVal: String?,
    attachmentSetIdVal: String?,
) {
    val caseRequest =
        AddCommunicationToCaseRequest {

```

```

        caseId = caseIdVal
        attachmentSetId = attachmentSetIdVal
        communicationBody = "Please refer to attachment for details."
    }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response = supportClient.addCommunicationToCase(caseRequest)
        if (response.result) {
            println("You have successfully added a communication to an AWS
Support case")
        } else {
            println("There was an error adding the communication to an AWS
Support case")
        }
    }
}

suspend fun addAttachment(fileAttachment: String): String? {
    val myFile = File(fileAttachment)
    val sourceBytes = (File(fileAttachment).readBytes())
    val attachmentVal =
        Attachment {
            fileName = myFile.name
            data = sourceBytes
        }

    val setRequest =
        AddAttachmentsToSetRequest {
            attachments = listOf(attachmentVal)
        }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response = supportClient.addAttachmentsToSet(setRequest)
        return response.attachmentSetId
    }
}

suspend fun getOpenCase() {
    // Specify the start and end time.
    val now = Instant.now()
    LocalDate.now()
    val yesterday = now.minus(1, ChronoUnit.DAYS)
    val describeCasesRequest =
        DescribeCasesRequest {

```

```

        maxResults = 20
        afterTime = yesterday.toString()
        beforeTime = now.toString()
    }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response = supportClient.describeCases(describeCasesRequest)
        response.cases?.forEach { sinCase ->
            println("The case status is ${sinCase.status}")
            println("The case Id is ${sinCase.caseId}")
            println("The case subject is ${sinCase.subject}")
        }
    }
}

suspend fun createSupportCase(
    sevCatListVal: List<String>,
    sevLevelVal: String,
): String? {
    val serCode = sevCatListVal[0]
    val caseCategory = sevCatListVal[1]
    val caseRequest =
        CreateCaseRequest {
            categoryCode = caseCategory.lowercase(Locale.getDefault())
            serviceCode = serCode.lowercase(Locale.getDefault())
            severityCode = sevLevelVal.lowercase(Locale.getDefault())
            communicationBody = "Test issue with
${serCode.lowercase(Locale.getDefault())}"
            subject = "Test case, please ignore"
            language = "en"
            issueType = "technical"
        }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response = supportClient.createCase(caseRequest)
        return response.caseId
    }
}

suspend fun displaySevLevels(): String {
    var levelName = ""
    val severityLevelsRequest =
        DescribeSeverityLevelsRequest {
            language = "en"

```



```

    }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response =
        supportClient.describeSeverityLevels(severityLevelsRequest)
        response.severityLevels?.forEach { sevLevel ->
            println("The severity level name is: ${sevLevel.name}")
            if (sevLevel.name == "High") {
                levelName = sevLevel.name!!
            }
        }
        return levelName
    }
}

// Return a List that contains a Service name and Category name.
suspend fun displayServices(): List<String> {
    var serviceCode = ""
    var catName = ""
    val sevCatList = mutableListOf<String>()
    val servicesRequest =
        DescribeServicesRequest {
            language = "en"
        }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response = supportClient.describeServices(servicesRequest)
        println("Get the first 10 services")
        var index = 1

        response.services?.forEach { service ->
            if (index == 11) {
                return@forEach
            }

            println("The Service name is ${service.name}")
            if (service.name == "Account") {
                serviceCode = service.code.toString()
            }

            // Get the categories for this service.
            service.categories?.forEach { cat ->
                println("The category name is ${cat.name}")
                if (cat.name == "Security") {

```

```
        catName = cat.name!!
    }
}
index++
}
}

// Push the two values to the list.
serviceCode.let { sevCatList.add(it) }
catName.let { sevCatList.add(it) }
return sevCatList
}
```

- For API details, see the following topics in *AWS SDK for Kotlin API reference*.

- [AddAttachmentsToSet](#)
- [AddCommunicationToCase](#)
- [CreateCase](#)
- [DescribeAttachment](#)
- [DescribeCases](#)
- [DescribeCommunications](#)
- [DescribeServices](#)
- [DescribeSeverityLevels](#)
- [ResolveCase](#)

Python

SDK for Python (Boto3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

Run an interactive scenario at a command prompt.

```
class SupportCasesScenario:
```

```

    """Runs an interactive scenario that shows how to get started using AWS
    Support."""

    def __init__(self, support_wrapper):
        """
        :param support_wrapper: An object that wraps AWS Support actions.
        """
        self.support_wrapper = support_wrapper

    def display_and_select_service(self):
        """
        Lists support services and prompts the user to select one.

        :return: The support service selected by the user.
        """
        print("-" * 88)
        services_list = self.support_wrapper.describe_services("en")
        print(f"AWS Support client returned {len(services_list)} services.")
        print("Displaying first 10 services:")

        service_choices = [svc["name"] for svc in services_list[:10]]
        selected_index = q.choose(
            "Select an example support service by entering a number from the
preceding list:",
            service_choices,
        )
        selected_service = services_list[selected_index]
        print("-" * 88)
        return selected_service

    def display_and_select_category(self, service):
        """
        Lists categories for a support service and prompts the user to select
one.

        :param service: The service of the categories.
        :return: The selected category.
        """
        print("-" * 88)
        print(
            f"Available support categories for Service {service['name']}
{len(service['categories'])}:"
        )

```

```

        categories_choices = [category["name"] for category in
service["categories"]]
        selected_index = q.choose(
            "Select an example support category by entering a number from the
preceding list:",
            categories_choices,
        )
        selected_category = service["categories"][selected_index]
        print("-" * 88)
        return selected_category

def display_and_select_severity(self):
    """
    Lists available severity levels and prompts the user to select one.

    :return: The selected severity level.
    """
    print("-" * 88)
    severity_levels_list =
self.support_wrapper.describe_severity_levels("en")
    print(f"Available severity levels:")
    severity_choices = [level["name"] for level in severity_levels_list]
    selected_index = q.choose(
        "Select an example severity level by entering a number from the
preceding list:",
        severity_choices,
    )
    selected_severity = severity_levels_list[selected_index]
    print("-" * 88)
    return selected_severity

def create_example_case(self, service, category, severity_level):
    """
    Creates an example support case with the user's selections.

    :param service: The service for the new case.
    :param category: The category for the new case.
    :param severity_level: The severity level for the new case.
    :return: The caseId of the new support case.
    """
    print("-" * 88)
    print(f"Creating new case for service {service['name']}.")
    case_id = self.support_wrapper.create_case(service, category,
severity_level)

```

```

        print(f"\tNew case created with ID {case_id}.")
        print("-" * 88)
        return case_id

def list_open_cases(self):
    """
    List the open cases for the current day.
    """
    print("-" * 88)
    print("Let's list the open cases for the current day.")
    start_time = str(datetime.utcnow().date())
    end_time = str(datetime.utcnow().date() + timedelta(days=1))
    open_cases = self.support_wrapper.describe_cases(start_time, end_time,
False)
    for case in open_cases:
        print(f"\tCase: {case['caseId']}: status {case['status']}")
    print("-" * 88)

def create_attachment_set(self):
    """
    Create an attachment set with a sample file.

    :return: The attachment set ID of the new attachment set.
    """
    print("-" * 88)
    print("Creating attachment set with a sample file.")
    attachment_set_id = self.support_wrapper.add_attachment_to_set()
    print(f"\tNew attachment set created with ID {attachment_set_id}.")
    print("-" * 88)
    return attachment_set_id

def add_communication(self, case_id, attachment_set_id):
    """
    Add a communication with an attachment set to the case.

    :param case_id: The ID of the case for the communication.
    :param attachment_set_id: The ID of the attachment set to
    add to the communication.
    """
    print("-" * 88)
    print(f"Adding a communication and attachment set to the case.")
    self.support_wrapper.add_communication_to_case(attachment_set_id,
case_id)
    print(

```

```

        f"Added a communication and attachment set {attachment_set_id} to the
        case {case_id}."
    )
    print("-" * 88)

def list_communications(self, case_id):
    """
    List the communications associated with a case.

    :param case_id: The ID of the case.
    :return: The attachment ID of an attachment.
    """
    print("-" * 88)
    print("Let's list the communications for our case.")
    attachment_id = ""
    communications =
self.support_wrapper.describe_all_case_communications(case_id)
    for communication in communications:
        print(
            f"\tCommunication created on {communication['timeCreated']} "
            f"has {len(communication['attachmentSet'])} attachments."
        )
        if len(communication["attachmentSet"]) > 0:
            attachment_id = communication["attachmentSet"][0]["attachmentId"]
    print("-" * 88)
    return attachment_id

def describe_case_attachment(self, attachment_id):
    """
    Describe an attachment associated with a case.

    :param attachment_id: The ID of the attachment.
    """
    print("-" * 88)
    print("Let's list the communications for our case.")
    attached_file = self.support_wrapper.describe_attachment(attachment_id)
    print(f"\tAttachment includes file {attached_file}.")
    print("-" * 88)

def resolve_case(self, case_id):
    """
    Shows how to resolve an AWS Support case by its ID.

    :param case_id: The ID of the case to resolve.

```

```

        """
        print("-" * 88)
        print(f"Resolving case with ID {case_id}.")
        case_status = self.support_wrapper.resolve_case(case_id)
        print(f"\tFinal case status is {case_status}.")
        print("-" * 88)

    def list_resolved_cases(self):
        """
        List the resolved cases for the current day.
        """
        print("-" * 88)
        print("Let's list the resolved cases for the current day.")
        start_time = str(datetime.utcnow().date())
        end_time = str(datetime.utcnow().date() + timedelta(days=1))
        resolved_cases = self.support_wrapper.describe_cases(start_time,
end_time, True)
        for case in resolved_cases:
            print(f"\tCase: {case['caseId']}: status {case['status']}.")
        print("-" * 88)

    def run_scenario(self):
        logging.basicConfig(level=logging.INFO, format="%(levelname)s:
%(message)s")

        print("-" * 88)
        print("Welcome to the AWS Support get started with support cases demo.")
        print("-" * 88)

        selected_service = self.display_and_select_service()
        selected_category = self.display_and_select_category(selected_service)
        selected_severity = self.display_and_select_severity()
        new_case_id = self.create_example_case(
            selected_service, selected_category, selected_severity
        )
        wait(10)
        self.list_open_cases()
        new_attachment_set_id = self.create_attachment_set()
        self.add_communication(new_case_id, new_attachment_set_id)
        new_attachment_id = self.list_communications(new_case_id)
        self.describe_case_attachment(new_attachment_id)
        self.resolve_case(new_case_id)
        wait(10)
        self.list_resolved_cases()

```

```

        print("\nThanks for watching!")
        print("-" * 88)

if __name__ == "__main__":
    try:
        scenario = SupportCasesScenario(SupportWrapper.from_client())
        scenario.run_scenario()
    except Exception:
        logging.exception("Something went wrong with the demo.")

```

Define a class that wraps support client actions.

```

class SupportWrapper:
    """Encapsulates Support actions."""

    def __init__(self, support_client):
        """
        :param support_client: A Boto3 Support client.
        """
        self.support_client = support_client

    @classmethod
    def from_client(cls):
        """
        Instantiates this class from a Boto3 client.
        """
        support_client = boto3.client("support")
        return cls(support_client)

    def describe_services(self, language):
        """
        Get the descriptions of AWS services available for support for a
        language.

        :param language: The language for support services.
        Currently, only "en" (English) and "ja" (Japanese) are supported.
        :return: The list of AWS service descriptions.
        """
        try:

```



```

        response = self.support_client.describe_services(language=language)
        services = response["services"]
    except ClientError as err:
        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                "examples."
            )
        else:
            logger.error(
                "Couldn't get Support services for language %s. Here's why:
%s: %s",
                language,
                err.response["Error"]["Code"],
                err.response["Error"]["Message"],
            )
            raise
    else:
        return services

def describe_severity_levels(self, language):
    """
    Get the descriptions of available severity levels for support cases for a
    language.

    :param language: The language for support severity levels.
    Currently, only "en" (English) and "ja" (Japanese) are supported.
    :return: The list of severity levels.
    """
    try:
        response =
self.support_client.describe_severity_levels(language=language)
        severity_levels = response["severityLevels"]
    except ClientError as err:
        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "

```

```

        "examples."
    )
else:
    logger.error(
        "Couldn't get severity levels for language %s. Here's why:
%s: %s",
        language,
        err.response["Error"]["Code"],
        err.response["Error"]["Message"],
    )
    raise
else:
    return severity_levels

def create_case(self, service, category, severity):
    """
    Create a new support case.

    :param service: The service to use for the new case.
    :param category: The category to use for the new case.
    :param severity: The severity to use for the new case.
    :return: The caseId of the new case.
    """
    try:
        response = self.support_client.create_case(
            subject="Example case for testing, ignore.",
            serviceCode=service["code"],
            severityCode=severity["code"],
            categoryCode=category["code"],
            communicationBody="Example support case body.",
            language="en",
            issueType="customer-service",
        )
        case_id = response["caseId"]
    except ClientError as err:
        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                "examples."
            )

```

```

        else:
            logger.error(
                "Couldn't create case. Here's why: %s: %s",
                err.response["Error"]["Code"],
                err.response["Error"]["Message"],
            )
            raise
    else:
        return case_id

def add_attachment_to_set(self):
    """
    Add an attachment to a set, or create a new attachment set if one does
    not exist.

    :return: The attachment set ID.
    """
    try:
        response = self.support_client.add_attachments_to_set(
            attachments=[
                {
                    "fileName": "attachment_file.txt",
                    "data": b"This is a sample file for attachment to a
support case.",
                }
            ]
        )
        new_set_id = response["attachmentSetId"]
    except ClientError as err:
        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                "examples."
            )
        else:
            logger.error(
                "Couldn't add attachment. Here's why: %s: %s",
                err.response["Error"]["Code"],
                err.response["Error"]["Message"],
            )

```

```

        raise
    else:
        return new_set_id

def add_communication_to_case(self, attachment_set_id, case_id):
    """
    Add a communication and an attachment set to a case.

    :param attachment_set_id: The ID of an existing attachment set.
    :param case_id: The ID of the case.
    """
    try:
        self.support_client.add_communication_to_case(
            caseId=case_id,
            communicationBody="This is an example communication added to a
support case.",
            attachmentSetId=attachment_set_id,
        )
    except ClientError as err:
        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                "examples."
            )
        else:
            logger.error(
                "Couldn't add communication. Here's why: %s: %s",
                err.response["Error"]["Code"],
                err.response["Error"]["Message"],
            )
            raise

def describe_all_case_communications(self, case_id):
    """
    Describe all the communications for a case using a paginator.

    :param case_id: The ID of the case.
    :return: The communications for the case.
    """

```

```

        try:
            communications = []
            paginator =
self.support_client.get_paginator("describe_communications")
            for page in paginator.paginate(caseId=case_id):
                communications += page["communications"]
        except ClientError as err:
            if err.response["Error"]["Code"] == "SubscriptionRequiredException":
                logger.info(
                    "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                    "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                    "examples."
                )
            else:
                logger.error(
                    "Couldn't describe communications. Here's why: %s: %s",
                    err.response["Error"]["Code"],
                    err.response["Error"]["Message"],
                )
                raise
        else:
            return communications

def describe_attachment(self, attachment_id):
    """
    Get information about an attachment by its attachmentID.

    :param attachment_id: The ID of the attachment.
    :return: The name of the attached file.
    """
    try:
        response = self.support_client.describe_attachment(
            attachmentId=attachment_id
        )
        attached_file = response["attachment"]["fileName"]
    except ClientError as err:
        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "

```

```

        "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
        "examples."
    )
    else:
        logger.error(
            "Couldn't get attachment description. Here's why: %s: %s",
            err.response["Error"]["Code"],
            err.response["Error"]["Message"],
        )
        raise
    else:
        return attached_file

def resolve_case(self, case_id):
    """
    Resolve a support case by its caseId.

    :param case_id: The ID of the case to resolve.
    :return: The final status of the case.
    """
    try:
        response = self.support_client.resolve_case(caseId=case_id)
        final_status = response["finalCaseStatus"]
    except ClientError as err:
        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                "examples."
            )
        else:
            logger.error(
                "Couldn't resolve case. Here's why: %s: %s",
                err.response["Error"]["Code"],
                err.response["Error"]["Message"],
            )
            raise
    else:
        return final_status

```

```

def describe_cases(self, after_time, before_time, resolved):
    """
    Describe support cases over a period of time, optionally filtering
    by status.

    :param after_time: The start time to include for cases.
    :param before_time: The end time to include for cases.
    :param resolved: True to include resolved cases in the results,
        otherwise results are open cases.
    :return: The final status of the case.
    """
    try:
        cases = []
        paginator = self.support_client.get_paginator("describe_cases")
        for page in paginator.paginate(
            afterTime=after_time,
            beforeTime=before_time,
            includeResolvedCases=resolved,
            language="en",
        ):
            cases += page["cases"]
    except ClientError as err:
        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                "examples."
            )
        else:
            logger.error(
                "Couldn't describe cases. Here's why: %s: %s",
                err.response["Error"]["Code"],
                err.response["Error"]["Message"],
            )
            raise
    else:
        if resolved:
            cases = filter(lambda case: case["status"] == "resolved", cases)
        return cases

```

- For API details, see the following topics in *AWS SDK for Python (Boto3) API Reference*.
 - [AddAttachmentsToSet](#)
 - [AddCommunicationToCase](#)
 - [CreateCase](#)
 - [DescribeAttachment](#)
 - [DescribeCases](#)
 - [DescribeCommunications](#)
 - [DescribeServices](#)
 - [DescribeSeverityLevels](#)
 - [ResolveCase](#)

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Actions for Support using AWS SDKs

The following code examples demonstrate how to perform individual Support actions with AWS SDKs. Each example includes a link to GitHub, where you can find instructions for setting up and running the code.

These excerpts call the Support API and are code excerpts from larger programs that must be run in context. You can see actions in context in [Scenarios for Support using AWS SDKs](#).

The following examples include only the most commonly used actions. For a complete list, see the [AWS Support API Reference](#).

Examples

- [Use AddAttachmentsToSet with an AWS SDK or CLI](#)
- [Use AddCommunicationToCase with an AWS SDK or CLI](#)
- [Use CreateCase with an AWS SDK or CLI](#)
- [Use DescribeAttachment with an AWS SDK or CLI](#)

- [Use DescribeCases with an AWS SDK or CLI](#)
- [Use DescribeCommunications with an AWS SDK or CLI](#)
- [Use DescribeServices with an AWS SDK or CLI](#)
- [Use DescribeSeverityLevels with an AWS SDK or CLI](#)
- [Use DescribeTrustedAdvisorCheckRefreshStatuses with a CLI](#)
- [Use DescribeTrustedAdvisorCheckResult with a CLI](#)
- [Use DescribeTrustedAdvisorCheckSummaries with a CLI](#)
- [Use DescribeTrustedAdvisorChecks with a CLI](#)
- [Use RefreshTrustedAdvisorCheck with a CLI](#)
- [Use ResolveCase with an AWS SDK or CLI](#)

Use AddAttachmentsToSet with an AWS SDK or CLI

The following code examples show how to use AddAttachmentsToSet.

Action examples are code excerpts from larger programs and must be run in context. You can see this action in context in the following code example:

- [Learn the basics](#)

.NET

SDK for .NET

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
/// <summary>
/// Add an attachment to a set, or create a new attachment set if one does
not exist.
/// </summary>
/// <param name="data">The data for the attachment.</param>
```

```

    /// <param name="fileName">The file name for the attachment.</param>
    /// <param name="attachmentSetId">Optional setId for the attachment. Creates
    a new attachment set if empty.</param>
    /// <returns>The setId of the attachment.</returns>
    public async Task<string> AddAttachmentToSet(MemoryStream data, string
    fileName, string? attachmentSetId = null)
    {
        var response = await _amazonSupport.AddAttachmentsToSetAsync(
            new AddAttachmentsToSetRequest
            {
                AttachmentSetId = attachmentSetId,
                Attachments = new List<Attachment>
                {
                    new Attachment
                    {
                        Data = data,
                        FileName = fileName
                    }
                }
            });
        return response.AttachmentSetId;
    }

```

- For API details, see [AddAttachmentsToSet](#) in *AWS SDK for .NET API Reference*.

CLI

AWS CLI

To add an attachment to a set

The following add-attachments-to-set example adds an image to a set that you can then specify for a support case in your AWS account.

```

aws support add-attachments-to-set \
    --attachment-set-id "as-2f5a6faa2a4a1e600-mu-nk5xQlBr70-
G1cUos5LZkd38K0AHZa9BMDVzNEXAMPLE" \
    --attachments fileName=troubleshoot-screenshot.png,data=base64-encoded-string

```

Output:

```
{
  "attachmentSetId": "as-2f5a6faa2a4a1e600-mu-nk5xQlBr70-
G1cUos5LZkd38K0AHZa9BMDVzNEXAMPLE",
  "expiryTime": "2020-05-14T17:04:40.790+0000"
}
```

For more information, see [Case management](#) in the *AWS Support User Guide*.

- For API details, see [AddAttachmentsToSet](#) in *AWS CLI Command Reference*.

Java

SDK for Java 2.x

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
public static String addAttachment(SupportClient supportClient, String
fileAttachment) {
    try {
        File myFile = new File(fileAttachment);
        InputStream sourceStream = new FileInputStream(myFile);
        SdkBytes sourceBytes = SdkBytes.fromInputStream(sourceStream);

        Attachment attachment = Attachment.builder()
            .fileName(myFile.getName())
            .data(sourceBytes)
            .build();

        AddAttachmentsToSetRequest setRequest =
        AddAttachmentsToSetRequest.builder()
            .attachments(attachment)
            .build();

        AddAttachmentsToSetResponse response =
        supportClient.addAttachmentsToSet(setRequest);
        return response.attachmentSetId();
    }
}
```

```
    } catch (SupportException | FileNotFoundException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
    return "";
}
```

- For API details, see [AddAttachmentsToSet](#) in *AWS SDK for Java 2.x API Reference*.

JavaScript

SDK for JavaScript (v3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
import { AddAttachmentsToSetCommand } from "@aws-sdk/client-support";

import { client } from "../libs/client.js";

export const main = async () => {
    try {
        // Create a new attachment set or add attachments to an existing set.
        // Provide an 'attachmentSetId' value to add attachments to an existing set.
        // Use AddCommunicationToCase or CreateCase to associate an attachment set
        with a support case.
        const response = await client.send(
            new AddAttachmentsToSetCommand({
                // You can add up to three attachments per set. The size limit is 5 MB
                per attachment.
                attachments: [
                    {
                        fileName: "example.txt",
                        data: new TextEncoder().encode("some example text"),
                    },
                ],
            })
        );
    }
```

```
// Use this ID in AddCommunicationToCase or CreateCase.
console.log(response.attachmentSetId);
return response;
} catch (err) {
  console.error(err);
}
};
```

- For API details, see [AddAttachmentsToSet](#) in *AWS SDK for JavaScript API Reference*.

Kotlin

SDK for Kotlin

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
suspend fun addAttachment(fileAttachment: String): String? {
    val myFile = File(fileAttachment)
    val sourceBytes = (File(fileAttachment).readBytes())
    val attachmentVal =
        Attachment {
            fileName = myFile.name
            data = sourceBytes
        }

    val setRequest =
        AddAttachmentsToSetRequest {
            attachments = listOf(attachmentVal)
        }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response = supportClient.addAttachmentsToSet(setRequest)
        return response.attachmentSetId
    }
}
```

- For API details, see [AddAttachmentsToSet](#) in *AWS SDK for Kotlin API reference*.

Python

SDK for Python (Boto3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
class SupportWrapper:
    """Encapsulates Support actions."""

    def __init__(self, support_client):
        """
        :param support_client: A Boto3 Support client.
        """
        self.support_client = support_client

    @classmethod
    def from_client(cls):
        """
        Instantiates this class from a Boto3 client.
        """
        support_client = boto3.client("support")
        return cls(support_client)

    def add_attachment_to_set(self):
        """
        Add an attachment to a set, or create a new attachment set if one does
        not exist.

        :return: The attachment set ID.
        """
        try:
            response = self.support_client.add_attachments_to_set(
                attachments=[
                    {
                        "fileName": "attachment_file.txt",
```

```

        "data": b"This is a sample file for attachment to a
support case.",
    }
]
)
new_set_id = response["attachmentSetId"]
except ClientError as err:
    if err.response["Error"]["Code"] == "SubscriptionRequiredException":
        logger.info(
            "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
            "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
            "examples."
        )
    else:
        logger.error(
            "Couldn't add attachment. Here's why: %s: %s",
            err.response["Error"]["Code"],
            err.response["Error"]["Message"],
        )
        raise
    else:
        return new_set_id

```

- For API details, see [AddAttachmentsToSet](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use AddCommunicationToCase with an AWS SDK or CLI

The following code examples show how to use AddCommunicationToCase.

Action examples are code excerpts from larger programs and must be run in context. You can see this action in context in the following code examples:

- [Learn the basics](#)
- [Getting started with technical support](#)

.NET

SDK for .NET

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
/// <summary>
/// Add communication to a case, including optional attachment set ID and CC
email addresses.
/// </summary>
/// <param name="caseId">Id for the support case.</param>
/// <param name="body">Body text of the communication.</param>
/// <param name="attachmentSetId">Optional Id for an attachment set.</param>
/// <param name="ccEmailAddresses">Optional list of CC email addresses.</
param>
/// <returns>True if successful.</returns>
public async Task<bool> AddCommunicationToCase(string caseId, string body,
    string? attachmentSetId = null, List<string>? ccEmailAddresses = null)
{
    var response = await _amazonSupport.AddCommunicationToCaseAsync(
        new AddCommunicationToCaseRequest()
        {
            CaseId = caseId,
            CommunicationBody = body,
            AttachmentSetId = attachmentSetId,
            CcEmailAddresses = ccEmailAddresses
        });
    return response.Result;
}
```

- For API details, see [AddCommunicationToCase](#) in *AWS SDK for .NET API Reference*.

CLI

AWS CLI

To add communication to a case

The following add-communication-to-case example adds communications to a support case in your AWS account.

```
aws support add-communication-to-case \  
  --case-id "case-12345678910-2013-c4c1d2bf33c5cf47" \  
  --communication-body "I'm attaching a set of images to this case." \  
  --cc-email-addresses "myemail@example.com" \  
  --attachment-set-id "as-2f5a6faa2a4a1e600-mu-nk5xQlBr70-  
G1cUos5LZkd38K0AHZa9BMDVzNEXAMPLE"
```

Output:

```
{  
  "result": true  
}
```

For more information, see [Case management](#) in the *AWS Support User Guide*.

- For API details, see [AddCommunicationToCase](#) in *AWS CLI Command Reference*.

Java

SDK for Java 2.x

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
public static void addAttachSupportCase(SupportClient supportClient, String  
caseId, String attachmentSetId) {  
    try {  
        AddCommunicationToCaseRequest caseRequest =  
        AddCommunicationToCaseRequest.builder()
```

```
        .caseId(caseId)
        .attachmentSetId(attachmentSetId)
        .communicationBody("Please refer to attachment for details.")
        .build();

        AddCommunicationToCaseResponse response =
supportClient.addCommunicationToCase(caseRequest);
        if (response.result())
            System.out.println("You have successfully added a communication
to an AWS Support case");
        else
            System.out.println("There was an error adding the communication
to an AWS Support case");

    } catch (SupportException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
}
```

- For API details, see [AddCommunicationToCase](#) in *AWS SDK for Java 2.x API Reference*.

JavaScript

SDK for JavaScript (v3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
import { AddCommunicationToCaseCommand } from "@aws-sdk/client-support";

import { client } from "../libs/client.js";

export const main = async () => {
    let attachmentSetId;

    try {
        // Add a communication to a case.
```

```
const response = await client.send(  
  new AddCommunicationToCaseCommand({  
    communicationBody: "Adding an attachment.",  
    // Set value to an existing support case id.  
    caseId: "CASE_ID",  
    // Optional. Set value to an existing attachment set id to add  
    attachments to the case.  
    attachmentSetId,  
  })),  
);  
console.log(response);  
return response;  
} catch (err) {  
  console.error(err);  
}  
};
```

- For API details, see [AddCommunicationToCase](#) in *AWS SDK for JavaScript API Reference*.

Kotlin

SDK for Kotlin

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
suspend fun addAttachSupportCase(  
  caseIdVal: String?,  
  attachmentSetIdVal: String?,  
) {  
  val caseRequest =  
    AddCommunicationToCaseRequest {  
      caseId = caseIdVal  
      attachmentSetId = attachmentSetIdVal  
      communicationBody = "Please refer to attachment for details."  
    }  
  
  SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
```

```
val response = supportClient.addCommunicationToCase(caseRequest)
if (response.result) {
    println("You have successfully added a communication to an AWS
Support case")
} else {
    println("There was an error adding the communication to an AWS
Support case")
}
}
```

- For API details, see [AddCommunicationToCase](#) in *AWS SDK for Kotlin API reference*.

PowerShell

Tools for PowerShell V4

Example 1: Adds the body of an email communication to the specified case.

```
Add-ASACommunicationToCase -CaseId "case-12345678910-2013-c4c1d2bf33c5cf47" -
CommunicationBody "Some text about the case"
```

Example 2: Adds the body of an email communication to the specified case plus one or more email addresses contained in the CC line of the email.

```
Add-ASACommunicationToCase -CaseId "case-12345678910-2013-c4c1d2bf33c5cf47" -
CcEmailAddress @"email1@address.com", "email2@address.com") -CommunicationBody
"Some text about the case"
```

- For API details, see [AddCommunicationToCase](#) in *AWS Tools for PowerShell Cmdlet Reference (V4)*.

Tools for PowerShell V5

Example 1: Adds the body of an email communication to the specified case.

```
Add-ASACommunicationToCase -CaseId "case-12345678910-2013-c4c1d2bf33c5cf47" -
CommunicationBody "Some text about the case"
```

Example 2: Adds the body of an email communication to the specified case plus one or more email addresses contained in the CC line of the email.

```
Add-ASACommunicationToCase -CaseId "case-12345678910-2013-c4c1d2bf33c5cf47" -
CcEmailAddress @"email1@address.com", "email2@address.com") -CommunicationBody
"Some text about the case"
```

- For API details, see [AddCommunicationToCase](#) in *AWS Tools for PowerShell Cmdlet Reference (V5)*.

Python

SDK for Python (Boto3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
class SupportWrapper:
    """Encapsulates Support actions."""

    def __init__(self, support_client):
        """
        :param support_client: A Boto3 Support client.
        """
        self.support_client = support_client

    @classmethod
    def from_client(cls):
        """
        Instantiates this class from a Boto3 client.
        """
        support_client = boto3.client("support")
        return cls(support_client)

    def add_communication_to_case(self, attachment_set_id, case_id):
        """
        Add a communication and an attachment set to a case.

        :param attachment_set_id: The ID of an existing attachment set.
        :param case_id: The ID of the case.
```

```

    """
    try:
        self.support_client.add_communication_to_case(
            caseId=case_id,
            communicationBody="This is an example communication added to a
support case.",
            attachmentSetId=attachment_set_id,
        )
    except ClientError as err:
        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                "examples."
            )
        else:
            logger.error(
                "Couldn't add communication. Here's why: %s: %s",
                err.response["Error"]["Code"],
                err.response["Error"]["Message"],
            )
            raise

```

- For API details, see [AddCommunicationToCase](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use CreateCase with an AWS SDK or CLI

The following code examples show how to use CreateCase.

Action examples are code excerpts from larger programs and must be run in context. You can see this action in context in the following code examples:

- [Learn the basics](#)

- [Getting started with technical support](#)

.NET

SDK for .NET

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
/// <summary>
/// Create a new support case.
/// </summary>
/// <param name="serviceCode">Service code for the new case.</param>
/// <param name="categoryCode">Category for the new case.</param>
/// <param name="severityCode">Severity code for the new case.</param>
/// <param name="subject">Subject of the new case.</param>
/// <param name="body">Body text of the new case.</param>
/// <param name="language">Optional language support for your case.
/// Currently Chinese ("zh"), English ("en"), Japanese ("ja") and Korean
/// ("ko") are supported.</param>
/// <param name="attachmentSetId">Optional Id for an attachment set for the
new case.</param>
/// <param name="issueType">Optional issue type for the new case. Options are
"customer-service" or "technical".</param>
/// <returns>The caseId of the new support case.</returns>
public async Task<string> CreateCase(string serviceCode, string categoryCode,
string severityCode, string subject,
string body, string language = "en", string? attachmentSetId = null,
string issueType = "customer-service")
{
    var response = await _amazonSupport.CreateCaseAsync(
        new CreateCaseRequest()
        {
            ServiceCode = serviceCode,
            CategoryCode = categoryCode,
            SeverityCode = severityCode,
            Subject = subject,
            Language = language,
```

```
        AttachmentSetId = attachmentSetId,  
        IssueType = issueType,  
        CommunicationBody = body  
    });  
    return response.CaseId;  
}
```

- For API details, see [CreateCase](#) in *AWS SDK for .NET API Reference*.

CLI

AWS CLI

To create a case

The following create-case example creates a support case for your AWS account.

```
aws support create-case \  
  --category-code "using-aws" \  
  --cc-email-addresses "myemail@example.com" \  
  --communication-body "I want to learn more about an AWS service." \  
  --issue-type "technical" \  
  --language "en" \  
  --service-code "general-info" \  
  --severity-code "low" \  
  --subject "Question about my account"
```

Output:

```
{  
  "caseId": "case-12345678910-2013-c4c1d2bf33c5cf47"  
}
```

For more information, see [Case management](#) in the *AWS Support User Guide*.

- For API details, see [CreateCase](#) in *AWS CLI Command Reference*.

Java

SDK for Java 2.x

 Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
public static String createSupportCase(SupportClient supportClient,
List<String> sevCatList, String sevLevel) {
    try {
        String serviceCode = sevCatList.get(0);
        String caseCat = sevCatList.get(1);
        CreateCaseRequest caseRequest = CreateCaseRequest.builder()
            .categoryCode(caseCat.toLowerCase())
            .serviceCode(serviceCode.toLowerCase())
            .severityCode(sevLevel.toLowerCase())
            .communicationBody("Test issue with " +
serviceCode.toLowerCase())
            .subject("Test case, please ignore")
            .language("en")
            .issueType("technical")
            .build();

        CreateCaseResponse response = supportClient.createCase(caseRequest);
        return response.caseId();

    } catch (SupportException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
    return "";
}
```

- For API details, see [CreateCase](#) in *AWS SDK for Java 2.x API Reference*.

JavaScript

SDK for JavaScript (v3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
import { CreateCaseCommand } from "@aws-sdk/client-support";

import { client } from "../libs/client.js";

export const main = async () => {
  try {
    // Create a new case and log the case id.
    // Important: This creates a real support case in your account.
    const response = await client.send(
      new CreateCaseCommand({
        // The subject line of the case.
        subject: "IGNORE: Test case",
        // Use DescribeServices to find available service codes for each service.
        serviceCode: "service-quicksight-end-user",
        // Use DescribeSecurityLevels to find available severity codes for your
        support plan.
        severityCode: "low",
        // Use DescribeServices to find available category codes for each
        service.
        categoryCode: "end-user-support",
        // The main description of the support case.
        communicationBody: "This is a test. Please ignore.",
      }),
    );
    console.log(response.caseId);
    return response;
  } catch (err) {
    console.error(err);
  }
};
```

- For API details, see [CreateCase](#) in *AWS SDK for JavaScript API Reference*.

Kotlin

SDK for Kotlin

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
suspend fun createSupportCase(
    sevCatListVal: List<String>,
    sevLevelVal: String,
): String? {
    val serCode = sevCatListVal[0]
    val caseCategory = sevCatListVal[1]
    val caseRequest =
        CreateCaseRequest {
            categoryCode = caseCategory.lowercase(Locale.getDefault())
            serviceCode = serCode.lowercase(Locale.getDefault())
            severityCode = sevLevelVal.lowercase(Locale.getDefault())
            communicationBody = "Test issue with
${serCode.lowercase(Locale.getDefault())}"
            subject = "Test case, please ignore"
            language = "en"
            issueType = "technical"
        }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response = supportClient.createCase(caseRequest)
        return response.caseId
    }
}
```

- For API details, see [CreateCase](#) in *AWS SDK for Kotlin API reference*.

PowerShell

Tools for PowerShell V4

Example 1: Creates a new case in the AWS Support Center. Values for the `-ServiceCode` and `-CategoryCode` parameters can be obtained using the `Get-ASAService` cmdlet. The value for the `-SeverityCode` parameter can be obtained using the `Get-ASASeverityLevel` cmdlet. The `-IssueType` parameter value can be either "customer-service" or "technical". If successful the AWS Support case number is output. By default the case will be handled in English, to use Japanese add the `-Language "ja"` parameter. The `-ServiceCode`, `-CategoryCode`, `-Subject` and `-CommunicationBody` parameters are mandatory.

```
New-ASACase -ServiceCode "amazon-cloudfront" -CategoryCode "APIs" -SeverityCode  
"low" -Subject "subject text" -CommunicationBody "description of the case" -  
CcEmailAddress @("email1@domain.com", "email2@domain.com") -IssueType "technical"
```

- For API details, see [CreateCase](#) in *AWS Tools for PowerShell Cmdlet Reference (V4)*.

Tools for PowerShell V5

Example 1: Creates a new case in the AWS Support Center. Values for the `-ServiceCode` and `-CategoryCode` parameters can be obtained using the `Get-ASAService` cmdlet. The value for the `-SeverityCode` parameter can be obtained using the `Get-ASASeverityLevel` cmdlet. The `-IssueType` parameter value can be either "customer-service" or "technical". If successful the AWS Support case number is output. By default the case will be handled in English, to use Japanese add the `-Language "ja"` parameter. The `-ServiceCode`, `-CategoryCode`, `-Subject` and `-CommunicationBody` parameters are mandatory.

```
New-ASACase -ServiceCode "amazon-cloudfront" -CategoryCode "APIs" -SeverityCode  
"low" -Subject "subject text" -CommunicationBody "description of the case" -  
CcEmailAddress @("email1@domain.com", "email2@domain.com") -IssueType "technical"
```

- For API details, see [CreateCase](#) in *AWS Tools for PowerShell Cmdlet Reference (V5)*.

Python

SDK for Python (Boto3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
class SupportWrapper:
    """Encapsulates Support actions."""

    def __init__(self, support_client):
        """
        :param support_client: A Boto3 Support client.
        """
        self.support_client = support_client

    @classmethod
    def from_client(cls):
        """
        Instantiates this class from a Boto3 client.
        """
        support_client = boto3.client("support")
        return cls(support_client)

    def create_case(self, service, category, severity):
        """
        Create a new support case.

        :param service: The service to use for the new case.
        :param category: The category to use for the new case.
        :param severity: The severity to use for the new case.
        :return: The caseId of the new case.
        """
        try:
            response = self.support_client.create_case(
                subject="Example case for testing, ignore.",
                serviceCode=service["code"],
                severityCode=severity["code"],
```

```

        categoryCode=category["code"],
        communicationBody="Example support case body.",
        language="en",
        issueType="customer-service",
    )
    case_id = response["caseId"]
except ClientError as err:
    if err.response["Error"]["Code"] == "SubscriptionRequiredException":
        logger.info(
            "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
            "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
            "examples."
        )
    else:
        logger.error(
            "Couldn't create case. Here's why: %s: %s",
            err.response["Error"]["Code"],
            err.response["Error"]["Message"],
        )
        raise
else:
    return case_id

```

- For API details, see [CreateCase](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use DescribeAttachment with an AWS SDK or CLI

The following code examples show how to use DescribeAttachment.

Action examples are code excerpts from larger programs and must be run in context. You can see this action in context in the following code example:

- [Learn the basics](#)

.NET

SDK for .NET

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
/// <summary>
/// Get description of a specific attachment.
/// </summary>
/// <param name="attachmentId">Id of the attachment, usually fetched by
describing the communications of a case.</param>
/// <returns>The attachment object.</returns>
public async Task<Attachment> DescribeAttachment(string attachmentId)
{
    var response = await _amazonSupport.DescribeAttachmentAsync(
        new DescribeAttachmentRequest()
        {
            AttachmentId = attachmentId
        });
    return response.Attachment;
}
```

- For API details, see [DescribeAttachment](#) in *AWS SDK for .NET API Reference*.

CLI

AWS CLI

To describe an attachment

The following describe-attachment example returns information about the attachment with the specified ID.

```
aws support describe-attachment \
```

```
--attachment-id "attachment-KBnjRNrePd9D6Jx0-Mm00xZuDEaL2JAj_0-  
gJv9qqDooTipsz3V1Nb19rCfkZneeQeDPgp8X1iVJyHH7UuhZDdNeqGoduZsPrAhyMakqlc60-  
iJjL5HqyYGiT1FG8EXAMPLE"
```

Output:

```
{  
  "attachment": {  
    "fileName": "troubleshoot-screenshot.png",  
    "data": "base64-blob"  
  }  
}
```

For more information, see [Case management](#) in the *AWS Support User Guide*.

- For API details, see [DescribeAttachment](#) in *AWS CLI Command Reference*.

Java

SDK for Java 2.x

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
public static void describeAttachment(SupportClient supportClient, String  
attachId) {  
    try {  
        DescribeAttachmentRequest attachmentRequest =  
DescribeAttachmentRequest.builder()  
            .attachmentId(attachId)  
            .build();  
  
        DescribeAttachmentResponse response =  
supportClient.describeAttachment(attachmentRequest);  
        System.out.println("The name of the file is " +  
response.attachment().fileName());  
  
    } catch (SupportException e) {
```



```
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
}
```

- For API details, see [DescribeAttachment](#) in *AWS SDK for Java 2.x API Reference*.

JavaScript

SDK for JavaScript (v3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
import { DescribeAttachmentCommand } from "@aws-sdk/client-support";

import { client } from "../libs/client.js";

export const main = async () => {
  try {
    // Get the metadata and content of an attachment.
    const response = await client.send(
      new DescribeAttachmentCommand({
        // Set value to an existing attachment id.
        // Use DescribeCommunications or DescribeCases to find an attachment id.
        attachmentId: "ATTACHMENT_ID",
      }),
    );
    console.log(response.attachment?.fileName);
    return response;
  } catch (err) {
    console.error(err);
  }
};
```

- For API details, see [DescribeAttachment](#) in *AWS SDK for JavaScript API Reference*.

Kotlin

SDK for Kotlin

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
suspend fun describeAttachment(attachId: String?) {
    val attachmentRequest =
        DescribeAttachmentRequest {
            attachmentId = attachId
        }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response = supportClient.describeAttachment(attachmentRequest)
        println("The name of the file is ${response.attachment?.fileName}")
    }
}
```

- For API details, see [DescribeAttachment](#) in *AWS SDK for Kotlin API reference*.

Python

SDK for Python (Boto3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
class SupportWrapper:
    """Encapsulates Support actions."""

    def __init__(self, support_client):
        """
```

```

        :param support_client: A Boto3 Support client.
        """
        self.support_client = support_client

    @classmethod
    def from_client(cls):
        """
        Instantiates this class from a Boto3 client.
        """
        support_client = boto3.client("support")
        return cls(support_client)

    def describe_attachment(self, attachment_id):
        """
        Get information about an attachment by its attachmentID.

        :param attachment_id: The ID of the attachment.
        :return: The name of the attached file.
        """
        try:
            response = self.support_client.describe_attachment(
                attachmentId=attachment_id
            )
            attached_file = response["attachment"]["fileName"]
        except ClientError as err:
            if err.response["Error"]["Code"] == "SubscriptionRequiredException":
                logger.info(
                    "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                    "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                    "examples."
                )
            else:
                logger.error(
                    "Couldn't get attachment description. Here's why: %s: %s",
                    err.response["Error"]["Code"],
                    err.response["Error"]["Message"],
                )
                raise
        else:
            return attached_file

```

- For API details, see [DescribeAttachment](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use DescribeCases with an AWS SDK or CLI

The following code examples show how to use DescribeCases.

Action examples are code excerpts from larger programs and must be run in context. You can see this action in context in the following code examples:

- [Learn the basics](#)
- [Getting started with technical support](#)

.NET

SDK for .NET

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
/// <summary>
/// Get case details for a list of case ids, optionally with date filters.
/// </summary>
/// <param name="caseIds">The list of case IDs.</param>
/// <param name="displayId">Optional display ID.</param>
/// <param name="includeCommunication">True to include communication.
Defaults to true.</param>
/// <param name="includeResolvedCases">True to include resolved cases.
Defaults to false.</param>
/// <param name="afterTime">The optional start date for a filtered search.</
param>
```

```

    /// <param name="beforeTime">The optional end date for a filtered search.</
param>
    /// <param name="language">Optional language support for your case.
    /// Currently Chinese ("zh"), English ("en"), Japanese ("ja") and Korean
    /// ("ko") are supported.</param>
    /// <returns>A list of CaseDetails.</returns>
    public async Task<List<CaseDetails>> DescribeCases(List<string> caseIds,
string? displayId = null, bool includeCommunication = true,
    bool includeResolvedCases = false, DateTime? afterTime = null, DateTime?
beforeTime = null,
    string language = "en")
    {
        var results = new List<CaseDetails>();
        var paginateCases = _amazonSupport.Paginators.DescribeCases(
            new DescribeCasesRequest()
            {
                CaseIdList = caseIds,
                DisplayId = displayId,
                IncludeCommunications = includeCommunication,
                IncludeResolvedCases = includeResolvedCases,
                AfterTime = afterTime?.ToString("s"),
                BeforeTime = beforeTime?.ToString("s"),
                Language = language
            });
        // Get the entire list using the paginator.
        await foreach (var cases in paginateCases.Cases)
        {
            results.Add(cases);
        }
        return results;
    }

```

- For API details, see [DescribeCases](#) in *AWS SDK for .NET API Reference*.

CLI

AWS CLI

To describe a case

The following `describe-cases` example returns information about the specified support case in your AWS account.

```
aws support describe-cases \  
  --display-id "1234567890" \  
  --after-time "2020-03-23T21:31:47.774Z" \  
  --include-resolved-cases \  
  --language "en" \  
  --no-include-communications \  
  --max-item 1
```

Output:

```
{  
  "cases": [  
    {  
      "status": "resolved",  
      "ccEmailAddresses": [],  
      "timeCreated": "2020-03-23T21:31:47.774Z",  
      "caseId": "case-12345678910-2013-c4c1d2bf33c5cf47",  
      "severityCode": "low",  
      "language": "en",  
      "categoryCode": "using-aws",  
      "serviceCode": "general-info",  
      "submittedBy": "myemail@example.com",  
      "displayId": "1234567890",  
      "subject": "Question about my account"  
    }  
  ]  
}
```

For more information, see [Case management](#) in the *AWS Support User Guide*.

- For API details, see [DescribeCases](#) in *AWS CLI Command Reference*.

Java

SDK for Java 2.x

 Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
public static void getOpenCase(SupportClient supportClient) {
    try {
        // Specify the start and end time.
        Instant now = Instant.now();
        java.time.LocalDate.now();
        Instant yesterday = now.minus(1, ChronoUnit.DAYS);

        DescribeCasesRequest describeCasesRequest =
        DescribeCasesRequest.builder()
            .maxResults(20)
            .afterTime(yesterday.toString())
            .beforeTime(now.toString())
            .build();

        DescribeCasesResponse response =
        supportClient.describeCases(describeCasesRequest);
        List<CaseDetails> cases = response.cases();
        for (CaseDetails sinCase : cases) {
            System.out.println("The case status is " + sinCase.status());
            System.out.println("The case Id is " + sinCase.caseId());
            System.out.println("The case subject is " + sinCase.subject());
        }

    } catch (SupportException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
}
```

- For API details, see [DescribeCases](#) in *AWS SDK for Java 2.x API Reference*.

JavaScript

SDK for JavaScript (v3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
import { DescribeCasesCommand } from "@aws-sdk/client-support";

import { client } from "../libs/client.js";

export const main = async () => {
  try {
    // Get all of the unresolved cases in your account.
    // Filter or expand results by providing parameters to the
    DescribeCasesCommand. Refer
    // to the TypeScript definition and the API doc for more information on
    possible parameters.
    // https://docs.aws.amazon.com/AWSJavaScriptSDK/v3/latest/clients/client-
    support/interfaces/describecasescommandinput.html
    const response = await client.send(new DescribeCasesCommand({}));
    const caseIds = response.cases.map((supportCase) => supportCase.caseId);
    console.log(caseIds);
    return response;
  } catch (err) {
    console.error(err);
  }
};
```

- For API details, see [DescribeCases](#) in *AWS SDK for JavaScript API Reference*.

Kotlin

SDK for Kotlin

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
suspend fun getOpenCase() {  
    // Specify the start and end time.  
    val now = Instant.now()  
    LocalDate.now()  
    val yesterday = now.minus(1, ChronoUnit.DAYS)  
    val describeCasesRequest =  
        DescribeCasesRequest {  
            maxResults = 20  
            afterTime = yesterday.toString()  
            beforeTime = now.toString()  
        }  
  
    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->  
        val response = supportClient.describeCases(describeCasesRequest)  
        response.cases?.forEach { sinCase ->  
            println("The case status is ${sinCase.status}")  
            println("The case Id is ${sinCase.caseId}")  
            println("The case subject is ${sinCase.subject}")  
        }  
    }  
}
```

- For API details, see [DescribeCases](#) in *AWS SDK for Kotlin API reference*.

PowerShell

Tools for PowerShell V4

Example 1: Returns the details of all support cases.

```
Get-ASACase
```

Example 2: Returns the details of all support cases since the specified date and time.

```
Get-ASACase -AfterTime "2013-09-10T03:06Z"
```

Example 3: Returns the details of the first 10 support cases, including those that have been resolved.

```
Get-ASACase -MaxResult 10 -IncludeResolvedCases $true
```

Example 4: Returns the details of the single specified support case.

```
Get-ASACase -CaseIdList "case-12345678910-2013-c4c1d2bf33c5cf47"
```

Example 5: Returns the details of specified support cases.

```
Get-ASACase -CaseIdList @("case-12345678910-2013-c4c1d2bf33c5cf47",  
"case-18929034710-2011-c4fdeabf33c5cf47")
```

- For API details, see [DescribeCases](#) in *AWS Tools for PowerShell Cmdlet Reference (V4)*.

Tools for PowerShell V5

Example 1: Returns the details of all support cases.

```
Get-ASACase
```

Example 2: Returns the details of all support cases since the specified date and time.

```
Get-ASACase -AfterTime "2013-09-10T03:06Z"
```

Example 3: Returns the details of the first 10 support cases, including those that have been resolved.

```
Get-ASACase -MaxResult 10 -IncludeResolvedCases $true
```

Example 4: Returns the details of the single specified support case.

```
Get-ASACase -CaseIdList "case-12345678910-2013-c4c1d2bf33c5cf47"
```

Example 5: Returns the details of specified support cases.

```
Get-ASACase -CaseIdList @("case-12345678910-2013-c4c1d2bf33c5cf47",
"case-18929034710-2011-c4fdeabf33c5cf47")
```

- For API details, see [DescribeCases](#) in *AWS Tools for PowerShell Cmdlet Reference (V5)*.

Python

SDK for Python (Boto3)**Note**

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
class SupportWrapper:
    """Encapsulates Support actions."""

    def __init__(self, support_client):
        """
        :param support_client: A Boto3 Support client.
        """
        self.support_client = support_client

    @classmethod
    def from_client(cls):
        """
        Instantiates this class from a Boto3 client.
        """
        support_client = boto3.client("support")
        return cls(support_client)

    def describe_cases(self, after_time, before_time, resolved):
```

```

"""
Describe support cases over a period of time, optionally filtering
by status.

:param after_time: The start time to include for cases.
:param before_time: The end time to include for cases.
:param resolved: True to include resolved cases in the results,
    otherwise results are open cases.
:return: The final status of the case.
"""
try:
    cases = []
    paginator = self.support_client.get_paginator("describe_cases")
    for page in paginator.paginate(
        afterTime=after_time,
        beforeTime=before_time,
        includeResolvedCases=resolved,
        language="en",
    ):
        cases += page["cases"]
except ClientError as err:
    if err.response["Error"]["Code"] == "SubscriptionRequiredException":
        logger.info(
            "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
            "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
            "examples."
        )
    else:
        logger.error(
            "Couldn't describe cases. Here's why: %s: %s",
            err.response["Error"]["Code"],
            err.response["Error"]["Message"],
        )
        raise
else:
    if resolved:
        cases = filter(lambda case: case["status"] == "resolved", cases)
    return cases

```

- For API details, see [DescribeCases](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use DescribeCommunications with an AWS SDK or CLI

The following code examples show how to use DescribeCommunications.

Action examples are code excerpts from larger programs and must be run in context. You can see this action in context in the following code examples:

- [Learn the basics](#)
- [Getting started with technical support](#)

.NET

SDK for .NET

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
/// <summary>
/// Describe the communications for a case, optionally with a date filter.
/// </summary>
/// <param name="caseId">The ID of the support case.</param>
/// <param name="afterTime">The optional start date for a filtered search.</
param>
/// <param name="beforeTime">The optional end date for a filtered search.</
param>
/// <returns>The list of communications for the case.</returns>
public async Task<List<Communication>> DescribeCommunications(string caseId,
DateTime? afterTime = null, DateTime? beforeTime = null)
{
    var results = new List<Communication>();
    var paginateCommunications =
        _amazonSupport.Paginators.DescribeCommunications(
```

```
new DescribeCommunicationsRequest()
{
    CaseId = caseId,
    AfterTime = afterTime?.ToString("s"),
    BeforeTime = beforeTime?.ToString("s")
});
// Get the entire list using the paginator.
await foreach (var communications in
paginateCommunications.Communications)
{
    results.Add(communications);
}
return results;
}
```

- For API details, see [DescribeCommunications](#) in *AWS SDK for .NET API Reference*.

CLI

AWS CLI

To describe the latest communication for a case

The following describe-communications example returns the latest communication for the specified support case in your AWS account.

```
aws support describe-communications \
  --case-id "case-12345678910-2013-c4c1d2bf33c5cf47" \
  --after-time "2020-03-23T21:31:47.774Z" \
  --max-item 1
```

Output:

```
{
  "communications": [
    {
      "body": "I want to learn more about an AWS service.",
      "attachmentSet": [],
      "caseId": "case-12345678910-2013-c4c1d2bf33c5cf47",
      "timeCreated": "2020-05-12T23:12:35.000Z",
```

```
        "submittedBy": "Amazon Web Services"
    },
    "NextToken":
    "eyJhZG9rZW4iOiBudWxsLCAiYm90b190cnVuY2F0ZV9hbW91bnQEXAMPLE=="
}
```

For more information, see [Case management](#) in the *AWS Support User Guide*.

- For API details, see [DescribeCommunications](#) in *AWS CLI Command Reference*.

Java

SDK for Java 2.x

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
public static String listCommunications(SupportClient supportClient, String
caseId) {
    try {
        String attachId = null;
        DescribeCommunicationsRequest communicationsRequest =
        DescribeCommunicationsRequest.builder()
            .caseId(caseId)
            .maxResults(10)
            .build();

        DescribeCommunicationsResponse response =
        supportClient.describeCommunications(communicationsRequest);
        List<Communication> communications = response.communications();
        for (Communication comm : communications) {
            System.out.println("the body is: " + comm.body());

            // Get the attachment id value.
            List<AttachmentDetails> attachments = comm.attachmentSet();
            for (AttachmentDetails detail : attachments) {
                attachId = detail.attachmentId();
            }
        }
    }
}
```

```
        }
        return attachId;

    } catch (SupportException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
    return "";
}
```

- For API details, see [DescribeCommunications](#) in *AWS SDK for Java 2.x API Reference*.

JavaScript

SDK for JavaScript (v3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
import { DescribeCommunicationsCommand } from "@aws-sdk/client-support";

import { client } from "../libs/client.js";

export const main = async () => {
    try {
        // Get all communications for the support case.
        // Filter results by providing parameters to the
        DescribeCommunicationsCommand. Refer
        // to the TypeScript definition and the API doc for more information on
        possible parameters.
        // https://docs.aws.amazon.com/AWSJavaScriptSDK/v3/latest/clients/client-
        support/interfaces/describecommunicationscommandinput.html
        const response = await client.send(
            new DescribeCommunicationsCommand({
                // Set value to an existing case id.
                caseId: "CASE_ID",
            }),
        );
    }
```



```
const text = response.communications.map((item) => item.body).join("\n");
console.log(text);
return response;
} catch (err) {
  console.error(err);
}
};
```

- For API details, see [DescribeCommunications](#) in *AWS SDK for JavaScript API Reference*.

Kotlin

SDK for Kotlin

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
suspend fun listCommunications(caseIdVal: String?): String? {
    val communicationsRequest =
        DescribeCommunicationsRequest {
            caseId = caseIdVal
            maxResults = 10
        }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response =
            supportClient.describeCommunications(communicationsRequest)
        response.communications?.forEach { comm ->
            println("the body is: " + comm.body)
            comm.attachmentSet?.forEach { detail ->
                return detail.attachmentId
            }
        }
    }
    return ""
}
```

- For API details, see [DescribeCommunications](#) in *AWS SDK for Kotlin API reference*.

PowerShell

Tools for PowerShell V4

Example 1: Returns all communications for the specified case.

```
Get-ASACommunication -CaseId "case-12345678910-2013-c4c1d2bf33c5cf47"
```

Example 2: Returns all communications since midnight UTC on January 1st 2012 for the specified case.

```
Get-ASACommunication -CaseId "case-12345678910-2013-c4c1d2bf33c5cf47" -AfterTime  
"2012-01-10T00:00Z"
```

- For API details, see [DescribeCommunications](#) in *AWS Tools for PowerShell Cmdlet Reference (V4)*.

Tools for PowerShell V5

Example 1: Returns all communications for the specified case.

```
Get-ASACommunication -CaseId "case-12345678910-2013-c4c1d2bf33c5cf47"
```

Example 2: Returns all communications since midnight UTC on January 1st 2012 for the specified case.

```
Get-ASACommunication -CaseId "case-12345678910-2013-c4c1d2bf33c5cf47" -AfterTime  
"2012-01-10T00:00Z"
```

- For API details, see [DescribeCommunications](#) in *AWS Tools for PowerShell Cmdlet Reference (V5)*.

Python

SDK for Python (Boto3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
class SupportWrapper:
    """Encapsulates Support actions."""

    def __init__(self, support_client):
        """
        :param support_client: A Boto3 Support client.
        """
        self.support_client = support_client

    @classmethod
    def from_client(cls):
        """
        Instantiates this class from a Boto3 client.
        """
        support_client = boto3.client("support")
        return cls(support_client)

    def describe_all_case_communications(self, case_id):
        """
        Describe all the communications for a case using a paginator.

        :param case_id: The ID of the case.
        :return: The communications for the case.
        """
        try:
            communications = []
            paginator =
self.support_client.get_paginator("describe_communications")
            for page in paginator.paginate(caseId=case_id):
                communications += page["communications"]
        except ClientError as err:
```

```
        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                "examples."
            )
        else:
            logger.error(
                "Couldn't describe communications. Here's why: %s: %s",
                err.response["Error"]["Code"],
                err.response["Error"]["Message"],
            )
            raise
    else:
        return communications
```

- For API details, see [DescribeCommunications](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use DescribeServices with an AWS SDK or CLI

The following code examples show how to use DescribeServices.

Action examples are code excerpts from larger programs and must be run in context. You can see this action in context in the following code examples:

- [Learn the basics](#)
- [Getting started with technical support](#)

.NET

SDK for .NET

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
/// <summary>
/// Get the descriptions of AWS services.
/// </summary>
/// <param name="name">Optional language for services.
/// Currently Chinese ("zh"), English ("en"), Japanese ("ja") and Korean
/// ("ko") are supported.</param>
/// <returns>The list of AWS service descriptions.</returns>
public async Task<List<Service>> DescribeServices(string language = "en")
{
    var response = await _amazonSupport.DescribeServicesAsync(
        new DescribeServicesRequest()
        {
            Language = language
        });
    return response.Services;
}
```

- For API details, see [DescribeServices](#) in *AWS SDK for .NET API Reference*.

CLI

AWS CLI

To list AWS services and service categories

The following `describe-services` example lists the available service categories for requesting general information.

```
aws support describe-services \
  --service-code-list general-info
```

Output:

```
{
  "services": [
    {
      "code": "general-info",
      "name": "General Info and Getting Started",
      "categories": [
        {
          "code": "charges",
          "name": "How Will I Be Charged?"
        },
        {
          "code": "gdpr-queries",
          "name": "Data Privacy Query"
        },
        {
          "code": "reserved-instances",
          "name": "Reserved Instances"
        },
        {
          "code": "resource",
          "name": "Where is my Resource?"
        },
        {
          "code": "using-aws",
          "name": "Using AWS & Services"
        },
        {
          "code": "free-tier",
          "name": "Free Tier"
        },
        {
          "code": "security-and-compliance",
          "name": "Security & Compliance"
        },
        {
          "code": "account-structure",
          "name": "Account Structure"
        }
      ]
    }
  ]
}
```

```

    ]
  }
]
}

```

For more information, see [Case management](#) in the *AWS Support User Guide*.

- For API details, see [DescribeServices](#) in *AWS CLI Command Reference*.

Java

SDK for Java 2.x

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```

// Return a List that contains a Service name and Category name.
public static List<String> displayServices(SupportClient supportClient) {
    try {
        DescribeServicesRequest servicesRequest =
DescribeServicesRequest.builder()
                        .language("en")
                        .build();

        DescribeServicesResponse response =
supportClient.describeServices(servicesRequest);
        String serviceCode = null;
        String catName = null;
        List<String> sevCatList = new ArrayList<>();
        List<Service> services = response.services();

        System.out.println("Get the first 10 services");
        int index = 1;
        for (Service service : services) {
            if (index == 11)
                break;

            System.out.println("The Service name is: " + service.name());
            if (service.name().compareTo("Account") == 0)

```

```
        serviceCode = service.code();

        // Get the Categories for this service.
        List<Category> categories = service.categories();
        for (Category cat : categories) {
            System.out.println("The category name is: " + cat.name());
            if (cat.name().compareTo("Security") == 0)
                catName = cat.name();
        }
        index++;
    }

    // Push the two values to the list.
    sevCatList.add(serviceCode);
    sevCatList.add(catName);
    return sevCatList;

} catch (SupportException e) {
    System.out.println(e.getLocalizedMessage());
    System.exit(1);
}
return null;
}
```

- For API details, see [DescribeServices](#) in *AWS SDK for Java 2.x API Reference*.

Kotlin

SDK for Kotlin

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
// Return a List that contains a Service name and Category name.
suspend fun displayServices(): List<String> {
    var serviceCode = ""
    var catName = ""
    val sevCatList = mutableListOf<String>()
```



```

val servicesRequest =
    DescribeServicesRequest {
        language = "en"
    }

SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
    val response = supportClient.describeServices(servicesRequest)
    println("Get the first 10 services")
    var index = 1

    response.services?.forEach { service ->
        if (index == 11) {
            return@forEach
        }

        println("The Service name is ${service.name}")
        if (service.name == "Account") {
            serviceCode = service.code.toString()
        }

        // Get the categories for this service.
        service.categories?.forEach { cat ->
            println("The category name is ${cat.name}")
            if (cat.name == "Security") {
                catName = cat.name!!
            }
        }
        index++
    }
}

// Push the two values to the list.
serviceCode.let { sevCatList.add(it) }
catName.let { sevCatList.add(it) }
return sevCatList
}

```

- For API details, see [DescribeServices](#) in *AWS SDK for Kotlin API reference*.

PowerShell

Tools for PowerShell V4

Example 1: Returns all available service codes, names and categories.

```
Get-ASAService
```

Example 2: Returns the name and categories for the service with the specified code.

```
Get-ASAService -ServiceCodeList "amazon-cloudfront"
```

Example 3: Returns the name and categories for the specified service codes.

```
Get-ASAService -ServiceCodeList @("amazon-cloudfront", "amazon-cloudwatch")
```

Example 4: Returns the name and categories (in Japanese) for the specified service codes. Currently English ("en") and Japanese ("ja") language codes are supported.

```
Get-ASAService -ServiceCodeList @("amazon-cloudfront", "amazon-cloudwatch") -  
Language "ja"
```

- For API details, see [DescribeServices](#) in *AWS Tools for PowerShell Cmdlet Reference (V4)*.

Tools for PowerShell V5

Example 1: Returns all available service codes, names and categories.

```
Get-ASAService
```

Example 2: Returns the name and categories for the service with the specified code.

```
Get-ASAService -ServiceCodeList "amazon-cloudfront"
```

Example 3: Returns the name and categories for the specified service codes.

```
Get-ASAService -ServiceCodeList @("amazon-cloudfront", "amazon-cloudwatch")
```

Example 4: Returns the name and categories (in Japanese) for the specified service codes. Currently English ("en") and Japanese ("ja") language codes are supported.

```
Get-ASAService -ServiceCodeList @("amazon-cloudfront", "amazon-cloudwatch") -
Language "ja"
```

- For API details, see [DescribeServices](#) in *AWS Tools for PowerShell Cmdlet Reference (V5)*.

Python

SDK for Python (Boto3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
class SupportWrapper:
    """Encapsulates Support actions."""

    def __init__(self, support_client):
        """
        :param support_client: A Boto3 Support client.
        """
        self.support_client = support_client

    @classmethod
    def from_client(cls):
        """
        Instantiates this class from a Boto3 client.
        """
        support_client = boto3.client("support")
        return cls(support_client)

    def describe_services(self, language):
        """
        Get the descriptions of AWS services available for support for a
        language.

        :param language: The language for support services.
        Currently, only "en" (English) and "ja" (Japanese) are supported.
        :return: The list of AWS service descriptions.
```

```

    """
    try:
        response = self.support_client.describe_services(language=language)
        services = response["services"]
    except ClientError as err:
        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                "examples."
            )
        else:
            logger.error(
                "Couldn't get Support services for language %s. Here's why:
%s: %s",
                language,
                err.response["Error"]["Code"],
                err.response["Error"]["Message"],
            )
            raise
    else:
        return services

```

- For API details, see [DescribeServices](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use DescribeSeverityLevels with an AWS SDK or CLI

The following code examples show how to use DescribeSeverityLevels.

Action examples are code excerpts from larger programs and must be run in context. You can see this action in context in the following code examples:

- [Learn the basics](#)
- [Getting started with technical support](#)

.NET

SDK for .NET

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
/// <summary>
/// Get the descriptions of support severity levels.
/// </summary>
/// <param name="name">Optional language for severity levels.
/// Currently Chinese ("zh"), English ("en"), Japanese ("ja") and Korean
/// ("ko") are supported.</param>
/// <returns>The list of support severity levels.</returns>
public async Task<List<SeverityLevel>> DescribeSeverityLevels(string language
= "en")
{
    var response = await _amazonSupport.DescribeSeverityLevelsAsync(
        new DescribeSeverityLevelsRequest()
        {
            Language = language
        });
    return response.SeverityLevels;
}
```

- For API details, see [DescribeSeverityLevels](#) in *AWS SDK for .NET API Reference*.

CLI

AWS CLI

To list the available severity levels

The following `describe-severity-levels` example lists the available severity levels for a support case.

aws support describe-severity-levels

Output:

```
{
  "severityLevels": [
    {
      "code": "low",
      "name": "Low"
    },
    {
      "code": "normal",
      "name": "Normal"
    },
    {
      "code": "high",
      "name": "High"
    },
    {
      "code": "urgent",
      "name": "Urgent"
    },
    {
      "code": "critical",
      "name": "Critical"
    }
  ]
}
```

For more information, see [Choosing a severity](#) in the *AWS Support User Guide*.

- For API details, see [DescribeSeverityLevels](#) in *AWS CLI Command Reference*.

Java

SDK for Java 2.x** Note**

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
public static String displaySevLevels(SupportClient supportClient) {
    try {
        DescribeSeverityLevelsRequest severityLevelsRequest =
DescribeSeverityLevelsRequest.builder()
            .language("en")
            .build();

        DescribeSeverityLevelsResponse response =
supportClient.describeSeverityLevels(severityLevelsRequest);
        List<SeverityLevel> severityLevels = response.severityLevels();
        String levelName = null;
        for (SeverityLevel sevLevel : severityLevels) {
            System.out.println("The severity level name is: " +
sevLevel.name());
            if (sevLevel.name().compareTo("High") == 0)
                levelName = sevLevel.name();
        }
        return levelName;

    } catch (SupportException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
    return "";
}
```

- For API details, see [DescribeSeverityLevels](#) in *AWS SDK for Java 2.x API Reference*.

JavaScript

SDK for JavaScript (v3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
import { DescribeSeverityLevelsCommand } from "@aws-sdk/client-support";
```

```
import { client } from "../libs/client.js";

export const main = async () => {
  try {
    // Get the list of severity levels.
    // The available values depend on the support plan for the account.
    const response = await client.send(new DescribeSeverityLevelsCommand({}));
    console.log(response.severityLevels);
    return response;
  } catch (err) {
    console.error(err);
  }
};
```

- For API details, see [DescribeSeverityLevels](#) in *AWS SDK for JavaScript API Reference*.

Kotlin

SDK for Kotlin

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
suspend fun displaySevLevels(): String {
    var levelName = ""
    val severityLevelsRequest =
        DescribeSeverityLevelsRequest {
            language = "en"
        }

    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->
        val response =
            supportClient.describeSeverityLevels(severityLevelsRequest)
        response.severityLevels?.forEach { sevLevel ->
            println("The severity level name is: ${sevLevel.name}")
            if (sevLevel.name == "High") {
                levelName = sevLevel.name!!
            }
        }
    }
}
```



```
        }  
        return levelName  
    }  
}
```

- For API details, see [DescribeSeverityLevels](#) in *AWS SDK for Kotlin API reference*.

PowerShell

Tools for PowerShell V4

Example 1: Returns the list of severity levels that can be assigned to an AWS Support case.

```
Get-ASASeverityLevel
```

Example 2: Returns the list of severity levels that can be assigned to an AWS Support case. The names of the levels are returned in Japanese.

```
Get-ASASeverityLevel -Language "ja"
```

- For API details, see [DescribeSeverityLevels](#) in *AWS Tools for PowerShell Cmdlet Reference (V4)*.

Tools for PowerShell V5

Example 1: Returns the list of severity levels that can be assigned to an AWS Support case.

```
Get-ASASeverityLevel
```

Example 2: Returns the list of severity levels that can be assigned to an AWS Support case. The names of the levels are returned in Japanese.

```
Get-ASASeverityLevel -Language "ja"
```

- For API details, see [DescribeSeverityLevels](#) in *AWS Tools for PowerShell Cmdlet Reference (V5)*.

Python

SDK for Python (Boto3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
class SupportWrapper:
    """Encapsulates Support actions."""

    def __init__(self, support_client):
        """
        :param support_client: A Boto3 Support client.
        """
        self.support_client = support_client

    @classmethod
    def from_client(cls):
        """
        Instantiates this class from a Boto3 client.
        """
        support_client = boto3.client("support")
        return cls(support_client)

    def describe_severity_levels(self, language):
        """
        Get the descriptions of available severity levels for support cases for a
        language.

        :param language: The language for support severity levels.
        Currently, only "en" (English) and "ja" (Japanese) are supported.
        :return: The list of severity levels.
        """
        try:
            response =
self.support_client.describe_severity_levels(language=language)
            severity_levels = response["severityLevels"]
        except ClientError as err:
```

```

        if err.response["Error"]["Code"] == "SubscriptionRequiredException":
            logger.info(
                "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
                "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
                "examples."
            )
        else:
            logger.error(
                "Couldn't get severity levels for language %s. Here's why:
%s: %s",
                language,
                err.response["Error"]["Code"],
                err.response["Error"]["Message"],
            )
            raise
    else:
        return severity_levels

```

- For API details, see [DescribeSeverityLevels](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use DescribeTrustedAdvisorCheckRefreshStatuses with a CLI

The following code examples show how to use `DescribeTrustedAdvisorCheckRefreshStatuses`.

CLI

AWS CLI

To list the refresh statuses of AWS Trusted Advisor checks

The following `describe-trusted-advisor-check-refresh-statuses` example lists the refresh statuses for two Trusted Advisor checks: Amazon S3 Bucket Permissions and IAM Use.

```
aws support describe-trusted-advisor-check-refresh-statuses \
  --check-id "Pfx0RwqBli" "zXCkfM1nI3"
```

Output:

```
{
  "statuses": [
    {
      "checkId": "Pfx0RwqBli",
      "status": "none",
      "millisUntilNextRefreshable": 0
    },
    {
      "checkId": "zXCkfM1nI3",
      "status": "none",
      "millisUntilNextRefreshable": 0
    }
  ]
}
```

For more information, see [AWS Trusted Advisor](#) in the *AWS Support User Guide*.

- For API details, see [DescribeTrustedAdvisorCheckRefreshStatuses](#) in *AWS CLI Command Reference*.

PowerShell

Tools for PowerShell V4

Example 1: Returns the current status of refresh requests for the specified checks. Request-ASATrustedAdvisorCheckRefresh can be used to request that the status information of the checks be refreshed.

```
Get-ASATrustedAdvisorCheckRefreshStatus -CheckId @("checkid1", "checkid2")
```

- For API details, see [DescribeTrustedAdvisorCheckRefreshStatuses](#) in *AWS Tools for PowerShell Cmdlet Reference (V4)*.

Tools for PowerShell V5

Example 1: Returns the current status of refresh requests for the specified checks. Request-ASATrustedAdvisorCheckRefresh can be used to request that the status information of the checks be refreshed.

```
Get-ASATrustedAdvisorCheckRefreshStatus -CheckId @("checkid1", "checkid2")
```

- For API details, see [DescribeTrustedAdvisorCheckRefreshStatuses](#) in *AWS Tools for PowerShell Cmdlet Reference (V5)*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use DescribeTrustedAdvisorCheckResult with a CLI

The following code examples show how to use DescribeTrustedAdvisorCheckResult.

CLI

AWS CLI

To list the results of an AWS Trusted Advisor check

The following describe-trusted-advisor-check-result example lists the results of the IAM Use check.

```
aws support describe-trusted-advisor-check-result \  
  --check-id "zXCkfM1nI3"
```

Output:

```
{  
  "result": {  
    "checkId": "zXCkfM1nI3",  
    "timestamp": "2020-05-13T21:38:05Z",  
    "status": "ok",  
    "resourcesSummary": {  
      "resourcesProcessed": 1,  
      "resourcesFlagged": 0,  
    }  
  }  
}
```

```

        "resourcesIgnored": 0,
        "resourcesSuppressed": 0
    },
    "categorySpecificSummary": {
        "costOptimizing": {
            "estimatedMonthlySavings": 0.0,
            "estimatedPercentMonthlySavings": 0.0
        }
    },
    "flaggedResources": [
        {
            "status": "ok",
            "resourceId": "47DEQpj8HBSa-_TImW-5JCeuQeRkm5NMpJWZEXAMPLE",
            "isSuppressed": false
        }
    ]
}

```

For more information, see [AWS Trusted Advisor](#) in the *AWS Support User Guide*.

- For API details, see [DescribeTrustedAdvisorCheckResult](#) in *AWS CLI Command Reference*.

PowerShell

Tools for PowerShell V4

Example 1: Returns the results of a Trusted Advisor check. The list of available Trusted Advisor checks can be obtained using `Get-ASATrustedAdvisorChecks`. The output is the overall status of the check, the timestamp at which the check was last run and the unique checkid for the specific check. To have the results output in Japanese, add the `-Language "ja"` parameter.

```
Get-ASATrustedAdvisorCheckResult -CheckId "checkid1"
```

- For API details, see [DescribeTrustedAdvisorCheckResult](#) in *AWS Tools for PowerShell Cmdlet Reference (V4)*.

Tools for PowerShell V5

Example 1: Returns the results of a Trusted Advisor check. The list of available Trusted Advisor checks can be obtained using `Get-ASATrustedAdvisorChecks`. The output is

the overall status of the check, the timestamp at which the check was last run and the unique checkid for the specific check. To have the results output in Japanese, add the -Language "ja" parameter.

```
Get-ASATrustedAdvisorCheckResult -CheckId "checkid1"
```

- For API details, see [DescribeTrustedAdvisorCheckResult](#) in *AWS Tools for PowerShell Cmdlet Reference (V5)*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use DescribeTrustedAdvisorCheckSummaries with a CLI

The following code examples show how to use DescribeTrustedAdvisorCheckSummaries.

CLI

AWS CLI

To list the summaries of AWS Trusted Advisor checks

The following describe-trusted-advisor-check-summaries example lists the results for two Trusted Advisor checks: Amazon S3 Bucket Permissions and IAM Use.

```
aws support describe-trusted-advisor-check-summaries \
  --check-ids "Pfx0RwqBli" "zXCkfM1nI3"
```

Output:

```
{
  "summaries": [
    {
      "checkId": "Pfx0RwqBli",
      "timestamp": "2020-05-13T21:38:12Z",
      "status": "ok",
      "hasFlaggedResources": true,
      "resourcesSummary": {
```

```

        "resourcesProcessed": 44,
        "resourcesFlagged": 0,
        "resourcesIgnored": 0,
        "resourcesSuppressed": 0
    },
    "categorySpecificSummary": {
        "costOptimizing": {
            "estimatedMonthlySavings": 0.0,
            "estimatedPercentMonthlySavings": 0.0
        }
    }
},
{
    "checkId": "zXCkfM1nI3",
    "timestamp": "2020-05-13T21:38:05Z",
    "status": "ok",
    "hasFlaggedResources": true,
    "resourcesSummary": {
        "resourcesProcessed": 1,
        "resourcesFlagged": 0,
        "resourcesIgnored": 0,
        "resourcesSuppressed": 0
    },
    "categorySpecificSummary": {
        "costOptimizing": {
            "estimatedMonthlySavings": 0.0,
            "estimatedPercentMonthlySavings": 0.0
        }
    }
}
]
}

```

For more information, see [AWS Trusted Advisor](#) in the *AWS Support User Guide*.

- For API details, see [DescribeTrustedAdvisorCheckSummaries](#) in *AWS CLI Command Reference*.

PowerShell

Tools for PowerShell V4

Example 1: Returns the latest summary for the specified Trusted Advisor check.


```
Get-ASATrustedAdvisorCheckSummary -CheckId "checkid1"
```

Example 2: Returns the latest summaries for the specified Trusted Advisor checks.

```
Get-ASATrustedAdvisorCheckSummary -CheckId @("checkid1", "checkid2")
```

- For API details, see [DescribeTrustedAdvisorCheckSummaries](#) in *AWS Tools for PowerShell Cmdlet Reference (V4)*.

Tools for PowerShell V5

Example 1: Returns the latest summary for the specified Trusted Advisor check.

```
Get-ASATrustedAdvisorCheckSummary -CheckId "checkid1"
```

Example 2: Returns the latest summaries for the specified Trusted Advisor checks.

```
Get-ASATrustedAdvisorCheckSummary -CheckId @("checkid1", "checkid2")
```

- For API details, see [DescribeTrustedAdvisorCheckSummaries](#) in *AWS Tools for PowerShell Cmdlet Reference (V5)*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use DescribeTrustedAdvisorChecks with a CLI

The following code examples show how to use DescribeTrustedAdvisorChecks.

CLI

AWS CLI

To list the available AWS Trusted Advisor checks

The following describe-trusted-advisor-checks example lists the available Trusted Advisor checks in your AWS account. This information includes the check name, ID, description, category, and metadata. Note that the output is shortened for readability.

```
aws support describe-trusted-advisor-checks \
  --language "en"
```

Output:

```
{
  "checks": [
    {
      "id": "zXCkfM1nI3",
      "name": "IAM Use",
      "description": "Checks for your use of AWS Identity and Access
Management (IAM). You can use IAM to create users, groups, and roles in
AWS, and you can use permissions to control access to AWS resources. \n<br>
\n<br>\n<b>Alert Criteria</b><br>\nYellow: No IAM users have been created
for this account.\n<br>\n<br>\n<b>Recommended Action</b><br>\nCreate one or
more IAM users and groups in your account. You can then create additional
users whose permissions are limited to perform specific tasks in your AWS
environment. For more information, see <a href=\"https://docs.aws.amazon.com/
IAM/latest/UserGuide/IAMGettingStarted.html\" target=\"_blank\">Getting
Started</a>. \n<br><br>\n<b>Additional Resources</b><br>\n<a href=\"https://
docs.aws.amazon.com/IAM/latest/UserGuide/IAM_Introduction.html\" target=\"_blank
\">What Is IAM?</a>",
      "category": "security",
      "metadata": []
    }
  ]
}
```

For more information, see [AWS Trusted Advisor](#) in the *AWS Support User Guide*.

- For API details, see [DescribeTrustedAdvisorChecks](#) in *AWS CLI Command Reference*.

PowerShell

Tools for PowerShell V4

Example 1: Returns the collection of Trusted Advisor checks. You must specify the Language parameter which can accept either "en" for English output or "ja" for Japanese output.

```
Get-ASATrustedAdvisorCheck -Language "en"
```

- For API details, see [DescribeTrustedAdvisorChecks](#) in *AWS Tools for PowerShell Cmdlet Reference (V4)*.

Tools for PowerShell V5

Example 1: Returns the collection of Trusted Advisor checks. You must specify the Language parameter which can accept either "en" for English output or "ja" for Japanese output.

```
Get-ASATrustedAdvisorCheck -Language "en"
```

- For API details, see [DescribeTrustedAdvisorChecks](#) in *AWS Tools for PowerShell Cmdlet Reference (V5)*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use RefreshTrustedAdvisorCheck with a CLI

The following code examples show how to use RefreshTrustedAdvisorCheck.

CLI

AWS CLI

To refresh an AWS Trusted Advisor check

The following refresh-trusted-advisor-check example refreshes the Amazon S3 Bucket Permissions Trusted Advisor check in your AWS account.

```
aws support refresh-trusted-advisor-check \  
  --check-id "Pfx0RwqBli"
```

Output:

```
{  
  "status": {  
    "checkId": "Pfx0RwqBli",  
    "status": "enqueued",  
    "millisUntilNextRefreshable": 3599992  
  }  
}
```

```
}  
}
```

For more information, see [AWS Trusted Advisor](#) in the *AWS Support User Guide*.

- For API details, see [RefreshTrustedAdvisorCheck](#) in *AWS CLI Command Reference*.

PowerShell

Tools for PowerShell V4

Example 1: Requests a refresh for the specified Trusted Advisor check.

```
Request-ASATrustedAdvisorCheckRefresh -CheckId "checkid1"
```

- For API details, see [RefreshTrustedAdvisorCheck](#) in *AWS Tools for PowerShell Cmdlet Reference (V4)*.

Tools for PowerShell V5

Example 1: Requests a refresh for the specified Trusted Advisor check.

```
Request-ASATrustedAdvisorCheckRefresh -CheckId "checkid1"
```

- For API details, see [RefreshTrustedAdvisorCheck](#) in *AWS Tools for PowerShell Cmdlet Reference (V5)*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Use ResolveCase with an AWS SDK or CLI

The following code examples show how to use ResolveCase.

Action examples are code excerpts from larger programs and must be run in context. You can see this action in context in the following code examples:

- [Learn the basics](#)
- [Getting started with technical support](#)

.NET

SDK for .NET

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
/// <summary>
/// Resolve a support case by caseId.
/// </summary>
/// <param name="caseId">Id for the support case.</param>
/// <returns>The final status of the case after resolving.</returns>
public async Task<string> ResolveCase(string caseId)
{
    var response = await _amazonSupport.ResolveCaseAsync(
        new ResolveCaseRequest()
        {
            CaseId = caseId
        });
    return response.FinalCaseStatus;
}
```

- For API details, see [ResolveCase](#) in *AWS SDK for .NET API Reference*.

CLI

AWS CLI

To resolve a support case

The following `resolve-case` example resolves a support case in your AWS account.

```
aws support resolve-case \
  --case-id "case-12345678910-2013-c4c1d2bf33c5cf47"
```

Output:

```
{
  "finalCaseStatus": "resolved",
  "initialCaseStatus": "work-in-progress"
}
```

For more information, see [Case management](#) in the *AWS Support User Guide*.

- For API details, see [ResolveCase](#) in *AWS CLI Command Reference*.

Java**SDK for Java 2.x****Note**

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
public static void resolveSupportCase(SupportClient supportClient, String
caseId) {
    try {
        ResolveCaseRequest caseRequest = ResolveCaseRequest.builder()
            .caseId(caseId)
            .build();

        ResolveCaseResponse response =
supportClient.resolveCase(caseRequest);
        System.out.println("The status of case " + caseId + " is " +
response.finalCaseStatus());

    } catch (SupportException e) {
        System.out.println(e.getLocalizedMessage());
        System.exit(1);
    }
}
```

- For API details, see [ResolveCase](#) in *AWS SDK for Java 2.x API Reference*.

JavaScript

SDK for JavaScript (v3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
import { ResolveCaseCommand } from "@aws-sdk/client-support";

import { client } from "../libs/client.js";

const main = async () => {
  try {
    const response = await client.send(
      new ResolveCaseCommand({
        caseId: "CASE_ID",
      }),
    );

    console.log(response.finalCaseStatus);
    return response;
  } catch (err) {
    console.error(err);
  }
};
```

- For API details, see [ResolveCase](#) in *AWS SDK for JavaScript API Reference*.

Kotlin

SDK for Kotlin

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
suspend fun resolveSupportCase(caseIdVal: String) {  
    val caseRequest =  
        ResolveCaseRequest {  
            caseId = caseIdVal  
        }  
    SupportClient.fromEnvironment { region = "us-west-2" }.use { supportClient ->  
        val response = supportClient.resolveCase(caseRequest)  
        println("The status of case $caseIdVal is ${response.finalCaseStatus}")  
    }  
}
```

- For API details, see [ResolveCase](#) in *AWS SDK for Kotlin API reference*.

PowerShell

Tools for PowerShell V4

Example 1: Returns the initial state of the specified case and the current state after the call to resolve it is completed.

```
Resolve-ASACase -CaseId "case-12345678910-2013-c4c1d2bf33c5cf47"
```

- For API details, see [ResolveCase](#) in *AWS Tools for PowerShell Cmdlet Reference (V4)*.

Tools for PowerShell V5

Example 1: Returns the initial state of the specified case and the current state after the call to resolve it is completed.

```
Resolve-ASACase -CaseId "case-12345678910-2013-c4c1d2bf33c5cf47"
```

- For API details, see [ResolveCase](#) in *AWS Tools for PowerShell Cmdlet Reference (V5)*.

Python

SDK for Python (Boto3)

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [AWS Code Examples Repository](#).

```
class SupportWrapper:
    """Encapsulates Support actions."""

    def __init__(self, support_client):
        """
        :param support_client: A Boto3 Support client.
        """
        self.support_client = support_client

    @classmethod
    def from_client(cls):
        """
        Instantiates this class from a Boto3 client.
        """
        support_client = boto3.client("support")
        return cls(support_client)

    def resolve_case(self, case_id):
        """
        Resolve a support case by its caseId.

        :param case_id: The ID of the case to resolve.
        :return: The final status of the case.
        """
        try:
            response = self.support_client.resolve_case(caseId=case_id)
            final_status = response["finalCaseStatus"]
        except ClientError as err:
            if err.response["Error"]["Code"] == "SubscriptionRequiredException":
                logger.info(
```

```

        "You must have a Business, Enterprise On-Ramp, or Enterprise
Support "
        "plan to use the AWS Support API. \n\tPlease upgrade your
subscription to run these "
        "examples."
    )
    else:
        logger.error(
            "Couldn't resolve case. Here's why: %s: %s",
            err.response["Error"]["Code"],
            err.response["Error"]["Message"],
        )
        raise
    else:
        return final_status

```

- For API details, see [ResolveCase](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Scenarios for Support using AWS SDKs

The following code examples show you how to implement common scenarios in Support with AWS SDKs. These scenarios show you how to accomplish specific tasks by calling multiple functions within Support or combined with other AWS services. Each scenario includes a link to the complete source code, where you can find instructions on how to set up and run the code.

Scenarios target an intermediate level of experience to help you understand service actions in context.

Examples

- [Getting started with technical support](#)

Getting started with technical support

The following code example shows how to:

- Check available services and severity levels
- Create a support case
- Add communications to a case

Bash

AWS CLI with Bash script

Note

There's more on GitHub. Find the complete example and learn how to set up and run in the [Sample developer tutorials](#) repository.

```
#!/bin/bash

# AWS Support CLI Tutorial Script
# This script demonstrates how to use AWS Support API through AWS CLI

set -euo pipefail

# Security: Validate script location and permissions
SCRIPT_DIR="$(cd "$(dirname "${BASH_SOURCE[0]}")" && pwd)"
readonly SCRIPT_DIR

# Security: Use secure temporary directory
readonly TEMP_DIR="$(mktemp -d)" || { echo "Failed to create temp directory";
  exit 1; }
trap "rm -rf '$TEMP_DIR'" EXIT

# Set up logging with secure permissions
LOG_FILE="${TEMP_DIR}/aws-support-tutorial.log"
touch "$LOG_FILE"
chmod 600 "$LOG_FILE"

# Security: Validate AWS CLI is available
if ! command -v aws &> /dev/null; then
  echo "ERROR: AWS CLI is not installed or not in PATH" >&2
  exit 1
fi
```

```

# Security: Check AWS credentials are configured
if ! aws sts get-caller-identity &> /dev/null; then
    echo "ERROR: AWS credentials are not properly configured" >&2
    exit 1
fi

{
    echo "Starting AWS Support Tutorial at $(date)"
    echo "Script: $0"
    echo "User: $(whoami)"
    echo "---"
} >> "$LOG_FILE"

# Function to log commands and their outputs securely
log_cmd() {
    local cmd="$1"
    echo "$(date): Running command" >> "$LOG_FILE"
    # Don't echo the actual command to prevent credential leakage
    eval "$cmd" 2>&1 | tee -a "$LOG_FILE"
    return ${PIPESTATUS[0]}
}

# Function to check for errors in command output
check_error() {
    local cmd_output="$1"
    local cmd_status="$2"
    local error_msg="$3"
    local is_fatal="${4:-true}"

    if [[ $cmd_status -ne 0 || "$cmd_output" =~ [Ee][Rr][Rr][Oo][Rr] ]]; then
        echo "ERROR: $error_msg" | tee -a "$LOG_FILE"
        echo "Command returned status: $cmd_status" >> "$LOG_FILE"

        # Check for subscription error
        if [[ "$cmd_output" =~ "SubscriptionRequiredException" ]]; then
            {
                echo ""
                echo "=====
                echo "IMPORTANT: This account does not have the required AWS
Support plan."
                echo "You need a Business, Enterprise On-Ramp, or Enterprise
Support plan"
                echo "to use the AWS Support API."
                echo ""
            }
        fi
    fi
}

```

```

        echo "This script will now demonstrate the commands that would be
run"

        echo "if you had the appropriate support plan, but will not
execute them."

        echo "=====
    } | tee -a "$LOG_FILE"

    DEMO_MODE=true
    return 0
fi

    if [[ "$is_fatal" == "true" ]]; then
        cleanup_resources
        exit 1
    fi
fi
}

# Function to clean up resources
cleanup_resources() {
    echo "Cleaning up resources..." | tee -a "$LOG_FILE"
    echo "No persistent resources were created that need cleanup." | tee -a
"$LOG_FILE"
}

# Function to run a command in demo mode
demo_cmd() {
    local cmd="$1"
    local description="$2"

    {
        echo ""
        echo "DEMO: $description"
        echo "Command that would be executed:"
        echo "  [Command hidden for security]"
        echo ""
    } | tee -a "$LOG_FILE"
}

# Function to safely extract JSON values
extract_json_value() {
    local json_output="$1"
    local key="$2"

```

```

    echo "$json_output" | grep -o "\"$key\": \"[^\"]*\" | head -1 | cut -d'"' -
f4 || echo ""
}

# Array to track created resources
declare -a CREATED_RESOURCES

# Initialize demo mode flag
DEMO_MODE=false

# Security: Validate input parameters
if [[ $# -gt 0 ]]; then
    echo "ERROR: This script does not accept parameters" >&2
    exit 1
fi

{
    echo "====="
    echo "AWS Support CLI Tutorial"
    echo "====="
    echo "This script demonstrates how to use AWS Support API"
    echo "Note: You must have a Business, Enterprise On-Ramp,"
    echo "or Enterprise Support plan to use the AWS Support API."
    echo "====="
    echo ""
} | tee -a "$LOG_FILE"

# Step 1: Check available services
echo "Step 1: Checking available AWS Support services..." | tee -a "$LOG_FILE"
SERVICES_OUTPUT=$(log_cmd "aws support describe-services --language en" 2>&1) ||
SERVICES_OUTPUT=""
check_error "$SERVICES_OUTPUT" $? "Failed to retrieve AWS Support services"

# If we're in demo mode, set default values
if [[ "$DEMO_MODE" == "true" ]]; then
    SERVICE_CODE="general-info"
    echo "Using demo service code: $SERVICE_CODE" | tee -a "$LOG_FILE"
else
    # Extract a service code for demonstration using safer method
    SERVICE_CODE=$(extract_json_value "$SERVICES_OUTPUT" "code") ||
SERVICE_CODE=""
    if [[ -z "$SERVICE_CODE" ]]; then
        SERVICE_CODE="general-info"
        echo "Using default service code: $SERVICE_CODE" | tee -a "$LOG_FILE"
    fi
fi

```

```

        else
            echo "Found service code: $SERVICE_CODE" | tee -a "$LOG_FILE"
        fi
    fi

# Step 2: Check available severity levels
echo "Step 2: Checking available severity levels..." | tee -a "$LOG_FILE"
if [[ "$DEMO_MODE" == "true" ]]; then
    demo_cmd "aws support describe-severity-levels --language en" "Check
    available severity levels"
    SEVERITY_CODE="low"
    echo "Using demo severity code: $SEVERITY_CODE" | tee -a "$LOG_FILE"
else
    SEVERITY_OUTPUT=$(log_cmd "aws support describe-severity-levels --language
    en" 2>&1) || SEVERITY_OUTPUT=""
    check_error "$SEVERITY_OUTPUT" $? "Failed to retrieve severity levels"

    SEVERITY_CODE=$(extract_json_value "$SEVERITY_OUTPUT" "code") ||
    SEVERITY_CODE=""
    if [[ -z "$SEVERITY_CODE" ]]; then
        SEVERITY_CODE="low"
        echo "Using default severity code: $SEVERITY_CODE" | tee -a "$LOG_FILE"
    else
        echo "Found severity code: $SEVERITY_CODE" | tee -a "$LOG_FILE"
    fi
fi

# Step 3: Create a test support case
{
    echo ""
    echo "=====
    echo "SUPPORT CASE CREATION"
    echo "=====
} | tee -a "$LOG_FILE"

if [[ "$DEMO_MODE" == "true" ]]; then
{
    echo "DEMO MODE: The following steps would create and manage a support
    case"
    echo "if you had a Business, Enterprise On-Ramp, or Enterprise Support
    plan."
    echo ""
} | tee -a "$LOG_FILE"

```

```

USER_EMAIL="example@example.com"

demo_cmd "aws support create-case" "Create a support case"

CASE_ID="case-12345678910-2013-c4c1d2bf33c5cf47"
echo "Demo case ID: $CASE_ID" | tee -a "$LOG_FILE"

demo_cmd "aws support describe-cases" "List support cases"
demo_cmd "aws support add-communication-to-case" "Add communication to case"
demo_cmd "aws support describe-communications" "View case communications"
demo_cmd "aws support resolve-case" "Resolve the support case"

else
  echo "Creating a test support case..." | tee -a "$LOG_FILE"

  USER_EMAIL="example@example.com"
  CC_EMAIL_PARAM="--cc-email-addresses $USER_EMAIL"

  # Create the case
  CASE_OUTPUT=$(log_cmd "aws support create-case --subject 'AWS CLI Tutorial
Test Case' --service-code '$SERVICE_CODE' --category-code 'using-aws' --
communication-body 'This is a test case created as part of an AWS CLI tutorial.'
--severity-code '$SEVERITY_CODE' --language 'en' $CC_EMAIL_PARAM" 2>&1) ||
CASE_OUTPUT=""

  check_error "$CASE_OUTPUT" $? "Failed to create support case"

  # Extract the case ID safely
  CASE_ID=$(extract_json_value "$CASE_OUTPUT" "caseId") || CASE_ID=""

  if [[ -n "$CASE_ID" ]]; then
    echo "Successfully created support case with ID: $CASE_ID" | tee -a
"$LOG_FILE"
    CREATED_RESOURCES+=("Support Case: $CASE_ID")

    # Step 4: List the case we just created
    echo "" | tee -a "$LOG_FILE"
    echo "Step 4: Listing the support case we just created..." | tee -a
"$LOG_FILE"
    CASES_OUTPUT=$(log_cmd "aws support describe-cases --case-id-
list '$CASE_ID' --include-resolved-cases false --language 'en'" 2>&1) ||
CASES_OUTPUT=""

    check_error "$CASES_OUTPUT" $? "Failed to retrieve case details"

```



```

# Step 5: Add a communication to the case
echo "" | tee -a "$LOG_FILE"
echo "Step 5: Adding a communication to the support case..." | tee -a
"$LOG_FILE"
COMM_OUTPUT=$(log_cmd "aws support add-communication-to-case --case-id
'$CASE_ID' --communication-body 'This is an additional communication for the
test case.' $CC_EMAIL_PARAM" 2>&1) || COMM_OUTPUT=""

check_error "$COMM_OUTPUT" $? "Failed to add communication to case"

# Step 6: View communications for the case
echo "" | tee -a "$LOG_FILE"
echo "Step 6: Viewing communications for the support case..." | tee -a
"$LOG_FILE"
COMMS_OUTPUT=$(log_cmd "aws support describe-communications --case-id
'$CASE_ID' --language 'en'" 2>&1) || COMMS_OUTPUT=""

check_error "$COMMS_OUTPUT" $? "Failed to retrieve case communications"

# Step 7: Resolve the case
{
    echo ""
    echo "=====
    echo "CASE RESOLUTION"
    echo "=====
    echo "Resolving the support case..."
} | tee -a "$LOG_FILE"

RESOLVE_OUTPUT=$(log_cmd "aws support resolve-case --case-id '$CASE_ID'"
2>&1) || RESOLVE_OUTPUT=""

check_error "$RESOLVE_OUTPUT" $? "Failed to resolve case"
echo "Successfully resolved support case: $CASE_ID" | tee -a "$LOG_FILE"
else
    echo "Could not extract case ID from the response." | tee -a "$LOG_FILE"
fi
fi

# Display summary of created resources
{
    echo ""
    echo "=====
    echo "TUTORIAL SUMMARY"

```

```

    echo "======"
} | tee -a "$LOG_FILE"

if [[ "$DEMO_MODE" == "true" ]]; then
{
    echo "This was a demonstration in DEMO MODE."
    echo "No actual AWS Support cases were created."
    echo "To use the AWS Support API, you need a Business, Enterprise On-
Ramp,"
    echo "or Enterprise Support plan."
} | tee -a "$LOG_FILE"
else
{
    echo "Resources created during this tutorial:"
    if [[ ${#CREATED_RESOURCES[@]} -eq 0 ]]; then
        echo "No resources were created."
    else
        for resource in "${CREATED_RESOURCES[@]}${CREATED_RESOURCES[@]}";
do
            echo "- $resource"
        done
    fi
} | tee -a "$LOG_FILE"
fi

{
    echo ""
    echo "Tutorial completed successfully!"
    echo "Log file: $LOG_FILE"
    echo "======"
} | tee -a "$LOG_FILE"

# Display log file path to user
echo "Log file: $LOG_FILE"

```

- For API details, see the following topics in *AWS CLI Command Reference*.

- [AddCommunicationToCase](#)
- [CreateCase](#)
- [DescribeCases](#)
- [DescribeCommunications](#)
- [DescribeServices](#)

- [DescribeSeverityLevels](#)
- [ResolveCase](#)

For a complete list of AWS SDK developer guides and code examples, see [Using AWS Support with an AWS SDK](#). This topic also includes information about getting started and details about previous SDK versions.

Monitoring and logging for AWS Support

Monitoring is an important part of maintaining the reliability, availability, and performance of Support and your other AWS solutions. AWS provides the following monitoring tools to watch Support, report when something is wrong, and take automatic actions when appropriate:

- *Amazon EventBridge* delivers a near real-time stream of system events that describe changes in AWS resources. EventBridge enables automated event-driven computing, as you can write rules that watch for certain events and trigger automated actions in other AWS services when these events happen. For more information, see the [Amazon EventBridge User Guide](#).
- *AWS CloudTrail* captures API calls and related events made by or on behalf of your AWS account and delivers the log files to an Amazon S3 bucket that you specify. You can identify which users and accounts called AWS, the source IP address from which the calls were made, and when the calls occurred. For more information, see the [AWS CloudTrail User Guide](#).

Topics

- [Integrating AWS Support into event-driven applications using Amazon EventBridge](#)
- [Logging AWS Support API calls with AWS CloudTrail](#)
- [Logging AWS Support App in Slack API calls using AWS CloudTrail](#)

Integrating AWS Support into event-driven applications using Amazon EventBridge

You can incorporate AWS Support into event-driven applications (EDAs) that use events that occur in AWS Support to communicate between application components and initiate downstream processes.

For example, you can get notified whenever the following AWS Support events occur in your account:

- A support case is created, resolved, or reopened
- A correspondence is added to an existing support case


You do this by using Amazon EventBridge to route events from AWS Support to other software components. Amazon EventBridge is a serverless service that uses events to connect application components together, making it easier for you to integrate AWS services like AWS Support into event-driven architectures without additional code and operations.

How EventBridge routes AWS Support events

Here's how EventBridge works with AWS Support events:

As with many AWS services, AWS Support generates and sends events to the EventBridge default *event bus*. An event bus is a router that receives events and routes them to the destinations, or *targets*, that you specify. Targets can include other AWS services, custom applications, and SaaS partner applications.

EventBridge routes events according to *rules* you create on the event bus. For each rule, you specify a filter, or *event pattern*, to select only the events you want. Whenever an event is sent to the event bus, EventBridge compares it against each rule. If the event matches the rule, EventBridge routes the event to the specified target(s).



AWS Support events

AWS Support sends the following events to the default EventBridge event bus automatically.

Event detail type	Description
Support Case Update	Represents a change in a support case.

Event structure

All events from AWS services contain two types of data:

- A common set of fields containing metadata about the event, such as the AWS service that is the source of the event, the time the event was generated, the account and region in which the event took place, and others. For definitions of these general fields, see [Event structure](#) in the *Amazon EventBridge Events Reference*.
- A `detail` field that contains data specific to that particular service event.

AWS Support event delivery via AWS CloudTrail

AWS services can send events directly to the EventBridge default event bus. In addition, AWS CloudTrail sends events originating from numerous AWS services to EventBridge as well. These events can include API calls, console signins and actions, service events, and CloudTrail Insights. For more information, see [AWS service events delivered via AWS CloudTrail](#) in the *EventBridge User Guide*.

For a list of AWS Support events sent to EventBridge, refer to the AWS Support topic in the [EventBridge Events Reference](#).

Creating event patterns that match AWS Support events

Event patterns are filters where specify the data that the events you want to select should have.

Each event pattern is a JSON object that contains:

- A `source` attribute that identifies the service sending the event. For AWS Support events, the source is `aws.support`.
- (Optional): A `detail-type` attribute that contains an array of the event names to match.
- (Optional): A `detail` attribute containing any other event data on which to match.

For example, the following event pattern would select all Support Case Update events from AWS Support:

```
{
  "source": ["aws.support"],
```

```
"detail-type": ["Support Case Update"]
}
```

You can get more specific in your event selection by including values in the event itself. For example, the following event pattern matches Support Case Update events that represent a case being reopened:

```
{
  "source": ["aws.support"],
  "detail-type": ["Support Case Update"],
  "detail": {
    "event-name": "ReopenCase"
  }
}
```

For more information on writing event patterns, see [Event patterns](#) in the *EventBridge User Guide*.

See also

For more information about how to use EventBridge with AWS Support, see the following resources:

- [How to automate AWS Support API with Amazon EventBridge](#)
- [AWS Support case activity notifier](#) on GitHub

Support Case Update event

Below are the detail fields for the Support Case Update event.

The source and detail-type fields are included below because they contain specific values for AWS Support events. For definitions of the other metadata fields that are included in all events, see [Event structure](#) in the *Amazon EventBridge Events Reference*.

```
{
  . . . ,
  "detail-type": "Support Case Update",
  "source": "aws.support",
  . . . ,
  "detail": {
```

```
"case-id" : "string",  
"display-id" : "string",  
"communication-id" : "string",  
"event-name" : "string",  
"origin" : "string"  
}  
}
```

detail-type

Identifies the type of event.

For this event, this value is `Support Case Update`.

source

Identifies the service that generated the event. For AWS Support events, this value is `aws.support`.

detail

A JSON object that contains information about the event. The service generating the event determines the content of this field.

For this event, this data includes:

case-id

The support case ID. The case ID is an alphanumeric string in the following format:
`case-12345678910-2013-c4c1d2bf33c5cf47`.

display-id

The identifier for the case on pages in the AWS Support Center.

communication-id

The communication ID.

event-name

Valid values: `CreateCase` | `AddCommunicationToCase` | `ResolveCase` | `ReopenCase`

Specifies the type of support case event.

origin

Valid values: AWS | CUSTOMER

Specifies whether you or an AWS Support agent added a case correspondence to a support case.

Currently, only events with an event-name of `AddCommunicationToCase` will contain have this value.

Example Support Case Update event example: Support case created

```
{
  "version": "0",
  "id": "3433df007-9285-55a3-f6d1-536944be45d7",
  "detail-type": "Support Case Update",
  "source": "aws.support",
  "account": "111122223333",
  "time": "2022-02-21T15:51:19Z",
  "region": "us-east-1",
  "resources": [],
  "detail": {
    "case-id": "case-111122223333-muen-2022-7118885805350839",
    "display-id": "1234563851",
    "communication-id": "",
    "event-name": "CreateCase",
    "origin": ""
  }
}
```

Example Support Case Update event example: AWS Support replies to a support case

```
{
  "version": "0",
  "id": "f90cb8cb-32be-1c91-c0ba-d50b4ca5e51b",
  "detail-type": "Support Case Update",
  "source": "aws.support",
  "account": "111122223333",
  "time": "2022-02-21T15:51:31Z",
  "region": "us-east-1",
  "resources": [],
  "detail": {
```

```
    "case-id": "case-111122223333-muen-2022-7118885805350839",
    "display-id": "1234563851",
    "communication-id": "ekko:us-east-1:12345678-268a-424b-be08-54613cab84d2",
    "event-name": "AddCommunicationToCase",
    "origin": "AWS"
  }
}
```

Example Support Case Update event example: Support case resolved

```
{
  "version": "0",
  "id": "1aa4458d-556f-732e-ddc1-4a5b2fbd14a5",
  "detail-type": "Support Case Update",
  "source": "aws.support",
  "account": "111122223333",
  "time": "2022-02-21T15:51:31Z",
  "region": "us-east-1",
  "resources": [],
  "detail": {
    "case-id": "case-111122223333-muen-2022-7118885805350839",
    "display-id": "1234563851",
    "communication-id": "",
    "event-name": "ResolveCase",
    "origin": ""
  }
}
```

Example Support Case Update event example: Support case reopened

```
{
  "version": "0",
  "id": "3bb9d8fe-6089-ad27-9508-804209b233ad",
  "detail-type": "Support Case Update",
  "source": "aws.support",
  "account": "111122223333",
  "time": "2022-02-21T15:47:19Z",
  "region": "us-east-1",
  "resources": [],
  "detail": {
    "case-id": "case-111122223333-muen-2021-27f40618fe0303ea",
    "display-id": "1234563851",
    "communication-id": "",
  }
}
```

```
    "event-name": "ReopenCase",  
    "origin": ""  
  }  
}
```

Logging AWS Support API calls with AWS CloudTrail

AWS Support is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in AWS Support. CloudTrail captures API calls for AWS Support as events. The calls captured include calls from the AWS Support console and code calls to the AWS Support API operations.

If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon Simple Storage Service (Amazon S3) bucket, including events for AWS Support. 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 the request that was made to AWS Support, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, including how to configure and enable it, see the [AWS CloudTrail User Guide](#).

AWS Support information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When supported event activity occurs in AWS Support, 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 AWS Support, 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 AWS Regions. 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](#)

All AWS Support API operations are logged by CloudTrail and are documented in the [AWS Support API Reference](#).

For example, calls to the `CreateCase`, `DescribeCases` and `ResolveCase` operations generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail userIdentity element](#).

You can also aggregate AWS Support log files from multiple AWS Regions and multiple AWS accounts into a single Amazon S3 bucket.

AWS Trusted Advisor information in CloudTrail logging

Trusted Advisor is an AWS Support service that you can use to check your AWS account for ways to save costs, improve security, and optimize your account.

All Trusted Advisor API operations are logged by CloudTrail and are documented in the [AWS Support API Reference](#).

For example, calls to the `DescribeTrustedAdvisorCheckRefreshStatuses`, `DescribeTrustedAdvisorCheckResult` and `RefreshTrustedAdvisorCheck` operations generate entries in the CloudTrail log files.

Note

CloudTrail also logs Trusted Advisor console actions. See [Logging AWS Trusted Advisor console actions with AWS CloudTrail](#).

Understanding AWS Support 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. It includes information about the requested operation, the date and time of the operation, 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.

Example: Log entry for CreateCase

The following example shows a CloudTrail log entry for the [CreateCase](#) operation.

```
{
  "Records": [
    {
      "eventVersion": "1.04",
      "userIdentity": {
        "type": "IAMUser",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",
        "arn": "arn:aws:iam::111122223333:user/janedoe",
        "accountId": "111122223333",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "userName": "janedoe",
        "sessionContext": {
          "attributes": {
            "mfaAuthenticated": "false",
            "creationDate": "2016-04-13T17:51:37Z"
          }
        }
      },
      "invokedBy": "signin.amazonaws.com"
    },
    {
      "eventTime": "2016-04-13T18:05:53Z",
      "eventSource": "support.amazonaws.com",
      "eventName": "CreateCase",
      "awsRegion": "us-east-1",
      "sourceIPAddress": "198.51.100.15",
```

```

    "userAgent": "signin.amazonaws.com",
    "requestParameters": {
      "severityCode": "low",
      "categoryCode": "other",
      "language": "en",
      "serviceCode": "support-api",
      "issueType": "technical"
    },
    "responseElements": {
      "caseId": "case-111122223333-muen-2016-c3f2077e504940f2"
    },
    "requestID": "58c257ef-01a2-11e6-be2a-01c031063738",
    "eventID": "5aa34bfc-ad5b-4fb1-8a55-2277c86e746a",
    "eventType": "AwsApiCall",
    "recipientAccountId": "111122223333"
  }
],
...
}

```

Example: Log entry for RefreshTrustedAdvisorCheck

The following example shows a CloudTrail log entry for the [RefreshTrustedAdvisorCheck](#) operation.

```

{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:iam::111122223333:user/Admin",
    "accountId": "111122223333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "userName": "Admin"
  },
  "eventTime": "2020-10-21T16:34:13Z",
  "eventSource": "support.amazonaws.com",
  "eventName": "RefreshTrustedAdvisorCheck",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "72.21.198.67",
  "userAgent": "signin.amazonaws.com",
  "requestParameters": {
    "checkId": "Pfx0RwqBli"
  },
  "responseElements": null,
}

```

```
"requestID": "4c4d5fc8-c403-4f82-9544-41f820e0fa01",  
"eventID": "2f4630ac-5c27-4f0d-b93f-63742d6fc85e",  
"eventType": "AwsApiCall",  
"recipientAccountId": "111122223333"  
}
```

Logging AWS Support App in Slack API calls using AWS CloudTrail

The AWS Support App in Slack is integrated with AWS CloudTrail. CloudTrail provides a record of actions taken by a user, role, or an AWS service in the AWS Support App. To create this record, CloudTrail captures all public API calls for AWS Support App as events. These captured calls include calls from the AWS Support App console, and code calls to the AWS Support App public API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket. These include events for AWS Support App. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. You can use the information that CloudTrail collects to determine that the request that was made to AWS Support App. You can also learn the IP address where the call originated, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the [AWS CloudTrail User Guide](#).

AWS Support App information in CloudTrail

When you create your AWS account, this activates CloudTrail on the account. When public API activity occurs in the AWS Support App, 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 AWS Support App, create a *trail*. By default, when you create a trail in the console, the trail applies to all AWS Regions. 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 analyze further the event data collected in CloudTrail logs and act upon the data. 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](#)

CloudTrail logs all public AWS Support App actions. These actions are also documented in the [AWS Support App in Slack API Reference](#). For example, calls to the `CreateSlackChannelConfiguration`, `GetAccountAlias` and `UpdateSlackChannelConfiguration` actions generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail `userIdentity` element](#).

Understanding AWS Support App 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. This means that the logs don't appear in any specific order.

Example: Log example for `CreateSlackChannelConfiguration`

The following example shows a CloudTrail log entry for the [CreateSlackChannelConfiguration](#) operation.

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE:JaneDoe",
    "arn": "arn:aws:sts::111122223333:assumed-role/Administrator/JaneDoe",
    "accountId": "111122223333",
```



```

    "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",
        "arn": "arn:aws:iam::111122223333:role/Administrator",
        "accountId": "111122223333",
        "userName": "Administrator"
      },
      "webIdFederationData": {},
      "attributes": {
        "creationDate": "2022-02-26T01:37:57Z",
        "mfaAuthenticated": "false"
      }
    },
    "eventTime": "2022-02-26T01:48:20Z",
    "eventSource": "supportapp.amazonaws.com",
    "eventName": "CreateSlackChannelConfiguration",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "205.251.233.183",
    "userAgent": "aws-cli/1.3.23 Python/2.7.6 Linux/2.6.18-164.el5",
    "requestParameters": {
      "notifyOnCreateOrReopenCase": true,
      "teamId": "T012ABCDEFGF",
      "notifyOnAddCorrespondenceToCase": true,
      "notifyOnCaseSeverity": "all",
      "channelName": "troubleshooting-channel",
      "notifyOnResolveCase": true,
      "channelId": "C01234A5BCD",
      "channelRoleArn": "arn:aws:iam::111122223333:role/AWSSupportAppRole"
    },
    "responseElements": null,
    "requestID": "d06df6ca-c233-4ffb-bbff-63470c5dc255",
    "eventID": "0898ce29-a396-444a-899d-b068f390c361",
    "readOnly": false,
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "eventCategory": "Management"
  }
}

```

Example: Log example for ListSlackChannelConfigurations

The following example shows a CloudTrail log entry for the [ListSlackChannelConfigurations](#) operation.

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE:AWSSupportAppRole",
    "arn": "arn:aws:sts::111122223333:assumed-role/AWSSupportAppRole",
    "accountId": "111122223333",
    "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",
        "arn": "arn:aws:iam::111122223333:role/AWSSupportAppRole",
        "accountId": "111122223333",
        "userName": "AWSSupportAppRole"
      },
      "webIdFederationData": {},
      "attributes": {
        "creationDate": "2022-03-01T20:06:32Z",
        "mfaAuthenticated": "false"
      }
    }
  },
  "eventTime": "2022-03-01T20:06:46Z",
  "eventSource": "supportapp.amazonaws.com",
  "eventName": "ListSlackChannelConfigurations",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "72.21.217.131",
  "userAgent": "aws-cli/1.3.23 Python/2.7.6 Linux/2.6.18-164.el5",
  "requestParameters": null,
  "responseElements": null,
  "requestID": "20f81d63-31c5-4351-bd02-9eda7f76e7b8",
  "eventID": "70acb7fe-3f84-47cd-8c28-cc148ad06d21",
  "readOnly": true,
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "111122223333",
  "eventCategory": "Management"
}
```

```
}
```

Example: Log example for GetAccountAlias

The following example shows a CloudTrail log entry for the [GetAccountAlias](#) operation.

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE:devdsk",
    "arn": "arn:aws:sts::111122223333:assumed-role/AWSSupportAppRole/devdsk",
    "accountId": "111122223333",
    "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",
        "arn": "arn:aws:iam::111122223333:role/AWSSupportAppRole",
        "accountId": "111122223333",
        "userName": "AWSSupportAppRole"
      },
      "webIdFederationData": {},
      "attributes": {
        "creationDate": "2022-03-01T20:31:27Z",
        "mfaAuthenticated": "false"
      }
    }
  },
  "eventTime": "2022-03-01T20:31:47Z",
  "eventSource": "supportapp.amazonaws.com",
  "eventName": "GetAccountAlias",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "72.21.217.142",
  "userAgent": "aws-cli/1.3.23 Python/2.7.6 Linux/2.6.18-164.el5",
  "requestParameters": null,
  "responseElements": null,
  "requestID": "a225966c-0906-408b-b8dd-f246665e6758",
  "eventID": "79ebba8d-3285-4023-831a-64af7de8d4ad",
  "readOnly": true,
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "111122223333",
  "eventCategory": "Management"
}
```

```
}
```

Monitoring and logging for AWS Support Plans

Monitoring is an important part of maintaining the reliability, availability, and performance of Support Plans and your other AWS solutions. AWS provides the following monitoring tools to watch Support Plans, report when something is wrong, and take automatic actions when appropriate:

- *AWS CloudTrail* captures API calls and related events made by or on behalf of your AWS account and delivers the log files to an Amazon S3 bucket that you specify. You can identify which users and accounts called AWS, the source IP address from which the calls were made, and when the calls occurred. For more information, see the [AWS CloudTrail User Guide](#).

Topics

- [Logging AWS Support Plans API calls with AWS CloudTrail](#)

Logging AWS Support Plans API calls with AWS CloudTrail

AWS Support Plans is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service. CloudTrail captures API calls for AWS Support Plans as events. The calls captured include calls from the AWS Support Plans console and code calls to the AWS Support Plans API operations.

If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon Simple Storage Service (Amazon S3) bucket, including events for AWS Support Plans. 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 the request that was made to AWS Support Plans, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, including how to configure and enable it, see the [AWS CloudTrail User Guide](#).

AWS Support Plans information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When supported event activity occurs in AWS Support Plans, 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 account. For more information, see [Viewing events with CloudTrail event history](#).

For an ongoing record of events in your account, including events for AWS Support Plans, 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 AWS Regions. 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](#)

All AWS Support Plans API operations are logged by CloudTrail. Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail userIdentity element](#).

You can also aggregate AWS Support Plans log files from multiple AWS Regions and multiple accounts into a single Amazon S3 bucket.

Understanding AWS Support Plans 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. It includes information about the requested operation, the date and time of the operation, 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.

Example: Log entry for GetSupportPlan

The following example shows a CloudTrail log entry for the GetSupportPlan operation.

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:sts::111122223333:user/janedoe",
    "accountId": "111122223333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",
        "arn": "arn:aws:iam::111122223333:role/Admin",
        "accountId": "111122223333",
        "userName": "Admin"
      },
      "webIdFederationData": {},
      "attributes": {
        "creationDate": "2022-06-29T16:30:04Z",
        "mfaAuthenticated": "false"
      }
    }
  },
  "eventTime": "2022-06-29T16:39:11Z",
  "eventSource": "supportplans.amazonaws.com",
  "eventName": "GetSupportPlan",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "205.251.233.183",
  "userAgent": "Mozilla/5.0 (Windows NT 10.0; Win64; x64; rv:91.0) Gecko/20100101 Firefox/91.0",
  "requestParameters": null,
```

```

    "responseElements": null,
    "requestID": "7665c39a-d6bf-4d0d-8010-2f59740b8ecb",
    "eventID": "b711bc30-16a5-4579-8f0d-9ada8fe6d1ce",
    "readOnly": true,
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "eventCategory": "Management"
  }

```

Example: Log entry for GetSupportPlanUpdateStatus

The following example shows a CloudTrail log entry for the GetSupportPlanUpdateStatus operation.

```

{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:sts::111122223333:user/janedoe",
    "accountId": "111122223333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",
        "arn": "arn:aws:iam::111122223333:role/Admin",
        "accountId": "111122223333",
        "userName": "Admin"
      },
      "webIdFederationData": {},
      "attributes": {
        "creationDate": "2022-06-29T16:30:04Z",
        "mfaAuthenticated": "false"
      }
    }
  },
  "eventTime": "2022-06-29T16:39:02Z",
  "eventSource": "supportplans.amazonaws.com",
  "eventName": "GetSupportPlanUpdateStatus",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "205.251.233.183",

```



```

    "userAgent": "Mozilla/5.0 (Windows NT 10.0; Win64; x64; rv:91.0) Gecko/20100101
    Firefox/91.0",
    "requestParameters": {
        "supportPlanUpdateArn":
        "arn:aws:supportplans::111122223333:supportplanupdate/7f03b7a233a0e87ebc79e56d4d2bcaf19e976c37
    },
    "responseElements": null,
    "requestID": "75e5c767-8703-4ed3-b01e-4dda28020322",
    "eventID": "28d1c0e3-ccb6-4fd1-8793-65be010114cc",
    "readOnly": true,
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "eventCategory": "Management"
}

```

Example: Log entry for StartSupportPlanUpdate

The following example shows a CloudTrail log entry for the StartSupportPlanUpdate operation.

```

{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",
        "arn": "arn:aws:sts::111122223333:user/janedoe",
        "accountId": "111122223333",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "sessionContext": {
            "sessionIssuer": {
                "type": "Role",
                "principalId": "AIDACKCEVSQ6C2EXAMPLE",
                "arn": "arn:aws:iam::111122223333:role/Admin",
                "accountId": "111122223333",
                "userName": "Admin"
            },
            "webIdFederationData": {},
            "attributes": {
                "creationDate": "2022-06-29T16:30:04Z",
                "mfaAuthenticated": "false"
            }
        }
    }
}

```

```

    },
    "eventTime": "2022-06-29T16:38:55Z",
    "eventSource": "supportplans.amazonaws.com",
    "eventName": "StartSupportPlanUpdate",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "205.251.233.183",
    "userAgent": "Mozilla/5.0 (Windows NT 10.0; Win64; x64; rv:91.0) Gecko/20100101
Firefox/91.0",
    "requestParameters": {
      "clientToken": "98add111-dcc9-464d-8722-438d697fe242",
      "update": {
        "supportLevel": "BASIC"
      }
    },
    "responseElements": {
      "Access-Control-Expose-Headers": "x-amzn-RequestId,x-amzn-ErrorType,x-amzn-
ErrorMessage,Date",
      "supportPlanUpdateArn":
"arn:aws:supportplans::111122223333:supportplanupdate/7f03b7a233a0e87ebc79e56d4d2bcaf19e976c37",
    },
    "requestID": "e5ff9382-5fb8-4764-9993-0f33fb0b1e17",
    "eventID": "5dba89f8-2e5b-42b9-9b8f-395580c52962",
    "readOnly": false,
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "eventCategory": "Management"
  }
}

```

Example: Log entry for CreateSupportPlanSchedule

The following example shows a CloudTrail log entry for the CreateSupportPlanSchedule operation.

```

{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:sts::111122223333:user/janedoe",
    "accountId": "111122223333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {

```

```

    "sessionIssuer": {
      "type": "Role",
      "principalId": "AIDACKCEVSQ6C2EXAMPLE",
      "arn": "arn:aws:iam::111122223333:role/Admin",
      "accountId": "111122223333",
      "userName": "Admin"
    },
    "webIdFederationData": {},
    "attributes": {
      "creationDate": "2023-05-09T16:30:04Z",
      "mfaAuthenticated": "false"
    }
  }
},
"eventTime": "2023-05-09T16:30:04Z",
"eventSource": "supportplans.amazonaws.com",
"eventName": "CreateSupportPlanSchedule",
"awsRegion": "us-west-2",
"sourceIPAddress": "205.251.233.183",
"userAgent": "Mozilla/5.0 (Windows NT 10.0; Win64; x64; rv:91.0) Gecko/20100101
Firefox/91.0",
"requestParameters": {
  "clientToken": "b998de5e-ad1c-4448-90db-2bf86d6d9e9a",
  "scheduleCreationDetails": {
    "startLevel": "BUSINESS",
    "startOffer": "TrialPlan7FB93B",
    "startTimestamp": "2023-06-03T17:23:56.109Z",
    "endLevel": "BUSINESS",
    "endOffer": "StandardPlan2074BB",
    "endTimestamp": "2023-09-03T17:23:55.109Z"
  }
},
"responseElements": {
  "Access-Control-Expose-Headers": "x-amzn-RequestId,x-amzn-ErrorType,x-amzn-
ErrorMessage,Date",
  "supportPlanUpdateArn":
  "arn:aws:supportplans::111122223333:supportplanschedule/
b9a9a4336a3974950a6e670f7dab79b77a4b104db548a0d57050ce4544721d4b"
},
"requestID": "150450b8-e61a-4b15-93a8-c3b557a1ca48",
"eventID": "a2a1ba44-610d-4dc8-bf16-29f1635b57a9",
"readOnly": false,
"eventType": "AwsApiCall",
"managementEvent": true,

```

```
"recipientAccountId": "111122223333",  
"eventCategory": "Management"  
}
```

Example: Log entry for ListSupportPlanModifiers

The following example shows a CloudTrail log entry for the ListSupportPlanModifiers operation.

```
{  
  "eventVersion": "1.09",  
  "userIdentity": {  
    "type": "AssumedRole",  
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",  
    "arn": "arn:aws:sts::111122223333:user/janedoe",  
    "accountId": "111122223333",  
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",  
    "sessionContext": {  
      "sessionIssuer": {  
        "type": "Role",  
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",  
        "arn": "arn:aws:sts::111122223333:user/janedoe",  
        "accountId": "111122223333",  
        "userName": "Admin"  
      },  
      "attributes": {  
        "creationDate": "2024-08-15T15:44:43Z",  
        "mfaAuthenticated": "false"  
      }  
    },  
  },  
  "eventTime": "2024-08-15T16:29:59Z",  
  "eventSource": "supportplans.amazonaws.com",  
  "eventName": "ListSupportPlanModifiers",  
  "awsRegion": "us-west-2",  
  "sourceIPAddress": "205.251.233.183",  
  "userAgent": "Mozilla/5.0 (Windows NT 10.0; Win64; x64; rv:91.0) Gecko/20100101  
Firefox/91.0",  
  "requestParameters": null,  
  "responseElements": null,  
  "requestID": "7665c39a-d6bf-4d0d-8010-2f59740b8ecb",  
  "eventID": "b711bc30-16a5-4579-8f0d-9ada8fe6d1ce",  
  "readOnly": true,  
}
```

```
"eventType": "AwsApiCall",
"managementEvent": true,
"recipientAccountId": "111122223333",
"eventCategory": "Management"
}
```

Logging changes to your Support plan

Important

As of August 3, 2022, the following operations are deprecated and won't appear in your new CloudTrail logs. For a list of supported operations, see [Understanding AWS Support Plans log file entries](#).

- `DescribeSupportLevelSummary` – This action appears in your log when you open the [Support plans](#) page.
- `UpdateProbationAutoCancellation` – After you sign up for Developer Support or Business Support and then try to cancel within 30 days, your plan will be automatically canceled at the end of that period. This action appears in your log when you choose **Opt-out of automatic cancellation** in the banner that appears on the [Support plans](#) page. You will resume your plan for Developer Support or Business Support.
- `UpdateSupportLevel` – This action appears in your log when you change your support plan.

Note

The `eventSource` field has the `support-subscription.amazonaws.com` namespace for these actions.

Example: Log entry for `DescribeSupportLevelSummary`

The following example shows a CloudTrail log entry for the `DescribeSupportLevelSummary` action.

```
{
  "eventVersion": "1.08",
  "userIdentity": {
```

```

    "type": "Root",
    "principalId": "111122223333",
    "arn": "arn:aws:iam::111122223333:root",
    "accountId": "111122223333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "sessionIssuer": {},
      "webIdFederationData": {},
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2021-01-07T22:08:05Z"
      }
    }
  },
  "eventTime": "2021-01-07T22:08:07Z",
  "eventSource": "support-subscription.amazonaws.com",
  "eventName": "DescribeSupportLevelSummary",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "100.127.8.67",
  "userAgent": "AWS-SupportPlansConsole, aws-internal/3",
  "requestParameters": {
    "lang": "en"
  },
  "responseElements": null,
  "requestID": "b423b84d-829b-4090-a239-2b639b123abc",
  "eventID": "e1eeda0e-d77c-487b-a7e5-4014f7123abc",
  "readOnly": true,
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "eventCategory": "Management",
  "recipientAccountId": "111122223333"
}

```

Example: Log entry for UpdateProbationAutoCancellation

The following example shows a CloudTrail log entry for the UpdateProbationAutoCancellation action.

```

{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "Root",
    "principalId": "111122223333",

```

```

    "arn": "arn:aws:iam::111122223333:root",
    "accountId": "111122223333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2021-01-07T23:28:43Z",
  "eventSource": "support-subscription.amazonaws.com",
  "eventName": "UpdateProbationAutoCancellation",
  "awsRegion": "us-east-1", "sourceIPAddress": "100.127.8.67",
  "userAgent": "AWS-SupportPlansConsole, aws-internal/3",
  "requestParameters": {
    "lang": "en"
  },
  "responseElements": null,
  "requestID": "5492206a-e200-4c33-9fcf-4162d4123abc",
  "eventID": "f4a58c09-0bb0-4ba2-a8d3-df6909123abc",
  "readOnly": false,
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "eventCategory": "Management",
  "recipientAccountId": "111122223333"
}

```

Example: Log entry for UpdateSupportLevel

The following example shows a CloudTrail log entry for the UpdateSupportLevel action to change to Developer Support.

```

{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "Root",
    "principalId": "111122223333",
    "arn": "arn:aws:iam::111122223333:root",
    "accountId": "111122223333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "sessionIssuer": {},
      "webIdFederationData": {},
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2021-01-07T22:08:05Z"
      }
    }
  }
}

```

```
{,
  "eventTime": "2021-01-07T22:08:43Z",
  "eventSource": "support-subscription.amazonaws.com",
  "eventName": "UpdateSupportLevel",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "100.127.8.247",
  "userAgent": "AWS-SupportPlansConsole, aws-internal/3",
  "requestParameters": {
    "supportLevel": "new_developer"
  },
  "responseElements": {
    "aispl": false,
    "supportLevel": "new_developer"
  },
  "requestID": "5df3da3a-61cd-4a3c-8f41-e5276b123abc",
  "eventID": "c69fb149-c206-47ce-8766-8df6ec123abc",
  "readOnly": false,
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "eventCategory": "Management",
  "recipientAccountId": "111122223333"
}
```


Monitoring and logging for AWS Trusted Advisor

Monitoring is an important part of maintaining the reliability, availability, and performance of Trusted Advisor and your other AWS solutions. AWS provides the following monitoring tools to watch Trusted Advisor, report when something is wrong, and take automatic actions when appropriate:

- *Amazon EventBridge* delivers a near real-time stream of system events that describe changes in AWS resources. EventBridge enables automated event-driven computing, as you can write rules that watch for certain events and trigger automated actions in other AWS services when these events happen.

For example, Trusted Advisor provides the **Amazon S3 Bucket Permissions** check. This check identifies if you have buckets that have open access permissions or allow access to any authenticated AWS user. If a bucket permission changes, the status changes for the Trusted Advisor check. EventBridge detects this event and then sends you a notification so that you can take action. For more information, see the [Amazon EventBridge User Guide](#).

- AWS Trusted Advisor checks identify ways for you to reduce cost, increase performance, and improve security for your AWS account. You can use EventBridge to monitor the status of Trusted Advisor checks. You can then use Amazon CloudWatch to create alarms on Trusted Advisor metrics. These alarms notify you when the status changes for a Trusted Advisor check, such as an updated resource or a service quota that is reached.
- *AWS CloudTrail* captures API calls and related events made by or on behalf of your AWS account and delivers the log files to an Amazon S3 bucket that you specify. You can identify which users and accounts called AWS, the source IP address from which the calls were made, and when the calls occurred. For more information, see the [AWS CloudTrail User Guide](#).

Topics

- [Monitoring AWS Trusted Advisor check results with Amazon EventBridge](#)
- [Creating Amazon CloudWatch alarms to monitor AWS Trusted Advisor metrics](#)
- [Logging AWS Trusted Advisor console actions with AWS CloudTrail](#)

Monitoring AWS Trusted Advisor check results with Amazon EventBridge

Important

End of Support Notice: Developer Support will be discontinued January 1, 2027. Customers with Developer Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information, see [Business Support+ plan details](#)

End of Support Notice: Business Support will be discontinued January 1, 2027. Customers with Business Support can continue using their existing plan or choose to upgrade to Business Support+ anytime before January 1, 2027. Business Support+ delivers AI-powered assistance that understands the context of your operations, with 24/7 access to AWS experts at \$29/month minimum per account. For more information see, [Business Support+ plan details](#)

End of Support Notice: On January 1, 2027, AWS will discontinue Enterprise On-Ramp. Throughout 2026, Enterprise On-Ramp customers will be automatically upgraded to AWS Enterprise Support during contract renewal or in periodic batches. Customers will receive an email notification a month before their upgrade. No further action is required. Enterprise Support provides designated TAM assignment, 15-minute response times, and AWS Security Incident Response available at no additional cost, all at a lower \$5,000 minimum (reduced from \$15,000). For more information, see [AWS Enterprise Support plan details](#).

For more information, see [Developer, Business, and Enterprise On-Ramp end of support](#). Developer Support, Business Support, and Enterprise On-Ramp will remain available in the AWS GovCloud (US) Region.

You can use EventBridge to detect when your checks for Trusted Advisor change status. Then, based on the rules that you create, EventBridge invokes one or more target actions when the status changes to a value that you specify in a rule.

Depending on the status change, you can send notifications, capture status information, take corrective action, initiate events, or take other actions. For example, you can specify the following

target types if a check changes status from no problems detected (green) to recommended action (red).

- Use an AWS Lambda function to pass a notification to a Slack channel.
- Push data about the check to an Amazon Kinesis stream to support comprehensive and real-time status monitoring.
- Send an Amazon Simple Notification Service topic to your email.
- Get notified with an Amazon CloudWatch alarm action.

For more information about on how to use EventBridge and Lambda functions to automate responses for Trusted Advisor, see [Trusted Advisor tools](#) in GitHub.

Notes

- Trusted Advisor delivers events on a best effort basis. Events are not always guaranteed to be delivered to EventBridge.
- You must have an AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations plan to create a rule for Trusted Advisor checks. For more information, see [Change AWS Support Plans](#).
-
- As Trusted Advisor is a Global service, all Events are emitted to EventBridge in the US East (N. Virginia) Region.

Follow this procedure to create an EventBridge rule for Trusted Advisor. Before you create event rules, do the following:

- Familiarize yourself with events, rules, and targets in EventBridge. For more information, see [What is Amazon EventBridge?](#) in the *Amazon EventBridge User Guide*.
- Create the target that you will use in your event rule.

To create an EventBridge rule for Trusted Advisor

1. Open the Amazon EventBridge console at <https://console.aws.amazon.com/events/>.

2. To change the Region, use the **Region selector** in the upper-right corner of the page and choose **US East (N. Virginia)**.
3. In the navigation pane, choose **Rules**.
4. Choose **Create rule**.
5. On the **Define rule detail** page, enter a name and description for your rule.
6. Keep the default values for **Event bus** and **Rule type**, and then choose **Next**.
7. On the **Build event pattern** page, for **Event source**, choose **AWS events or EventBridge partner events**.
8. Under **Event pattern**, keep the default value for **AWS services**.
9. For **AWS service**, choose **Trusted Advisor**.
10. For **Event type**, choose **Check Item Refresh Status**.
11. Choose one of the following options for check statuses:
 - Choose **Any status** to create a rule that monitors for any status change.
 - Choose **Specific status(es)**, and then choose the values that you want your rule to monitor.
 - **ERROR** – Trusted Advisor recommends an action for the check.
 - **INFO** – Trusted Advisor can't determine the status of the check.
 - **OK** – Trusted Advisor doesn't detect an issue for the check.
 - **WARN** – Trusted Advisor detects a possible issue for the check and recommends investigation.
12. Choose one of the following options for your checks:
 - Choose **Any check**.
 - Choose **Specific check(s)**, and then choose one or more check names from the list.
13. Choose one of the following options for AWS resources:
 - Choose **Any resource ID** to create a rule that monitors all resources.
 - Choose **Specific resource ID(s) by ARN**, and then enter the Amazon Resource Names (ARNs) that you want.
14. Choose **Next**.

15. In the **Select target(s)** page, choose the target type that you created for this rule, and then configure any additional options that are required for that type. For example, you might send the event to an Amazon SQS queue or an Amazon SNS topic.
16. Choose **Next**.
17. (Optional) On the **Configure tags** page, add any tags and then choose **Next**.
18. On the **Review and create** page, review your rule setup and ensure that it meets your event monitoring requirements.
19. Choose **Create rule**. Your rule will now monitor for Trusted Advisor checks and then send the event to the target that you specified.

Creating Amazon CloudWatch alarms to monitor AWS Trusted Advisor metrics

When AWS Trusted Advisor refreshes your checks, Trusted Advisor publishes metrics about your check results to CloudWatch. You can view the metrics in CloudWatch. You can also create alarms to detect status changes to Trusted Advisor checks and status changes for resources, and service quota usage (formerly referred to as limits).

Follow this procedure to create a CloudWatch alarm for a specific Trusted Advisor metric.

Topics

- [Prerequisites](#)
- [CloudWatch metrics for Trusted Advisor](#)
- [Trusted Advisor metrics and dimensions](#)

Prerequisites

Before you create CloudWatch alarms for Trusted Advisor metrics, review the following information:

- Understand how CloudWatch uses metrics and alarms. For more information, see [How CloudWatch works](#) in the *Amazon CloudWatch User Guide*.
- Use the Trusted Advisor console or the AWS Support API to refresh your checks and get the latest check results. For more information, see [Refresh check results](#).

To create a CloudWatch alarm for Trusted Advisor metrics

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Use the **Region selector** and choose the **US East (N. Virginia)** AWS Region.
3. In the navigation pane, choose **Alarms**.
4. Choose **Create alarm**.
5. Choose **Select metric**.
6. For **Metrics**, enter one or more dimension values to filter the metric list. For example, you can enter the metric name **ServiceLimitUsage** or the dimension, such as the Trusted Advisor check name.

Tip

- You can search for **Trusted Advisor** to list all metrics for the service.
- For a list of metric and dimension names, see [Trusted Advisor metrics and dimensions](#).

7. In the results table, select the check box for the metric.

In the following example, the check name is **IAM Access Key Rotation** and the metric name is **YellowResources**.

N. Virginia ▾		All > TrustedAdvisor > Check Metrics		Trusted ✕	Advisor ✕	IAM ✕	Access ✕	Key ✕
☐	CheckName (2)	Metric Name						
☐	IAM Access Key Rotation	RedResources						
☒	IAM Access Key Rotation	YellowResources						

8. Choose **Select metric**.
9. On the **Specify metric and conditions** page, verify that the **Metric name** and **CheckName** that you chose appear on the page.
10. For **Period**, you can specify the time period that you want the alarm to start when the check status changes, such as 5 minutes.
11. Under **Conditions**, choose **Static**, and then specify the alarm condition for when the alarm should start.

For example, if you choose **Greater/Equal $\geq$ threshold** and enter **1** for the threshold value, this means that the alarm starts when Trusted Advisor detects at least one IAM access key that hasn't been rotated in the last 90 days.

Notes

- For the **GreenChecks**, **RedChecks**, **YellowChecks**, **RedResources**, and **YellowResources** metrics, you can specify a threshold that is any whole number greater than or equal to zero.
- Trusted Advisor doesn't send metrics for **GreenResources**, which are resources for which Trusted Advisor hasn't detected any issues.

12. Choose **Next**.
13. On the **Configure actions** page, for **Alarm state trigger**, choose **In alarm**.
14. For **Select an SNS topic**, choose an existing Amazon Simple Notification Service (Amazon SNS) topic or create one.

Notification

Alarm state trigger
Define the alarm state that will trigger this action.

☒ **In alarm**
The metric or expression is outside of the defined threshold.

☐ **OK**
The metric or expression is within the defined threshold.

☐ **Insufficient data**
The alarm has just started or not enough data is available.

Remove

Select an SNS topic
Define the SNS (Simple Notification Service) topic that will receive the notification.

☒ **Select an existing SNS topic**

☐ Create new topic

☐ Use topic ARN

Send a notification to...

X

Only email lists for this account are available.

Email (endpoints)
janedoe@example.com - [View in SNS Console](#)

Add notification

15. Choose **Next**.

16. For **Name and description**, enter a name and description for your alarm.

17. Choose **Next**.

18. On the **Preview and create** page, review your alarm details, and then choose **Create alarm**.

When the status for the **IAM Access Key Rotation** check changes to red for 5 minutes, your alarm will send a notification to your SNS topic.

Example: Email notification for a CloudWatch alarm

The following email message shows that an alarm detected a change for the **IAM Access Key Rotation** check.

You are receiving this email because your Amazon CloudWatch Alarm "IAMAccessKeyRotationCheckAlarm" in the US East (N. Virginia) region has entered the ALARM state, because "Threshold Crossed: 1 out of the last 1 datapoints [9.0 (26/03/21 22:44:00)] was greater than or equal to the threshold (1.0) (minimum 1 datapoint for OK -> ALARM transition)." at "Friday 26 March, 2021 22:49:42 UTC".

View this alarm in the AWS Management Console:

<https://us-east-1.console.aws.amazon.com/cloudwatch/home?region=us-east-1#s=Alarms&alarm=IAMAccessKeyRotationCheckAlarm>

Alarm Details:

- Name: IAMAccessKeyRotationCheckAlarm
- Description: This alarm starts when one or more AWS access keys in my AWS account have not been rotated in the last 90 days.
- State Change: INSUFFICIENT_DATA -> ALARM
- Reason for State Change: Threshold Crossed: 1 out of the last 1 datapoints [9.0 (26/03/21 22:44:00)] was greater than or equal to the threshold (1.0) (minimum 1 datapoint for OK -> ALARM transition).
- Timestamp: Friday 26 March, 2021 22:49:42 UTC
- AWS Account: 123456789012
- Alarm Arn: arn:aws:cloudwatch:us-east-1:123456789012:alarm:IAMAccessKeyRotationCheckAlarm

Threshold:

- The alarm is in the ALARM state when the metric is GreaterThanOrEqualToThreshold 1.0 for 300 seconds.

Monitored Metric:

- MetricNamespace: AWS/TrustedAdvisor
- MetricName: RedResources
- Dimensions: [CheckName = IAM Access Key Rotation]
- Period: 300 seconds
- Statistic: Average
- Unit: not specified
- TreatMissingData: missing

State Change Actions:

- OK:
- ALARM: [arn:aws:sns:us-east-1:123456789012:Default_CloudWatch_Alarms_Topic]
- INSUFFICIENT_DATA:

CloudWatch metrics for Trusted Advisor

You can use the CloudWatch console or the AWS Command Line Interface (AWS CLI) to find the metrics available for Trusted Advisor.

For a list of the namespaces, metrics, and dimensions for all services that publish metrics, see [AWS services that publish CloudWatch metrics](#) in the *Amazon CloudWatch User Guide*.


View Trusted Advisor metrics (console)

You can sign in to the CloudWatch console and view the available metrics for Trusted Advisor.

To view available Trusted Advisor metrics (console)

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Use the **Region selector** and choose the **US East (N. Virginia)** AWS Region.
3. In the navigation pane, choose **Metrics**.
4. Enter a metric namespace, such as **TrustedAdvisor**.
5. Choose a metric dimension, such as **Check Metrics**.
6. The **All metrics** tab shows metrics for that dimension in the namespace. You can do the following:
 - a. To sort the table, choose the column heading.
 - b. To graph a metric, select the check box next to the metric. To select all metrics, select the check box in the heading row of the table.
 - c. To filter by metric, choose the metric name, and then choose **Add to search**.

The following example shows the results for the **Security Groups - Specific Ports Unrestricted** check. The check identifies 13 resources that are yellow. Trusted Advisor recommends that you investigate checks that are yellow.



- (Optional) To add this graph to a CloudWatch dashboard, choose **Actions**, and then choose **Add to dashboard**.

For more information about creating a graph to view your metrics, see [Graphing a metric](#) in the *Amazon CloudWatch User Guide*.

View Trusted Advisor metrics (CLI)

You can use the [list-metrics](#) AWS CLI command to view available metrics for Trusted Advisor.

Example: List all metrics for Trusted Advisor

The following example specifies the `AWS/TrustedAdvisor` namespace to view all metrics for Trusted Advisor.

```
aws cloudwatch list-metrics --namespace AWS/TrustedAdvisor
```

Your output might look like the following.

```
{
```

```

"Metrics": [
  {
    "Namespace": "AWS/TrustedAdvisor",
    "Dimensions": [
      {
        "Name": "ServiceName",
        "Value": "EBS"
      },
      {
        "Name": "ServiceLimit",
        "Value": "Magnetic (standard) volume storage (TiB)"
      },
      {
        "Name": "Region",
        "Value": "ap-northeast-2"
      }
    ],
    "MetricName": "ServiceLimitUsage"
  },
  {
    "Namespace": "AWS/TrustedAdvisor",
    "Dimensions": [
      {
        "Name": "CheckName",
        "Value": "Overutilized Amazon EBS Magnetic Volumes"
      }
    ],
    "MetricName": "YellowResources"
  },
  {
    "Namespace": "AWS/TrustedAdvisor",
    "Dimensions": [
      {
        "Name": "ServiceName",
        "Value": "EBS"
      },
      {
        "Name": "ServiceLimit",
        "Value": "Provisioned IOPS"
      },
      {
        "Name": "Region",
        "Value": "eu-west-1"
      }
    ]
  }
]

```

```

    ],
    "MetricName": "ServiceLimitUsage"
  },
  {
    "Namespace": "AWS/TrustedAdvisor",
    "Dimensions": [
      {
        "Name": "ServiceName",
        "Value": "EBS"
      },
      {
        "Name": "ServiceLimit",
        "Value": "Provisioned IOPS"
      },
      {
        "Name": "Region",
        "Value": "ap-south-1"
      }
    ],
    "MetricName": "ServiceLimitUsage"
  },
  ...
]
}

```

Example: List all metrics for a dimension

The following example specifies the `AWS/TrustedAdvisor` namespace and the `Region` dimension to view the metrics available for the specified AWS Region.

```
aws cloudwatch list-metrics --namespace AWS/TrustedAdvisor --dimensions
Name=Region,Value=us-east-1
```

Your output might look like the following.

```

{
  "Metrics": [
    {
      "Namespace": "AWS/TrustedAdvisor",
      "Dimensions": [
        {
          "Name": "ServiceName",
          "Value": "SES"
        }
      ]
    }
  ]
}

```

```

        },
        {
            "Name": "ServiceLimit",
            "Value": "Daily sending quota"
        },
        {
            "Name": "Region",
            "Value": "us-east-1"
        }
    ],
    "MetricName": "ServiceLimitUsage"
},
{
    "Namespace": "AWS/TrustedAdvisor",
    "Dimensions": [
        {
            "Name": "ServiceName",
            "Value": "AutoScaling"
        },
        {
            "Name": "ServiceLimit",
            "Value": "Launch configurations"
        },
        {
            "Name": "Region",
            "Value": "us-east-1"
        }
    ],
    "MetricName": "ServiceLimitUsage"
},
{
    "Namespace": "AWS/TrustedAdvisor",
    "Dimensions": [
        {
            "Name": "ServiceName",
            "Value": "CloudFormation"
        },
        {
            "Name": "ServiceLimit",
            "Value": "Stacks"
        },
        {
            "Name": "Region",
            "Value": "us-east-1"
        }
    ]
}

```

```

        }
      ],
      "MetricName": "ServiceLimitUsage"
    },
    ...
  ]
}

```

Example: List metrics for a specific metric name

The following example specifies the AWS/TrustedAdvisor namespace and the RedResources metric name to view the results for only this specific metric.

```
aws cloudwatch list-metrics --namespace AWS/TrustedAdvisor --metric-name RedResources
```

Your output might look like the following.

```

{
  "Metrics": [
    {
      "Namespace": "AWS/TrustedAdvisor",
      "Dimensions": [
        {
          "Name": "CheckName",
          "Value": "Amazon RDS Security Group Access Risk"
        }
      ],
      "MetricName": "RedResources"
    },
    {
      "Namespace": "AWS/TrustedAdvisor",
      "Dimensions": [
        {
          "Name": "CheckName",
          "Value": "Exposed Access Keys"
        }
      ],
      "MetricName": "RedResources"
    },
    {
      "Namespace": "AWS/TrustedAdvisor",
      "Dimensions": [
        {

```

```
        "Name": "CheckName",
        "Value": "Large Number of Rules in an EC2 Security Group"
      }
    ],
    "MetricName": "RedResources"
  },
  {
    "Namespace": "AWS/TrustedAdvisor",
    "Dimensions": [
      {
        "Name": "CheckName",
        "Value": "Auto Scaling Group Health Check"
      }
    ],
    "MetricName": "RedResources"
  },
  ...
]
```

Trusted Advisor metrics and dimensions

See the following tables for the Trusted Advisor metrics and dimensions that you can use for your CloudWatch alarms and graphs.

Trusted Advisor check-level metrics

You can use the following metrics for Trusted Advisor checks.

Metric	Description
RedResources	The number of resources that are in a red state (action recommended).
YellowResources	The number of resources that are in a yellow state (investigation recommended).

Trusted Advisor service quota-level metrics

You can use the following metrics for AWS service quotas.

Metric	Description
ServiceLimitUsage	The percentage of resource usage against a service quota (formerly referred to as limits).

Dimensions for check-level metrics

You can use the following dimension for Trusted Advisor checks.

Dimension	Description
CheckName	The name of the Trusted Advisor check. You can find all check names in the Trusted Advisor console or the AWS Trusted Advisor check reference.

Dimensions for service quota metrics

You can use the following dimensions for Trusted Advisor service quota metrics.

Dimension	Description
Region	The AWS Region for a service quota.
ServiceName	The name of the AWS service.
ServiceLimit	The name of the service quota. For more information about service quotas, see AWS service quotas in the <i>AWS General Reference</i>.

Logging AWS Trusted Advisor console actions with AWS CloudTrail

Trusted Advisor is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Trusted Advisor. CloudTrail captures actions for Trusted Advisor

as events. The calls captured include calls from the Trusted Advisor console. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon Simple Storage Service (Amazon S3) bucket, including events for Trusted Advisor. 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 the request that was made to Trusted Advisor, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, including how to configure and enable it, see the [AWS CloudTrail User Guide](#).

Trusted Advisor information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When supported event activity occurs in the Trusted Advisor console, 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 Trusted Advisor, 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 AWS Regions. 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](#)

Trusted Advisor supports logging a subset of the Trusted Advisor console actions as events in CloudTrail log files. CloudTrail logs the following actions:

- [BatchUpdateRecommendationResourceExclusion](#)
- CreateEngagement
- CreateEngagementAttachment

- CreateEngagementCommunication
- CreateExcelReport
- DescribeAccount
- DescribeAccountAccess
- DescribeCheckItems
- DescribeCheckRefreshStatuses
- DescribeCheckSummaries
- DescribeChecks
- DescribeNotificationPreferences
- DescribeOrganization
- DescribeOrganizationAccounts
- DescribeReports
- DescribeServiceMetadata
- ExcludeCheckItems
- GenerateReport
- GetEngagement
- GetEngagementAttachment
- GetEngagementType
- GetExcelReport
- [GetOrganizationRecommendation](#)
- [GetRecommendation](#)
- IncludeCheckItems
- ListAccountsForParent
- [ListChecks](#)
- ListEngagementCommunications
- ListEngagementTypes
- ListEngagements
- [ListOrganizationRecommendationAccounts](#)
- [ListOrganizationRecommendationResources](#)
- [ListOrganizationRecommendations](#)

- ListOrganizationalUnitsForParent
- [ListRecommendationResources](#)
- [ListRecommendations](#)
- ListRoots
- RefreshCheck
- SetAccountAccess
- SetOrganizationAccess
- UpdateEngagement
- UpdateEngagementStatus
- UpdateNotificationPreferences
- [UpdateOrganizationRecommendationLifecycle](#)
- [UpdateRecommendationLifecycle](#)

For a complete list of Trusted Advisor console actions, see [Trusted Advisor actions](#).

 Note

CloudTrail also logs the Trusted Advisor API operations in the [AWS Support API Reference](#). For more information, see [Logging AWS Support API calls with AWS CloudTrail](#).

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail userIdentity Element](#).

Example: Trusted Advisor 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.

Example: Log entry for RefreshCheck

The following example shows a CloudTrail log entry that demonstrates the RefreshCheck action for the Amazon S3 Bucket Versioning check (ID R365s2Qddf).

```
{
  "eventVersion": "1.04",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:iam:123456789012:user/janedoe",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "userName": "janedoe",
    "sessionContext": {
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2020-10-21T22:06:18Z"
      }
    }
  },
  "eventTime": "2020-10-21T22:06:33Z",
  "eventSource": "trustedadvisor.amazonaws.com",
  "eventName": "RefreshCheck",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "100.127.34.136",
  "userAgent": "signin.amazonaws.com",
  "requestParameters": {
    "checkId": "R365s2Qddf"
  },
  "responseElements": {
    "status": {
      "checkId": "R365s2Qddf",
      "status": "enqueued",
      "millisUntilNextRefreshable": 3599993
    }
  },
  "requestID": "d23ec729-8995-494c-8054-dedeaEXAMPLE",
  "eventID": "a49d5202-560f-4a4e-b38a-02f1cEXAMPLE",
}
```

```
"eventType": "AwsApiCall",
"recipientAccountId": "123456789012"
}
```

Example: Log entry for UpdateNotificationPreferences

The following example shows a CloudTrail log entry that demonstrates the UpdateNotificationPreferences action.

```
{
  "eventVersion": "1.04",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:iam::123456789012:user/janedoe",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "userName": "janedoe",
    "sessionContext": {
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2020-10-21T22:06:18Z"
      }
    }
  },
  "eventTime": "2020-10-21T22:09:49Z",
  "eventSource": "trustedadvisor.amazonaws.com",
  "eventName": "UpdateNotificationPreferences",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "100.127.34.167",
  "userAgent": "signin.amazonaws.com",
  "requestParameters": {
    "contacts": [
      {
        "id": "billing",
        "type": "email",
        "active": false
      },
      {
        "id": "operational",
        "type": "email",
        "active": false
      }
    ]
  }
}
```

```

    },
    {
      "id": "security",
      "type": "email",
      "active": false
    }
  ],
  "language": "en"
},
"responseElements": null,
"requestID": "695295f3-c81c-486e-9404-fa148EXAMPLE",
"eventID": "5f923d8c-d210-4037-bd32-997c6EXAMPLE",
"eventType": "AwsApiCall",
"recipientAccountId": "123456789012"
}

```

Example: Log entry for GenerateReport

The following example shows a CloudTrail log entry that demonstrates the GenerateReport action. This action creates a report for your AWS organization.

```

{
  "eventVersion": "1.04",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:iam::123456789012:user/janedoe",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "userName": "janedoe",
    "sessionContext": {
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2020-11-03T13:03:10Z"
      }
    }
  },
  "eventTime": "2020-11-03T13:04:29Z",
  "eventSource": "trustedadvisor.amazonaws.com",
  "eventName": "GenerateReport",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "100.127.36.171",

```

```
"userAgent":"signin.amazonaws.com",
"requestParameters":{
  "refresh":false,
  "includeSuppressedResources":false,
  "language":"en",
  "format":"JSON",
  "name":"organizational-view-report",
  "preference":{
    "accounts":[

    ],
    "organizationalUnitIds":[
      "r-j134"
    ],
    "preferenceName":"organizational-view-report",
    "format":"json",
    "language":"en"
  }
},
"responseElements":{
  "status":"ENQUEUED"
},
"requestID":"bb866dc1-60af-47fd-a660-21498EXAMPLE",
"eventID":"2606c89d-c107-47bd-a7c6-ec92fEXAMPLE",
"eventType":"AwsApiCall",
"recipientAccountId":"123456789012"
}
```


Troubleshooting resources

For answers to common troubleshooting questions, see the [AWS Support Knowledge Center](#).

For Amazon EC2 Windows instances, you can use EC2Rescue to examine your instances to help identify common problems, collect log files, and help Support troubleshoot your issues. You can also use EC2Rescue to analyze boot volumes from non-functional instances. For more information, see [How can I use EC2Rescue to troubleshoot and fix common issues on my EC2 Windows instance?](#)

Service-specific troubleshooting

Most AWS service documentation contains troubleshooting topics that can get you started before contacting AWS Support. The following table provides links to troubleshooting topics, arranged by service.

Note

The following table provides a list of the most common services. To search for other troubleshooting topics, use the search text box on the [AWS Documentation landing page](#).

Service	Link
Amazon Web Services	Troubleshooting AWS Signature Version 4 errors
Amazon API Gateway	Troubleshooting issues with HTTP APIs
Amazon AppStream	Troubleshoot Amazon AppStream
Amazon Athena	Troubleshoot in Athena
Amazon Aurora MySQL	Troubleshoot for Amazon Aurora
Amazon Aurora PostgreSQL	Troubleshoot for Amazon Aurora
Amazon EC2 Auto Scaling	Troubleshooting Auto Scaling
AWS Certificate Manager (ACM)	Troubleshooting

Service	Link
AWS CloudFormation	Troubleshooting AWS CloudFormation
Amazon CloudFront	Troubleshooting Troubleshooting RTMP distributions
AWS CloudHSM	Troubleshooting
Amazon CloudSearch	Troubleshooting Amazon CloudSearch
AWS CodeDeploy	Troubleshooting AWS CodeDeploy
Amazon CloudWatch	Troubleshooting
AWS Database Migration Service	Troubleshooting migration tasks in AWS Database Migration Service
AWS Data Pipeline	Troubleshooting
AWS Direct Connect	Troubleshooting AWS Direct Connect
AWS Directory Service	Troubleshooting AWS Directory Service administration issues
Amazon DynamoDB	Troubleshooting Troubleshooting SSL/TLS connection establishment issues
AWS Elastic Beanstalk	Troubleshooting
Amazon Elastic Compute Cloud (Amazon EC2)	Troubleshooting instances Troubleshooting Windows instances Troubleshooting VM Import/Export Troubleshooting API request errors
Amazon Elastic Container Service (Amazon ECS)	Amazon ECS troubleshooting
Amazon Elastic Kubernetes Service (Amazon EKS)	Amazon EKS troubleshooting
Elastic Load Balancing	Troubleshoot your application load balancers Troubleshoot your Classic Load Balancer

Service	Link
Amazon ElastiCache (Memcached)	Troubleshooting applications
Amazon ElastiCache (Redis OSS)	Troubleshooting applications
Amazon EMR	Troubleshoot a cluster
AWS Flow Framework	Troubleshooting and debugging tips
AWS Glue	Troubleshooting AWS Glue
AWS Glue DataBrew	Troubleshooting identity and access in AWS Glue DataBrew
AWS GovCloud (US)	Troubleshooting
AWS Identity and Access Management (IAM)	Troubleshooting IAM
Amazon Keyspaces (for Apache Cassandra)	Troubleshooting Amazon Keyspaces (for Apache Cassandra)
Amazon Kinesis Data Streams	Troubleshooting Amazon Kinesis Data Streams producers Troubleshooting Amazon Kinesis Data Streams consumers
Amazon Managed Service for Apache Flink	Troubleshooting Performance Troubleshooting Amazon Managed Service for Apache Flink for SQL Applications
Amazon Data Firehose	Troubleshooting Amazon Data Firehose
AWS Lambda	Troubleshooting and monitoring AWS Lambda functions with CloudWatch
Amazon OpenSearch Service	Troubleshooting Amazon OpenSearch Service
Amazon Personalize	Troubleshooting
Amazon Quick	Troubleshooting Amazon Quick Troubleshooting skipped row errors

Service	Link
AWS Resource Access Manager (AWS RAM)	Troubleshooting issues with AWS RAM
Amazon Redshift	Troubleshooting queries Troubleshooting data loads Troubleshooting connection issues in Amazon Redshift Troubleshooting Amazon Redshift audit logging Troubleshooting queries in Amazon Redshift Spectrum
Amazon Relational Database Service (Amazon RDS)	Troubleshooting Troubleshooting applications on Amazon RDS Troubleshooting DB issues for Amazon RDS Custom
Amazon Route 53	Troubleshooting Amazon Route 53
Amazon SageMaker AI	Troubleshoot errors Troubleshooting Amazon SageMaker AI Studio
Amazon Silk	Troubleshooting
Amazon Simple Email Service (Amazon SES)	Troubleshooting Amazon SES
Amazon Simple Storage Service (Amazon S3)	Troubleshooting
Amazon Simple Workflow Service (Amazon SWF)	AWS flow framework for Java: Troubleshooting and debugging tips AWS flow framework for Ruby: Troubleshooting and debugging workflows
AWS Storage Gateway	Troubleshooting your gateway
AWS Systems Manager	Troubleshooting SSM Agent
Amazon Virtual Private Cloud (Amazon VPC)	Troubleshooting
AWS Virtual Private Network (Site-to-Site VPN)	Troubleshooting your customer gateway device

Service	Link
AWS WAF	Testing and tuning your AWS WAF protections
Amazon WorkMail	Troubleshooting the Amazon WorkMail web application
Amazon WorkSpaces	Troubleshooting Amazon WorkSpaces issues Troubleshooting Amazon WorkSpaces client issues

Document history

The following table describes the important changes to the documentation since the last release of the AWS Support service.

- **AWS Support API version:** 2013-04-15
- **AWS Support App API version:** 2021-08-20

The following table describes important updates to the AWS Support and AWS Trusted Advisor documentation, beginning May 10, 2021. You can subscribe to the RSS feed to receive notifications about the updates.

Change	Description	Date
Updated Trusted Advisor check: Amazon ECS AWSLogs driver in blocking mode	Updated the description for the c1dvkm4z6b: Amazon ECS AWSLogs driver in blocking mode check to indicate that it evaluates only Amazon ECS task definitions that are actively in use by running tasks or that are among the 5 most recent revisions per task definition family.	September 19, 2026
Multi-session support available for the AWS Support App	The AWS Support App now supports multi-session. You no longer need to turn off multi-session support to use the AWS Support App. For more information, see AWS Support App in Slack .	September 18, 2026

[Added iam:PassRole requirement for the AWS Support App](#)

Added information about the iam:PassRole permission that you need to create or update a Slack channel configuration for the AWS Support App. For more information, see [Managing access to the AWS Support App](#).

September 15, 2026

[New Trusted Advisor check: NitroV4 Amazon EC2 instances running outdated ENA driver](#)

Added a new operational excellence check that identifies Amazon EC2 instances running on Nitro system version 4 that use outdated Elastic Network Adapter (ENA) drivers. For more information, see [Operational Excellence](#) in the Trusted Advisor check reference.

September 8, 2026

[Support for large attachments in the AWS Support API](#)

The AWS Support API now supports uploading and downloading attachments up to 150 MB. For more information, see [Uploading attachments](#).

August 31, 2026

[Updated AWS Security Hub CSPM controls in Trusted Advisor](#)

Trusted Advisor supports only AWS Security Hub CSPM controls available through July 8, 2026. For more information, see [Viewing AWS Security Hub CSPM controls in AWS Trusted Advisor](#).

August 27, 2026

[Updated AWS Support Plans managed policies](#)

Added support agreement actions to the `AWSSupportPlansFullAccess` and `AWSSupportPlansReadOnlyAccess` managed policies, and added service-linked role actions to the `AWSSupportPlansFullAccess` managed policy. For more information, see [AWS managed policies for AWS Support Plans](#).

August 27, 2026

[New: Self-service subscription sign-up for AWS Enterprise Support](#)

You can now subscribe to AWS Enterprise Support directly from the AWS Support Plans console without contacting sales. For more information, see [Sign up for AWS Enterprise Support using self-service subscription](#).

August 27, 2026

[Updated documentation for `AWSSupportServiceRolePolicy`](#)

Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see [AWS managed policy: `AWSSupportServiceRolePolicy`](#).

August 19, 2026

[New: Service-linked role for AWS Support Plans](#)

Added documentation for the `AWSServiceRoleForSupportPlans` service-linked role used by AWS Support Plans. For more information, see [Using service-linked roles for AWS Support Plans](#).

July 27, 2026

[Added guidance on removing AI-enhanced troubleshooting operations](#)

Added a new topic describing how to deny AI-enhanced troubleshooting operations by removing support interaction and Amazon Q Developer conversation actions from your IAM policy. For more information, see [Set up permissions to use AI-enhanced troubleshooting](#).

July 27, 2026

[Multi-session support not available for AWS Support App](#)

Added guidance on disabling multi-session support for the AWS Support App. For more information, see [AWS Support App in Slack](#).

July 22, 2026

[Attach files up to 150 MB to support cases](#)

You can now attach up to 10 files, each up to 150 MB, when creating or replying to a support case in the Support Center Console. For more information, see [Managing attachments](#).

July 21, 2026

[Added email authentication guidance](#)

Added guidance on configuring SPF, DKIM, and DMARC email authentication to ensure email deliverability when corresponding with AWS Support. For more information, see [About the AWS Support API](#).

July 15, 2026

[Updated support plan documentation for invited accounts](#)

Clarified that accounts invited to an organization are automatically subscribed to the organization's support plan, in addition to newly created accounts.

July 2, 2026

[New: AWS Support authorization](#)

Added documentation for AWS Support authorization, which gives you auditable control over AWS Support access to information about your services during support case resolution. For more information, see [AWS Support authorization](#).

June 30, 2026

[Updated CreateContact permission description](#)

Clarified that the CreateContact permission covers contact form submissions such as the Request to remove email sending limitations, Report Abuse, and Simulated Events forms. For more information, see [Adding IAM policies for the Support Center Console API operation](#).

June 26, 2026

[Updated documentation for
AWSSupportServiceRolePolicy](#)

Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see [AWS managed policy: AWSSupportServiceRolePolicy](#).

June 22, 2026

[New: Using AWS DevOps Agent with your AWS Support plan](#)

Added documentation for enabling AWS DevOps Agent from your AWS Support Center Console if your account is on an eligible Support plan. For more information, see [Using AWS DevOps Agent with your AWS Support plan](#).

June 8, 2026

[New: AWS DevOps Agent integration to AI-enhanced troubleshooting](#)

Added documentation for AWS DevOps Agent integration with AI-enhanced troubleshooting in the Support Center Console, including autonomous investigation capabilities for operational issues, required IAM permissions, and managed policy details. For more information, see [AI-enhanced troubleshooting in the Support Center Console](#).

June 8, 2026

Updated procedure to request 5-minute incident response	Corrected the Category step in the procedure to request 5-minute incident response for AWS Unified Operations. For more information, see How to request 5-minute incident response .	May 18, 2026
Updated documentation for AWSTrustedAdvisorServiceRolePolicy	Added new IAM actions <code>ecs:ListClusters</code> , <code>ecs:ListTasks</code> , <code>ecs:DescribeTasks</code> , and <code>ecs:ListTaskDefinitionFamilies</code> to onboard new checks. For more information, see AWS managed policy: AWSTrustedAdvisorServiceRolePolicy .	May 14, 2026
Updated steps for testing Support Center Console API calls	Updated the procedure for testing Support Center Console API calls to include steps for identifying missing IAM permissions using AWS CloudTrail. For more information, see Testing Support Center Console API calls .	May 8, 2026
Updated deadline for Support Center Console API IAM policies	Updated the deadline from November 2, 2026 to November 16, 2026 for adding Support Center Console operations to IAM policies. For more information, see Adding IAM policies for the Support Center Console API operations .	May 8, 2026

[Updated procedure to request 5-minute incident response](#)

Updated the procedure to request 5-minute incident response for AWS Unified Operations. For more information, see [How to request 5-minute incident response](#).

May 4, 2026

[New: Service-linked role for AWS Support Plans](#)

Added documentation for the new `AWSServiceRoleForSupportPlans` service-linked role and the `AWSSupportPlansServiceRolePolicy` managed policy. For more information, see [Using service-linked roles for AWS Support Plans](#).

April 24, 2026

[Updated documentation for `AWSSupportServiceRolePolicy`](#)

Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see [AWS managed policy: `AWSSupportServiceRolePolicy`](#).

April 21, 2026

[Added Turkish language support information](#)

Added information about Turkish language support availability for support cases. For more information, see [Understanding AWS Support response times](#).

April 14, 2026

Updated deadline for Support Center Console API IAM policies	Updated the deadline from June 1, 2026 to November 2, 2026 for adding Support Center Console operations to IAM policies. For more information, see Adding IAM policies for the Support Center Console API operations .	March 27, 2026
Updated AWS Unified Operations team role title	Updated the title from Migration Specialist to Migrations and Events Engineer in the Unified Operations team section. For more information, see Migrations and Events Engineer .	March 19, 2026
Updated AWS Support App documentation for authorizing multiple accounts	Updated the manual steps for authorizing multiple accounts without using CloudFormation templates. For more information, see Authorize multiple accounts and Creating resources with CloudFormation .	February 27, 2026
New Idle NAT gateways check	Added a new check to identify idle or underutilized NAT gateways that might be candidates for cost optimization. For more information, see Change log for AWS Trusted Advisor checks .	February 19, 2026

Updated Manage access to AWS Support Center	Added a reference to the Support Center Console API operations topic to help customers find information about additional operations for their policies. For more information, see Manage access to AWS Support Center .	February 13, 2026
Updated note in Adding IAM policies for the Support Center Console API operations	Updated the note regarding using a custom VPN configuration. For more information, see Adding IAM policies for the Support Center Console API operations .	February 11, 2026
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	January 29, 2026
Added Best practices and target response times for Enterprise Support cases in Create a support case from a support interaction	Added a new drop down section detailing best practices for Enterprise Support customers when creating cases to AWS Support. For more information, see Create a support case from a support interaction .	January 12, 2026
Deprecated 5 Cost Optimization Trusted Advisor checks	For details, see the Change log for AWS Trusted Advisor checks .	January 9, 2026

Updated Change AWS Support Plans	Added a note detailing what happens if an account is added to or removed from an organization that's enrolled in AWS Business Support+. For more information, see Change AWS Support Plans .	January 8, 2026
New section Set up permissions to use AI-enhanced troubleshooting	Added a new section outlining how to configure the required permissions for AI-enhanced troubleshooting in Service Catalog. For more information, see Set up permissions to use AI-enhanced troubleshooting .	December 22, 2025
Updated Trusted Advisor check reference	For details, see the Change log for AWS Trusted Advisor checks .	December 18, 2025
Updated Trusted Advisor check reference	For details, see the Change log for AWS Trusted Advisor checks .	December 17, 2025
Added information regarding virtual meetings with AWS Support	For more information, see Virtual meetings with AWS Support .	December 9, 2025
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	December 8, 2025

Added information regarding new AWS Support plans	For more information, see AWS Support Plans .	December 2, 2025
Added information regarding AI-enhanced troubleshooting	For more information, see AI-enhanced troubleshooting in Support Center Console .	December 2, 2025
Added information a new section for AWS Unified Operations	For more information, see What is AWS Unified Operations .	December 2, 2025
Added description for Support Assistant APIs	Added AWS Support API Actions for Support Assistant : ListInteractions, ListInteractionEntries, and ResolveInteraction in Manage access to AWS Support Center .	December 2, 2025
Updated Trusted Advisor check reference	For details, see the Change log for AWS Trusted Advisor checks .	November 21, 2025
Updated Trusted Advisor check reference	For details, see the Change log for AWS Trusted Advisor checks .	November 17, 2025
Updated Editing and deleting a service-linked role for Support	Customers onboarded to AWS Organizations with an Enterprise Support plan can delete the AWSServiceRoleForSupport service-linked role. For more information, see Editing and deleting a service-linked role for Support .	October 31, 2025

Updated Trusted Advisor check reference	For details, see the Change log for AWS Trusted Advisor checks .	October 15, 2025
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	September 30, 2025
New topic	Added a new topic with information on how to enable promotional plan expiration notifications. For more information, see Configure promotional plan expiration notifications .	September 12, 2025
Updated Trusted Advisor check	For details, see the Change log for AWS Trusted Advisor checks .	September 11, 2025
Updated Testing Support Center Console API calls	Step 3 has been updated to indicate the correct event source, support-console.amazonaws.com . For more information see Testing Support Center Console API calls .	September 2, 2025

Updated Changing AWS Support plans	The steps to upgrade or downgrade your AWS Support subscription have been updated. For more information see Changing AWS Support plans .	August 27, 2025
Updated Trusted Advisor check	For details, see the Change log for AWS Trusted Advisor checks .	August 19, 2025
Removed Trusted Advisor Engage section	Removed Get started with AWS Trusted Advisor Engage (Preview) section.	August 7, 2025
Updated how to add the AWS Support App to a Slack channel	For details, see Configuring a Slack channel .	August 6, 2025
Two checks deprecated	For details, see the Change log for AWS Trusted Advisor checks .	July 22, 2025
Updated managed policies: AWSSupportAccess	For details, see AWS Support updates to AWS managed policies .	July 18, 2025
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	July 15, 2025
Updated check: Amazon EC2 instances with Ubuntu LTS end of standard support	For details, see the Change log for AWS Trusted Advisor checks .	July 3, 2025

Updated Trusted Advisor check: Amazon S3 Bucket Permissions	For details, see the Change log for AWS Trusted Advisor checks .	July 3, 2025
Updated check: MFA on root account	For details, see the Change log for AWS Trusted Advisor checks .	July 2, 2025
Updated check: Amazon ECS AWSLogs driver in blocking mode.	For details, see the Change log for AWS Trusted Advisor checks .	July 2, 2025
Updated Request a service quota increase	Added information on how to create a service quota increase request if your AWS service or AWS Region isn't supported in the Service Quotas console. For more information, see Request service quota increase .	July 2, 2025
Creating a service quota increase request in the Support Center Console is no longer supported.	For more information, see Creating a service quota increase .	June 23, 2025
Added description for UpdateInteraction in Support API	Added AWS Support API UpdateInteraction Actions description in Manage access to AWS Support Center .	June 23, 2025
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	June 17, 2025

New section: About the Support Center Console API	The Support Center Console API enhances your experience with the Support Center Console. For details, see About the Support Center Console API .	June 16, 2025
Updated AWS Trusted Advisor chapter introduction to reflect that Basic and Developer Support plans don't support automatic check refresh. You must manually refresh Security checks to see the most recent check status.	For details, see AWS Trusted Advisor .	June 11, 2025
Updated check: AWS STS global endpoint usage across AWS Regions is now available at all AWS Support plan tiers.	For details, see the Change log for AWS Trusted Advisor checks .	June 9, 2025
New check: Amazon Aurora cost optimization recommendations for DB cluster storage	For details, see the Change log for AWS Trusted Advisor checks .	June 9, 2025
New check: AWS STS global endpoint usage across AWS Regions	For details, see the Change log for AWS Trusted Advisor checks .	June 2, 2025
15 new AWS Cost Optimization Hub checks add to Trusted Advisor	For details, see the Change log for AWS Trusted Advisor checks .	May 30, 2025
Updated three Trusted Advisor checks	For details, see the Change log for AWS Trusted Advisor checks .	May 21, 2025

New feature: Update support case severity	For details, see the <i>Changing a support case severity level</i> section Creating support cases and case management .	May 21, 2025
Updated time you can view AWS Support case details.	For details, see Security for your AWS Support cases .	April 29, 2025
Updated two Trusted Advisor checks	For details, see the Change log for AWS Trusted Advisor checks .	April 2, 2025
Added description for Support API	Added AWS Support API Actions description in Manage access to AWS Support Center .	March 7, 2025
Deprecated 6 AWS Security Hub CSPM checks	For details, see the Change log for AWS Trusted Advisor checks .	March 5, 2025
Removed references to category-level metrics for Trusted Advisor	Category-level metrics for Trusted Advisor are deprecated. References to category-level metrics are removed from Creating Amazon CloudWatch alarms to monitor AWS Trusted Advisor metrics .	January 27, 2025
Updated documentation for Trusted Advisor	Added two new checks: AWS CloudTrail Management Events Logging and Amazon RDS Continuous Backups Not Enabled. For more information, see Change log for AWS Trusted Advisor checks .	December 23, 2024

Updated documentation for Trusted Advisor	Updated Auto Scaling Group Resources. For more information, see Change log for AWS Trusted Advisor checks .	December 23, 2024
Updated documentation for Trusted Advisor	Updated IAM Access Analyzer External Access check. For more information, see Change log for AWS Trusted Advisor checks .	December 23, 2024
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	November 25, 2024
Updated documentation for Trusted Advisor	Added 1 new Trusted Advisor check. For more information, see Change log for AWS Trusted Advisor checks .	November 22, 2024
Added documentation for AWS managed policies for AWS Partner-Led Support	Added documentation for a new AWS managed policy <code>AWSPartnerLedSupportReadOnlyAccess</code> . For more information, see AWS managed policies for AWS Partner-Led Support .	November 22, 2024
Updated documentation for Trusted Advisor	Updated 3 Trusted Advisor checks. For more information, see Change log for AWS Trusted Advisor checks .	November 7, 2024

Updated documentation for AWS Support Plans	Added a new log example for the <code>ListSupportPlanModifiers</code> operation to the Logging Support Plans API calls with AWS CloudTrail page.	November 6, 2024
Updated documentation for AWSTrustedAdvisorServiceRolePolicy	Added new IAM actions <code>elasticloadbalancing:DescribeListeners</code> and <code>elasticloadbalancing:DescribeRules</code> , to onboard a new security check. For more information, see AWS managed policy: AWSTrustedAdvisorServiceRolePolicy .	October 30, 2024
Updated documentation for Trusted Advisor	Added 4 new Trusted Advisor checks. For more information, see Change log for AWS Trusted Advisor checks .	October 11, 2024
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	October 8, 2024

Updated documentation for Trusted Advisor	Moved 1 Cost Optimization check under the Fault Tolerance pillar. Updated 1 Security check and 1 Fault Tolerance check. For more information, see Change log for AWS Trusted Advisor checks .	October 2, 2024
Updated AWS Trusted Advisor Engage section	Updated the AWS Trusted Advisor Engage section to reference AWS Countdown . For more information, see Get started with AWS Trusted Advisor Engage (Preview) .	September 16, 2024
Updated documentation for AWS Support Plans	Added a new permission and CloudTrail documentation for viewing a list of support plan modifiers. For more information, see Manage access to AWS Support Plans , AWS managed policies for AWS Support Plans and Logging AWS Support Plans API calls with AWS CloudTrail .	September 9, 2024
Updated documentation for Trusted Advisor	Trusted Advisor added 9 new checks on Aug 23rd. For more information, see Change log for AWS Trusted Advisor checks .	August 23, 2024

Updated documentation for Trusted Advisor	Updated 1 Trusted Advisor Operational Excellence check and added 1 new Trusted Advisor Security check. For more information, see Change log for AWS Trusted Advisor checks .	August 22, 2024
Updated documentation for Trusted Advisor	Updated 6 Trusted Advisor Security checks. For more information, see Change log for AWS Trusted Advisor checks .	August 20, 2024
Updated documentation for Trusted Advisor	Updated 2 Trusted Advisor checks. For more information, see Change log for AWS Trusted Advisor checks .	August 12, 2024
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	August 5, 2024
Updated documentation for Trusted Advisor	Updated 9 Trusted Advisor Checks. For more information, see Change log for AWS Trusted Advisor checks .	July 21, 2024

[Updated documentation for
AWSTrustedAdvisorServiceRolePolicy](#)

Added new IAM actions `access-analyzer:ListAnalyzers` , `cloudwatch:ListMetrics` , `dax:DescribeClusters` , `ec2:DescribeNatGateways` , `ec2:DescribeRouteTables` , `ec2:DescribeVpcEndpoints` , `ec2:GetManagedPrefixListEntries` , `elasticloadbalancing:DescribeTargetHealth` , `iam:ListSAMLProviders` , `kafka:DescribeClusterV2` , `network-firewall:ListFirewalls` , `network-firewall:DescribeFirewall` and `sqs:GetQueueAttributes` to onboard new checks. For more information, see [AWS managed policy: AWSTrustedAdvisorServiceRolePolicy](#).

June 11, 2024

[Removed 5 AWS Trusted Advisor checks from documentation](#)

Removed 5 AWS Trusted Advisor checks that are now deprecated. For more information, see [Change log for AWS Trusted Advisor checks](#).

May 15, 2024

Added 1 new AWS Trusted Advisor Security check to documentation	Added 1 new AWS Trusted Advisor Security check to documentation. For more information, see Change log for AWS Trusted Advisor checks .	May 15, 2024
Removed 3 Fault Tolerance checks from documentation	Removed 3 Fault Tolerance checks that are now deprecated. For more information, see Change log for AWS Trusted Advisor checks .	April 25, 2024
Updated Fault Tolerance and Security check documentation	Added 1 new fault tolerance check. Updated 1 fault tolerance and 1 security check. For more information, see Change log for AWS Trusted Advisor checks .	March 29, 2024
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	March 22, 2024
Updated documentation for Support plan	Updates to the Features of Support Plans. For more information, see Support plans .	March 11, 2024

Updated documentation for Trusted Advisor	Added 1 fault tolerance check. For more information, see Change log for AWS Trusted Advisor checks .	February 29, 2024
Updated documentation for Trusted Advisor	Added 1 fault tolerance check. For more information, see Change log for AWS Trusted Advisor checks .	January 31, 2024
Updated documentation for AWSTrustedAdvisorServiceRolePolicy	Added new IAM actions <code>cloudtrail:GetTrail</code> , <code>cloudtrail:ListTrails</code> , <code>cloudtrail:GetEventSelectors</code> , <code>outposts:GetOutpost</code> , <code>outposts:ListAssets</code> and <code>outposts:ListOutposts</code> to onboard new checks. For more information, see AWS managed policy: AWSTrustedAdvisorServiceRolePolicy .	January 18, 2024
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	January 17, 2024

Updated documentation for Trusted Advisor	Updated 1 fault tolerance check to amend title and description. For more information, see Change log for AWS Trusted Advisor checks .	January 8, 2024
Updated documentation for Trusted Advisor	Updated 1 security check to reflect change in deprecation period. For more information, see Change log for AWS Trusted Advisor checks .	December 21, 2023
Updated documentation for Trusted Advisor	Added 2 security checks and 2 performance checks. For more information, see Change log for AWS Trusted Advisor checks .	December 20, 2023
Updated documentation for Trusted Advisor	Added 1 security check. For more information, see Change log for AWS Trusted Advisor checks .	December 15, 2023
Updated documentation for Trusted Advisor Engage	Updated Trusted Advisor Engage documentation with changes for email notification option.	December 14, 2023
Updated documentation for Trusted Advisor Engage	Updated Trusted Advisor Engage documentation with changes for scheduled engagements.	December 11, 2023

Updated documentation for Trusted Advisor	Added 2 new fault tolerance checks and 1 cost optimization check. For more information, see Change log for AWS Trusted Advisor checks .	December 7, 2023
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	December 6, 2023
Updated AWS managed policies for Trusted Advisor	Updated the AWSTruste dAdvisorPriorityFu llAccess and AWSTruste dAdvisorPriorityRe adOnlyAccess AWS managed policies to include statement IDs. For more information, see AWS managed policies for AWS Trusted Advisor .	December 6, 2023
Updated documentation for Trusted Advisor	Added 3 new fault tolerance checks. For more information, see Change log for AWS Trusted Advisor checks .	November 17, 2023
Updated documentation for Trusted Advisor	Added 37 new checks for Amazon RDS. For more information, see Change log for AWS Trusted Advisor checks .	November 15, 2023

Updated documentation for AWSTrustedAdvisorServiceRolePolicy	Added new IAM actions <code>ec2:DescribeRegions</code> , <code>s3:GetLifecycleConfiguration</code> , <code>ecs:DescribeTaskDefinition</code> and <code>ecs:ListTaskDefinitions</code> to onboard new checks. For more information, see AWS managed policy: AWSTrustedAdvisorServiceRolePolicy .	November 9, 2023
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	October 27, 2023
Updated documentation for Trusted Advisor	Added 64 new checks integrated from AWS Config. For more information, see Change log for AWS Trusted Advisor checks .	October 26, 2023
Updated documentation for Trusted Advisor	Added six new fault tolerance checks in Trusted Advisor. For more information, see the Change log for AWS Trusted Advisor checks .	October 12, 2023

Updated documentation for AWSTrustedAdvisorServiceRolePolicy	Added new IAM actions <code>route53resolver:ListResolverEndpoints</code> , <code>route53resolver:ListResolverEndpointIpAddresses</code> , <code>ec2:DescribeSubnets</code> , <code>kafka:ListClustersV2</code> and <code>kafka:ListNodes</code> to onboard new resilience checks. For more information, see AWS managed policy: AWSTrustedAdvisorServiceRolePolicy .	September 14, 2023
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	August 28, 2023
Updated documentation for Trusted Advisor	Added 1 new service limits checks for AWS Lambda. For more information, see the Change log for AWS Trusted Advisor checks .	August 17, 2023
Updated documentation for Trusted Advisor	Added 1 new fault tolerance checks for Lambda. For more information, see the Change log for AWS Trusted Advisor checks .	August 3, 2023

[Updated documentation for
Trusted Advisor Engage](#)

Updated Trusted Advisor Engage documentation with changes to forms for creating and editing engagements. Added page with Example Service Control Policies for AWS Trusted Advisor.

July 27, 2023

[Updated documentation for
AWSSupportServiceRolePolicy](#)

Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see [AWS managed policy: AWSSupportServiceRolePolicy](#).

June 26, 2023

[Updated documentation for
Trusted Advisor](#)

Added two new fault tolerance checks for Amazon MQ. Added one new fault tolerance check and one new performance check for Amazon Elastic File System. For more information, see the [Change log for AWS Trusted Advisor checks](#).

June 1, 2023

[Updated documentation for
Trusted Advisor](#)

Added two new fault tolerance checks for NAT Gateway. For more information, see the [Change log for AWS Trusted Advisor checks](#).

May 16, 2023

[Updated documentation for AWS Support Plans](#)

Added a new permission and CloudTrail documentation for the creation of support plan schedules. For more information, see [Manage access to AWS Support Plans](#), [AWS managed policies for AWS Support Plans](#) and [Logging AWS Support Plans API calls with AWS CloudTrail](#).

May 8, 2023

[Updated documentation for AWSSupportServiceRolePolicy](#)

Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see [AWS managed policy: AWSSupportServiceRolePolicy](#).

May 2, 2023

[Updated documentation for Trusted Advisor Engage and Trusted Advisor Priority](#)

Clarified prerequisites for Trusted Advisor Engage and Trusted Advisor Priority. Added example IAM policy with ability to use Trusted Advisor Engage and to enable trusted access to Trusted Advisor.

April 28, 2023

[Updated documentation for Trusted Advisor](#)

Added two new fault tolerance checks for AWS Resilience Hub and Incident Manager. For more information, see the [Change log for AWS Trusted Advisor checks](#).

April 27, 2023

[Added documentation for
Trusted Advisor Engage](#)

You can use AWS Trusted Advisor Engage to get the most out of your AWS Support Plans by making it easy for you to see, request and track all your proactive engagements, and communicate with your AWS account team about ongoing engagements. For more information, see [Get started with AWS Trusted Advisor Engage](#).

April 6, 2023

[Updated documentation for
Trusted Advisor](#)

Added two new fault tolerance checks for Amazon ECS. For more information, see the [Change log for AWS Trusted Advisor checks](#).

March 30, 2023

[Updated documentation for
AWSSupportServiceRolePolicy](#)

Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see [AWS managed policy: AWSSupportServiceRolePolicy](#).

March 16, 2023

[Added documentation for Trusted Advisor Priority](#)

Updated the Trusted Advisor Priority console:

February 16, 2023

- The **Acknowledge** and **Dismiss** buttons have replaced the **Accept** and **Reject** buttons.
- You don't need to enter your job title or name to acknowledge, resolve, dismiss, or reopen recommendations.

For more information, see [Getting started with Trusted Advisor Priority](#).

[Updated code examples for Support](#)

Added .NET, Java, and Kotlin code examples that show how to use Support with an AWS software development kit (SDK). For more information, see [Code examples for Support using AWS SDKs](#).

January 16, 2023

[Updated documentation for AWSSupportServiceRolePolicy](#)

Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see [AWS managed policy: AWSSupportServiceRolePolicy](#).

January 10, 2023

[Updated documentation for AWS Support App](#)

You can search for support cases in Slack by using filter options or searching by case ID. For more information, see [Searching for support cases in Slack](#).

December 29, 2022

[Updated documentation for AWS Support App](#)

You can also use Terraform to create your resources for the AWS Support App. For more information, see [Create AWS Support App resources by using Terraform](#).

December 22, 2022

[Updated documentation for Trusted Advisor](#)

Added three new fault tolerance checks for Amazon MemoryDB, Amazon ElastiCache, and AWS CloudHSM. For more information, see the [Change log for AWS Trusted Advisor checks](#).

December 15, 2022

[Updated documentation for the AWS Support App in Slack](#)

You can now request live chat support for the following options:

December 14, 2022

- Account and billing support cases.
- Japanese language support for technical support cases.
- For more information, see [Creating support cases in a Slack channel](#).

[Updated documentation for AWS Support](#)

Added documentation about new endpoints for the Support API. For more information, see [About the AWS Support API](#).

December 14, 2022

[Added documentation for CloudFormation templates to use for the AWS Support App in Slack](#)

You can use CloudFormation templates to create Slack configuration workspaces and channels for AWS accounts in AWS Organizations. For more information, see [Creating AWS Support App resources with AWS CloudFormation](#).

December 5, 2022

[Updated documentation for Trusted Advisor](#)

Added two new fault tolerance checks for AWS Resilience Hub. For more information, see the [Change log for AWS Trusted Advisor checks](#).

November 17, 2022

[Added documentation for your AWS Security Hub CSPM findings in Trusted Advisor](#)

Your findings from Security Hub CSPM controls are removed from Trusted Advisor faster. For more information, see the [Change log for AWS Trusted Advisor checks](#).

November 17, 2022

[Updated documentation for AWS Trusted Advisor](#)

Added documentation for Trusted Advisor Recommendations. For more information, see the [Change log for AWS Trusted Advisor checks](#).

November 16, 2022

Updated documentation for the AWS Support App in Slack	Added documentation for Japanese language support. For more information, see Creating support cases in a Slack channel .	November 11, 2022
Updated documentation for AWS Support Plans	Added troubleshooting information to allow Support Plans access in an organization. For more information, see Troubleshooting .	November 9, 2022
Updated documentation for the AWS Support App in Slack	Added documentation for supportapp permissions. For more information, see Permissions required for the AWS Support App to connect to Slack .	November 1, 2022
Updated documentation for the AWS Support App in Slack	You can use the RegisterSlackWorkspaceForOrganization API operation to register a Slack workspace for your AWS account. To call this API, your account must be part of an organization in AWS Organizations. For more information, see the AWS Support App in Slack API Reference .	October 19, 2022

[Updated documentation for
AWSSupportServiceRolePolicy](#)

Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see [AWS managed policy: AWSSupportServiceRolePolicy](#).

October 4, 2022

[Updated documentation for
Support Plans](#)

You can now use AWS Identity and Access Management (IAM) to manage permissions to change the support plan for your AWS account. For more information, see the following topics:

September 29, 2022

- [Managing access for AWS Support Plans](#)
- [AWS managed policies for AWS Support Plans](#)
- [Changing AWS Support Plans](#)
- [Logging AWS Support Plans API calls with AWS CloudTrail](#)

[Updated documentation for
the AWS Support App in Slack](#)

Added documentation on how to configure a public or private channel to use with the AWS Support App. For more information, see [Configuring a Slack channel](#).

September 22, 2022

[Updated documentation for AWS Support](#)

Added a new section about security for your support cases. For more information, see [Security for your AWS Support cases](#).

September 9, 2022

[Updated documentation for Trusted Advisor](#)

Added a new security check for Amazon EC2. For more information, see the [Change log for AWS Trusted Advisor checks](#).

September 1, 2022

[Updated documentation for the AWS Support App in Slack](#)

See the following topics:

August 24, 2022

You can use the AWS Support App to manage your support cases, request service quota increases, and chat with support agents directly in your Slack channels. For more information, see the [AWS Support App in Slack documentation](#).

You can attach AWS managed policies to your IAM roles to use the AWS Support App. For more information, see [AWS managed policies for AWS Support App in Slack](#).

New API reference for the AWS Support App. See the [AWS Support App API Reference](#).

[Updated documentation for
AWSSupportServiceRolePolicy](#)

Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see [AWS managed policy: AWSSupportServiceRolePolicy](#).

August 17, 2022

[Added documentation for
Trusted Advisor Priority](#)

Trusted Advisor Priority adds support for the following features:

August 17, 2022

- Delegated administrators
- Daily and weekly email notifications for recommendation summaries
- Reopen resolved or rejected recommendations
- AWS managed policies

For more information, see [Getting started with Trusted Advisor Priority](#).

[Updated documentation for
Trusted Advisor](#)

The **Preferences** page in the Trusted Advisor console has been updated. For more information, see [Getting started with AWS Trusted Advisor](#).

July 15, 2022

[Updated documentation for Trusted Advisor](#)

Updated the checks to include the following information: July 7, 2022

- **Alert Criteria**
- **Recommended Action**
- **Additional Resources**
- **Report columns**

For more information, see the [AWS Trusted Advisor check reference](#).

[Updated documentation for AWS Support](#)

Added documentation that explains how to manage your support cases. June 28, 2022

- [Updating an existing support case](#)
- [Troubleshooting](#)

[Updated documentation for AWSSupportServiceRolePolicy](#)

Updated permissions to provide billing, administrative, and support services for the service-linked role. For more information, see [AWS managed policy: AWSSupportServiceRolePolicy](#). June 23, 2022

[Updated documentation for
Trusted Advisor](#)

Trusted Advisor supports additional AWS Foundational Security Best Practices security standard controls that are sourced from AWS Security Hub CSPM. For more information, see the [Change log for AWS Trusted Advisor checks](#).

June 23, 2022

[Updated documentation for
Trusted Advisor](#)

Added information about how to request service quota increases. For more information, see [Service limits](#).

June 21, 2022

[Updated documentation for
AWS Support](#)

The create case experience has been updated in the Support Center Console. For more information, see [Creating support cases and case management](#).

May 18, 2022

[Updated documentation for
Trusted Advisor](#)

Added four checks for Amazon EBS and AWS Lambda. For more information, see [Opt in AWS Compute Optimizer to add Trusted Advisor checks](#).

May 4, 2022

[Updated documentation for
AWSSupportServiceRolePolicy](#)

Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see [AWS managed policy: AWSSupportServiceRolePolicy](#).

April 27, 2022

Updated documentation for the Exposed Access Keys check	This check is now automatically refreshed for you. For more information, see Change log for AWS Trusted Advisor checks .	April 25, 2022
Updated documentation for Trusted Advisor	The AWS Direct Connect checks in the fault tolerance category are updated. For more information, see Change log for AWS Trusted Advisor checks .	March 29, 2022
Updated documentation for AWSSupportServiceRolePolicy	Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see AWS managed policy: AWSSupportServiceRolePolicy .	March 14, 2022
Added documentation for Trusted Advisor Priority	You can use Trusted Advisor Priority to view a list of prioritized recommendations from your technical account manager (TAM). For more information, see Getting started with Trusted Advisor Priority .	February 28, 2022
Updated documentation for using Amazon EventBridge for Trusted Advisor	You can create an EventBridge rule to monitor changes to your Trusted Advisor checks. For more information, see Monitoring AWS Trusted Advisor check results with EventBridge .	February 21, 2022

[New documentation for using Amazon EventBridge to monitor AWS Support cases](#)

You can create an EventBridge rule to monitor and receive notifications about your support cases. For more information, see [Monitoring Support cases with EventBridge](#).

February 21, 2022

[Updated documentation for AWSSupportServiceRolePolicy](#)

Added new permissions to provide billing, administrative, and support services for the service-linked role. For more information, see [AWS managed policy: AWSSupportServiceRolePolicy](#).

February 17, 2022

[Added documentation for integrating with AWS Security Hub CSPM](#)

In the Trusted Advisor console, you can now view the findings for your Security Hub CSPM controls that are part of the AWS Foundational Security Best Practices security standard. For more information, see [Viewing AWS Security Hub CSPM controls in the AWS Trusted Advisor console](#).

January 18, 2022

[Updated documentation for Trusted Advisor](#)

Added three new checks for Amazon EC2 instances that are running Microsoft SQL Server.

December 20, 2021

- Amazon EC2 instances consolidation for Microsoft SQL Server
- Amazon EC2 instances over-provisioned for Microsoft SQL Server
- Amazon EC2 instances with Microsoft SQL Server end of support

For more information, see the [AWS Trusted Advisor check reference](#).

[Updated documentation for Trusted Advisor](#)

Trusted Advisor added four new checks for AWS Well-Architected

December 20, 2021

- AWS Well-Architected high risk issues for cost optimization
- AWS Well-Architected high risk issues for performance
- AWS Well-Architected high risk issues for security
- AWS Well-Architected high risk issues for reliability

For more information, see the [AWS Trusted Advisor check reference](#).

[Updated documentation](#)

If you have an [Enterprise On-Ramp](#) Support plan, you have access to all Trusted Advisor checks and the AWS Support API.

November 24, 2021

[Updated documentation for Trusted Advisor](#)

Trusted Advisor added two new checks for Amazon Comprehend. For more information, see the [AWS Trusted Advisor check reference](#).

September 29, 2021

Updated documentation for Trusted Advisor	The check name for Amazon OpenSearch Service Reserved Instance Optimization was updated. For more information, see Change log for AWS Trusted Advisor checks .	September 8, 2021
Updated documentation for Trusted Advisor checks	Added a reference topic for all Trusted Advisor checks. For more information, see AWS Trusted Advisor check reference .	September 1, 2021
Updated documentation for Trusted Advisor managed policies	Updated documentation for the Trusted Advisor managed policies. For more information, see AWS managed policies for AWS Support and AWS Trusted Advisor .	August 10, 2021
Updated documentation for Trusted Advisor	Updated documentation for the Trusted Advisor console. For more information, see Get started with AWS Trusted Advisor .	July 16, 2021
Updated documentation for creating Support cases	Added documentation about how to create a related support case for cases that are permanently closed. For more information, see Reopening a closed case and Creating a related case .	June 8, 2021

[Updated documentation for Trusted Advisor](#)

Trusted Advisor added two new checks for Amazon Elastic Block Store (Amazon EBS) volume storage. For more information, see [Change log for AWS Trusted Advisor checks](#).

June 8, 2021

[Updated documentation](#)

The following topics are updated:

May 12, 2021

- Updated procedures and added content to the [Creating Amazon CloudWatch alarms to monitor AWS Trusted Advisor metrics](#) topic
- Added the [Service quotas for the AWS Support API](#) section

Earlier updates

Change	Description	Date
Updated documentation for Trusted Advisor	Added documentation to filter, refresh, and download check results. For more information, see the following sections: • Filter your checks • Refresh check results • Download check results	March 16, 2021
Updated documentation about AWS managed policies	Added information about the <code>AWSSupportServiceRolePolicy</code> AWS managed	March 16, 2021

Change	Description	Date
	policy. For more information, see Using service-linked roles for AWS Support .	
Added checks for AWS Lambda	Added four AWS Trusted Advisor checks for Lambda in the Change log for AWS Trusted Advisor .	March 8, 2021
Updated service limit checks for Amazon Elastic Block Store	Updated five AWS Trusted Advisor checks for Amazon EBS in the Change log for AWS Trusted Advisor .	March 5, 2021
Updated documentation for CloudTrail logging	CloudTrail supports logging for console actions when you change your AWS Support plan. For more information, see Logging changes to your Support plan .	February 9, 2021
Updated documentation for Trusted Advisor	Updated the Get started with Trusted Advisor Recommendations topic.	January 29, 2021
Updated documentation for Trusted Advisor reports	Added a Troubleshooting section for using Trusted Advisor reports with other AWS services.	December 4, 2020
Added AWS Trusted Advisor support for AWS CloudTrail logging	CloudTrail supports logging for a subset of Trusted Advisor console actions. For more information, see Logging AWS Trusted Advisor console actions with AWS CloudTrail .	November 23, 2020
Added a change log topic	View changes to AWS Trusted Advisor checks and categories in the Change log for AWS Trusted Advisor .	November 18, 2020
Added support for organizational units	You can now create reports for Trusted Advisor checks for organizational units (OUs). For more information, see Create organizational view reports .	November 17, 2020

Change	Description	Date
Updated the logging with AWS CloudTrail topic	Added an example log entry for a Trusted Advisor API operation. See AWS Trusted Advisor information in CloudTrail logging .	October 22, 2020
Added AWS Support quotas	Added information about the current quotas and restrictions for Support. See the Support endpoints and quotas in the <i>AWS General Reference</i> .	August 4, 2020
Organizational view for AWS Trusted Advisor	You can now create reports for Trusted Advisor checks for accounts that are part of AWS Organizations. See Organizational view for AWS Trusted Advisor .	July 17, 2020
Security and AWS Support	Updated information about security considerations when using AWS Support and Trusted Advisor. See Security in AWS Support	May 5, 2020
Security and AWS Support	Added information about security considerations when using AWS Support.	January 10, 2020
Using Trusted Advisor as a web service	Added updated instructions to refresh Trusted Advisor data after getting list of Trusted Advisor checks.	November 1, 2018
Using Service-linked roles	Added new section.	July 11, 2018
Getting Started: Troubleshooting	Added troubleshooting links for Route 53 and AWS Certificate Manager.	September 1, 2017
Case Management Example: Creating a Case	Added a note about the CC box for users who have the Basic support plan.	August 1, 2017

Change	Description	Date
Monitoring Trusted Advisor Check Results with CloudWatch Events	Added new section.	November 18, 2016
Case Management	Updated the names of case severity levels.	October 27, 2016
Logging AWS Support Calls with AWS CloudTrail	Added new section.	April 21, 2016
Getting Started: Troubleshooting	Added more troubleshooting links.	May 19, 2015
Getting Started: Troubleshooting	Added more troubleshooting links.	November 18, 2014
Getting Started: Case Management	Updated to reflect Service Catalog in the AWS Management Console.	October 30, 2014
Programming the Life of an AWS Support Case	Added information about new API elements for adding attachments to cases and for omitting case communications when retrieving case history.	July 16, 2014
Accessing AWS Support	Removed named support contacts as an access method.	May 28, 2014
Getting Started	Added the Getting Started section.	December 13, 2013
Initial publication	New AWS Support service released.	April 30, 2013