



User Guide

AWS AppConfig



AWS AppConfig: User Guide

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

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

Table of Contents

What is AWS AppConfig?	1
Get started with AWS AppConfig	1
AWS AppConfig use cases	1
Benefits overview	2
How AWS AppConfig works	3
Pricing for AWS AppConfig	5
AWS AppConfig quotas	5
Additional resources	5
Blogs	6
SDKs	6
Setting up AWS AppConfig	7
Sign up for an AWS account	7
Understanding IPv6 support	7
Configure permissions for automatic rollback	8
Step 1: Create the permission policy for rollback based on CloudWatch alarms	8
Step 2: Create the IAM role for rollback based on CloudWatch alarms	9
Step 3: Add a trust relationship	10
Creating	12
Understanding the configuration profile IAM role	13
Creating a namespace	16
Creating an AWS AppConfig application (console)	16
Creating an AWS AppConfig application (command line)	17
Creating environments	18
Creating an AWS AppConfig environment (console)	19
Creating an AWS AppConfig environment (command line)	20
Creating a configuration profile in AWS AppConfig	22
Creating a feature flag configuration profile	26
Creating a free form configuration profile	58
Creating a configuration profile for non-native data sources	73
Deploying	75
Working with deployment strategies	76
Using predefined deployment strategies	79
Create a deployment strategy	81
Deploying a configuration	85

Deploy a configuration (console)	86
Deploy a configuration (commandline)	87
Deploying with CodePipeline	91
How integration works	92
Reverting a configuration	92
Retrieving	94
What is AWS AppConfig Agent?	95
How to use AWS AppConfig Agent to retrieve configuration data	97
Using AWS AppConfig Agent for user- or entity-based gradual deployments	97
Using AWS AppConfig Agent with AWS Lambda	99
Using AWS AppConfig Agent with Amazon EC2 and on-premises machines	256
Using AWS AppConfig Agent with Amazon ECS and Amazon EKS	275
Retrieving feature flags	296
Using a manifest to enable additional retrieval features	300
Generating a client using the OpenAPI specification	310
Working with AWS AppConfig Agent local development mode	313
Browser and mobile use considerations	318
Retrieving configuration data without AWS AppConfig Agent	321
Extending AWS AppConfig workflows	325
Understanding AWS AppConfig extensions	325
Step 1: Determine what you want to do with extensions	326
Step 2: Determine when you want the extension to run	327
Step 3: Create an extension association	328
Step 4: Deploy a configuration and verify the extension actions are performed	328
Working with AWS authored extensions	329
Using the AWS AppConfig deployment events to Amazon EventBridge extension	329
Using the AWS AppConfig deployment events to Amazon SNS extension	332
Using the AWS AppConfig deployment events to Amazon SQS extension	335
Using the Jira extension	339
Walkthrough: Creating custom AWS AppConfig extensions	344
Step 1: Create a Lambda function for a custom AWS AppConfig extension	346
Step 2: Configure permissions for a custom AWS AppConfig extension	352
Step 3: Create a custom AWS AppConfig extension	353
Step 4: Create an extension association for a custom AWS AppConfig extension	357
AWS AppConfig experimentation	360
Experimentation terminology	360

Example experiment workflow	362
About experiments in AWS AppConfig	363
Documenting your hypothesis	364
Specifying the target audience for the experiment	364
Selecting the experiment feature flag	371
About controls and treatments	372
About running and monitoring an experiment	375
About data collection	378
Creating and running an experiment	381
Step 1: Configuring prerequisites	381
Step 2: Creating an experiment definition	382
Step 3: Starting an experiment	385
Step 4: Increasing exposure	385
Stopping an experiment	387
Promoting a winning treatment	388
Cleaning up an experiment	388
Code samples	390
Creating or updating a freeform configuration stored in the hosted configuration store	390
Creating a configuration profile for a secret stored in Secrets Manager	393
Deploying a configuration profile	394
Using AWS AppConfig Agent to read a freeform configuration profile	399
Using AWS AppConfig Agent to read a specific feature flag	401
Using AWS AppConfig Agent to retrieve a feature flag with variants	402
Using the GetLatestConfiguration API action to read a freeform configuration profile	405
Cleaning up your environment	415
Deletion protection	421
Bypassing or forcing a deletion protection check	422
Security	424
Implement least privilege access	424
Data encryption at rest for AWS AppConfig	425
AWS PrivateLink	430
Considerations	430
Create an interface endpoint	430
Create an endpoint policy	431
Secrets Manager key rotation	432
Setting up automatic rotation of Secrets Manager secrets deployed by AWS AppConfig ...	432

Monitoring	434
CloudTrail logs	435
AWS AppConfig data events in CloudTrail	436
AWS AppConfig management events in CloudTrail	438
AWS AppConfig event examples	438
Logging metrics for AWS AppConfig data plane calls	439
Creating an alarm for a CloudWatch metric	442
Monitoring deployments for automatic rollback	442
Recommended metrics to monitor for automatic rollback	443
Document history	450

What is AWS AppConfig?

AWS AppConfig helps you safely change application behavior in production without redeploying code. Using feature flags and dynamic free-form configurations, you can control how your application runs in real time. This approach reduces risk, accelerates releases, and enables faster responses to issues. You can gradually roll out new features to specific users, monitor their impact, and expand availability with confidence. You can also update block lists, allow lists, throttling limits, and logging levels instantly, allowing you to mitigate issues and fine-tune performance without a deployment.

AWS AppConfig experimentation enables software teams to run A/B tests and gradual feature rollouts using production traffic. Built on Amazon's 25+ years of software experimentation expertise (including on Amazon.com), the service integrates experimentation directly into your existing AWS AppConfig feature flag workflows. Teams can define hypotheses, target specific audiences, split traffic between treatments, and gradually increase exposure while monitoring operational metrics using CloudWatch or other preferred monitoring tools. AWS AppConfig experimentation works with any compute environment supported by the AWS AppConfig Agent, including Amazon EC2, AWS Lambda, Amazon ECS, and Amazon EKS. Teams can export experiment assignment data to data warehouses such as Amazon Redshift, Snowflake, or Databricks for results analysis. As AI continues to accelerate how teams build and iterate on software, AWS AppConfig experimentation provides a way to validate AI-driven changes with measurable production data and reduced operational risk.

Get started with AWS AppConfig

The following video can help you understand the capabilities of AWS AppConfig.

View more AWS videos on the [Amazon Web Services YouTube Channel](#).

AWS AppConfig use cases

AWS AppConfig supports a broad spectrum of use cases:

- **Application experimentation** – Safely evaluate application changes in production using controlled audience exposure and feature flag-based treatments.
- **Feature flags and toggles** – Gradually release new capabilities to targeted users, monitor impact, and instantly roll back changes if issues occur.

- **Application tuning** – Introduce changes safely in production, measure their effects, and refine behavior without redeploying code.
- **Allow list or block list** – Control access to features or restrict specific users in real time, without modifying application code.
- **Centralized configuration storage** – Manage configuration data consistently across workloads. AWS AppConfig can deploy configuration from the AWS AppConfig hosted configuration store, AWS Secrets Manager, Systems Manager Parameter Store, or Amazon S3.

Benefits overview

The following brief overview outlines the benefits of using AWS AppConfig.

Improve efficiency and release changes faster

Feature flags help you release changes to production more quickly and with less risk. Instead of managing long-lived development branches and complex merges, you can adopt trunk-based development and control feature availability through flags. This approach allows you to deploy code that remains hidden from users until you are ready to release it. When the feature is ready, enable it without redeploying code. If issues occur after release, you can disable the feature immediately without rolling back the deployment.

Avoid unintended changes or failures with built-in safety features

AWS AppConfig includes safety features that help prevent configuration changes from causing application failures:

- **Validators:** Ensure that configuration data is syntactically and semantically correct before deployment.
- **Deployment strategies:** Gradually roll out changes over a defined period of time to reduce risk.
- **Monitoring and automatic rollback:** AWS AppConfig integrates with Amazon CloudWatch to monitor application health. If a configuration change triggers an alarm, AWS AppConfig automatically rolls back the change to minimize impact.


Secure and scalable feature flag deployments

AWS AppConfig integrates with AWS Identity and Access Management (IAM) to provide fine-grain, role-based access control. It also integrates with AWS Key Management Service (AWS KMS) for encryption and AWS CloudTrail for auditing.

AWS developed and validated AWS AppConfig safety controls with internal teams that operate at scale before making them available to external customers. This foundation helps ensure the service is secure, reliable, and ready for production workloads.

How AWS AppConfig works

This section provides a high-level description of how AWS AppConfig works.



1. Identify configuration data to manage in AWS AppConfig

Before creating a configuration profile, identify the configuration data in your code that you want to manage dynamically using AWS AppConfig. Common examples include feature flags, allow and block lists, logging levels, service limits, and throttling rules. These values tend to change frequently and can cause issues if misconfigured.

If your configuration data already exists in cloud services such as SSM Parameter Store or Amazon S3, you can use AWS AppConfig to validate, deploy, and manage that data more effectively.

2. Create a configuration profile in AWS AppConfig

A configuration profile defines how AWS AppConfig locates and manages your configuration data. It includes a URI that points to the data source and a profile type.

AWS AppConfig supports two profile types

- **Feature flags** – Enable controlled feature releases, gradual rollouts, and testing in production.
- **Free-form configurations** – Store and retrieve configuration data from external sources and update it without redeploying code.

Both profile types help decouple configuration from code, support continuous delivery, and reduce deployment risk.

You can also add optional validators to ensure that configuration data is syntactically and semantically correct. During deployment, AWS AppConfig evaluates these validators and automatically rolls back changes if validation fails.

Each configuration profile is associated with an application, which acts as a logical container for your configuration resources. For more information about creating a configuration profile, see [Creating a configuration profile in AWS AppConfig](#).

3. Deploy configuration data

When you start a deployment, AWS AppConfig:

1. Retrieves configuration data from the source defined in the configuration profile
2. Validates the data using the configured validators
3. Delivers the validated configuration to AWS AppConfig Agent

The delivered configuration becomes the deployed version used by your application. For more information about deploying a configuration, see [Deploying feature flags and configuration data in AWS AppConfig](#).

4. Retrieve configuration data

Your application retrieves configuration data by calling a local endpoint exposed by AWS AppConfig Agent, which caches the deployed configuration. Retrieving data is a metered event. AWS AppConfig Agent supports a variety of use cases, as described in [How to use AWS AppConfig Agent to retrieve configuration data](#).

If the agent is not suitable for your use case, your application can retrieve configuration data directly from AWS AppConfig by calling the [StartConfigurationSession](#) and [GetLatestConfiguration](#) API actions.

For more information about retrieving a configuration, see [Retrieving feature flags and configuration data in AWS AppConfig](#).

Pricing for AWS AppConfig

Pricing for AWS AppConfig is pay-as-you-go based on configuration data and feature flag retrieval. We recommend using AWS AppConfig Agent to help optimize costs. For more information, see [AWS Systems Manager Pricing](#).

AWS AppConfig quotas

You can view information about AWS AppConfig endpoints and service quotas in the [Amazon Web Services General Reference](#).

Note

AWS AppConfig is a capability of AWS Systems Manager.

For information about quotas for services that store AWS AppConfig configurations, see [Understanding configuration store quotas and limitations](#).

Additional resources

The following resources can help you learn more about AWS AppConfig.

Blogs

The following blogs can help you learn more about AWS AppConfig and its capabilities:

- [Why you should use AWS AppConfig](#)
- [Unleash the power of feature flags with AWS AppConfig](#)
- [Using AWS AppConfig feature flags](#)
- [Best practices for validating AWS AppConfig feature flags and configuration data](#)

SDKs

For information about AWS AppConfig language-specific SDKs, see the following resources:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

Setting up AWS AppConfig

Sign up for an AWS account

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

Understanding IPv6 support

All AWS AppConfig APIs fully support IPv4 and IPv6 calls.

Control plane APIs

Use the following endpoint for IPv4 and IPv6 dual-stack calls to the [control plane](#):

```
appconfig.Region.api.aws
```

For example: appconfig.us-east-1.api.aws

For IPv4 only, use the following URL:

```
appconfig.Region.amazonaws.com
```

Data plane APIs

For dual-stack calls to the [data plane](#), use the following endpoint:

```
appconfigdata.Region.api.aws
```

For example: appconfig.us-east-1.api.aws

For IPv4 only, use the following URL:

```
appconfigdata.Region.amazonaws.com
```

Note

For more information, see [AWS AppConfig endpoints and quotas](#) in the *AWS General Reference*.

Configure permissions for automatic rollback

You can configure AWS AppConfig to roll back to a previous version of a configuration in response to one or more Amazon CloudWatch alarms. When you configure a deployment to respond to CloudWatch alarms, you specify an AWS Identity and Access Management (IAM) role. AWS AppConfig requires this role so that it can monitor CloudWatch alarms. This procedure is optional, but highly recommended.

Note

Note the following information.

- The IAM role must belong to the current account. By default, AWS AppConfig can only monitor alarms owned by the current account.
- For information about metrics to monitor and how to configure AWS AppConfig for automatic rollback, see [Monitoring deployments for automatic rollback](#).

Use the following procedures to create an IAM role that enables AWS AppConfig to rollback based on CloudWatch alarms. This section includes the following procedures.

1. [Step 1: Create the permission policy for rollback based on CloudWatch alarms](#)
2. [Step 2: Create the IAM role for rollback based on CloudWatch alarms](#)
3. [Step 3: Add a trust relationship](#)

Step 1: Create the permission policy for rollback based on CloudWatch alarms

Use the following procedure to create an IAM policy that gives AWS AppConfig permission to call the `DescribeAlarms` API action.

To create an IAM permission policy for rollback based on CloudWatch alarms

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Policies**, and then choose **Create policy**.
3. On the **Create policy** page, choose the **JSON** tab.
4. Replace the default content on the JSON tab with the following permission policy, and then choose **Next: Tags**.

Note

To return information about CloudWatch composite alarms, the [DescribeAlarms](#) API operation must be assigned * permissions, as shown here. You can't return information about composite alarms if DescribeAlarms has a narrower scope.

JSON

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

5. Enter tags for this role, and then choose **Next: Review**.
6. On the **Review** page, enter **SSMCloudWatchAlarmDiscoveryPolicy** in the **Name** field.
7. Choose **Create policy**. The system returns you to the **Policies** page.

Step 2: Create the IAM role for rollback based on CloudWatch alarms

Use the following procedure to create an IAM role and assign the policy you created in the previous procedure to it.

To create an IAM role for rollback based on CloudWatch alarms

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Roles**, and then choose **Create role**.
3. Under **Select type of trusted entity**, choose **AWS service**.
4. Immediately under **Choose the service that will use this role**, choose **EC2: Allows EC2 instances to call AWS services on your behalf**, and then choose **Next: Permissions**.
5. On the **Attached permissions policy** page, search for **SSMCloudWatchAlarmDiscoveryPolicy**.
6. Choose this policy and then choose **Next: Tags**.
7. Enter tags for this role, and then choose **Next: Review**.
8. On the **Create role** page, enter **SSMCloudWatchAlarmDiscoveryRole** in the **Role name** field, and then choose **Create role**.
9. On the **Roles** page, choose the role you just created. The **Summary** page opens.

Step 3: Add a trust relationship

Use the following procedure to configure the role you just created to trust AWS AppConfig.

To add a trust relationship for AWS AppConfig

1. In the **Summary** page for the role you just created, choose the **Trust Relationships** tab, and then choose **Edit Trust Relationship**.
2. Edit the policy to include only "appconfig.amazonaws.com", as shown in the following example:

JSON

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

```
]
}
```

3. Choose **Update Trust Policy**.

Creating feature flags and free form configuration data in AWS AppConfig

The topics in this section help you complete the following tasks in AWS AppConfig. These tasks create important artifacts for deploying configuration data.

1. [Create an application namespace](#)

To create an application namespace, you create an AWS AppConfig artifact called an application. An application is simply an organizational construct like a folder.

2. [Create environments](#)

For each AWS AppConfig application, you define one or more environments. An environment is a logical deployment group of AWS AppConfig targets, such as applications in a Beta or Production environment. You can also define environments for application subcomponents, such as AWS Lambda functions, Containers, Web, Mobile, and Back-end.

You can configure Amazon CloudWatch alarms for each environment to automatically rollback problematic configuration changes. The system monitors alarms during a configuration deployment. If an alarm is triggered, the system rolls back the configuration.

3. [Create a configuration profile](#)

Configuration data is a collection of settings that influence the behavior of your application. A *configuration profile* includes, among other things, a URI that enables AWS AppConfig to locate your configuration data in its stored location and a configure type. AWS AppConfig supports the following types of configuration profiles:

- **Feature flags:** You can use feature flags to enable or disable features within your applications or to configure different characteristics of your application features using flag attributes. AWS AppConfig stores feature flag configurations in the AWS AppConfig hosted configuration store in a feature flag format that contains data and metadata about your flags and the flag attributes. The URI for feature flag configurations is simply hosted.
- **Freeform configurations:** A freeform configuration can store data in any of the following AWS services and Systems Manager tools:
 - AWS AppConfig hosted configuration store
 - Amazon Simple Storage Service
 - AWS CodePipeline

- AWS Secrets Manager
- AWS Systems Manager (SSM) Parameter Store
- SSM Document Store

 Note

If possible, we recommend hosting your configuration data in the AWS AppConfig hosted configuration store as it offers the most features and enhancements.

4. (Optional, but recommended) [Create multi-variant feature flags](#)

AWS AppConfig offers basic feature flags, which (if enabled) return a specific set of configuration data per request. To better support user segmentation and traffic splitting use cases, AWS AppConfig also offers multi-variant feature flags, which enable you to define a set of possible flag values to return for a request. You can also configure different statuses (enabled or disabled) for multi-variant flags. When requesting a flag configured with variants, your application provides context that AWS AppConfig evaluates against a set of user-defined rules. Depending on the context specified in the request and the rules defined for the variant, AWS AppConfig returns different flag values to the application.

Topics

- [Understanding the configuration profile IAM role](#)
- [Creating a namespace for your application in AWS AppConfig](#)
- [Creating environments for your application in AWS AppConfig](#)
- [Creating a configuration profile in AWS AppConfig](#)

Understanding the configuration profile IAM role

You can create the IAM role that provides access to the configuration data by using AWS AppConfig. Or you can create the IAM role yourself. If you create the role by using AWS AppConfig, the system creates the role and specifies one of the following permissions policies, depending on which type of configuration source you choose.

Configuration source is a Secrets Manager secret

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:us-
east-1:111122223333:secret:secret_name-a1b2c3"
      ]
    }
  ]
}
```

Configuration source is a Parameter Store parameter

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ssm:GetParameter"
      ],
      "Resource": [
        "arn:aws:ssm:us-east-1:111122223333:parameter/parameter_name"
      ]
    }
  ]
}
```

Configuration source is an SSM document

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ssm:GetDocument"
      ],
      "Resource": [
        "arn:aws:ssm:us-east-1:111122223333:document/document_name"
      ]
    }
  ]
}
```

If you create the role by using AWS AppConfig, the system also creates the following trust relationship for the role.

JSON

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

Creating a namespace for your application in AWS AppConfig

The procedures in this section help you create an AWS AppConfig artifact called an application. An application is simply an organizational construct like a folder that identifies the namespace of your application. This organizational construct has a relationship with some unit of executable code. For example, you could create an application called `MyMobileApp` to organize and manage configuration data for a mobile application installed by your users. You must create these artifacts before you can use AWS AppConfig to deploy and retrieve feature flags or free form configuration data.

The following procedure gives you the option to associate an extension with a feature flag configuration profile. An *extension* augments your ability to inject logic or behavior at different points during the AWS AppConfig workflow of creating or deploying a configuration. For more information, see [Understanding AWS AppConfig extensions](#).

Note

You can use AWS CloudFormation to create AWS AppConfig artifacts, including applications, environments, configuration profiles, deployments, deployment strategies, and hosted configuration versions. For more information, see [AWS AppConfig resource type reference](#) in the *AWS CloudFormation User Guide*.

Topics

- [Creating an AWS AppConfig application \(console\)](#)
- [Creating an AWS AppConfig application \(command line\)](#)

Creating an AWS AppConfig application (console)

Use the following procedure to create an AWS AppConfig application by using the AWS Systems Manager console.

To create an application

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Applications**, and then choose **Create application**.

3. For **Name**, enter a name for the application.
4. For **Description**, enter information about the application.
5. (Optional) In the **Extensions** section, choose an extension from the list. For more information, see [Understanding AWS AppConfig extensions](#).
6. (Optional) In the **Tags** section, enter a key and an optional value. You can specify a maximum of 50 tags for a resource.
7. Choose **Create application**.

AWS AppConfig creates the application and then displays the **Environments** tab. Proceed to [Creating environments for your application in AWS AppConfig](#).

Creating an AWS AppConfig application (command line)

The following procedure describes how to use the AWS CLI (on Linux or Windows) or AWS Tools for PowerShell to create an AWS AppConfig application.

To create an application step by step

1. Open the AWS CLI.
2. Run the following command to create an application.

Linux

```
aws appconfig create-application \  
  --name A_name_for_the_application \  
  --description A_description_of_the_application \  
  --tags User_defined_key_value_pair_metadata_for_the_application
```

Windows

```
aws appconfig create-application ^  
  --name A_name_for_the_application ^  
  --description A_description_of_the_application ^  
  --tags User_defined_key_value_pair_metadata_for_the_application
```

PowerShell

```
New-APPCApplication `
```

```
-Name Name_for_the_application `
-Description Description_of_the_application `
-Tag Hashtable_type_user_defined_key_value_pair_metadata_for_the_application
```

The system returns information like the following.

Linux

```
{
  "Id": "Application ID",
  "Name": "Application name",
  "Description": "Description of the application"
}
```

Windows

```
{
  "Id": "Application ID",
  "Name": "Application name",
  "Description": "Description of the application"
}
```

PowerShell

```
ContentLength      : Runtime of the command
Description        : Description of the application
HttpStatusCode     : HTTP Status of the runtime
Id                 : Application ID
Name               : Application name
ResponseMetadata   : Runtime Metadata
```

Creating environments for your application in AWS AppConfig

For each AWS AppConfig application, you define one or more environments. An environment is a logical deployment group of AppConfig targets, such as applications in a Beta or Production environment, AWS Lambda functions, or containers. You can also define environments for application subcomponents, such as the Web, Mobile, and Back-end. You can configure Amazon

CloudWatch alarms for each environment. The system monitors alarms during a configuration deployment. If an alarm is triggered, the system rolls back the configuration.

Before You Begin

If you want to enable AWS AppConfig to roll back a configuration in response to a CloudWatch alarm, then you must configure an AWS Identity and Access Management (IAM) role with permissions to enable AWS AppConfig to respond to CloudWatch alarms. You choose this role in the following procedure. For more information, see [Configure permissions for automatic rollback](#).

Topics

- [Creating an AWS AppConfig environment \(console\)](#)
- [Creating an AWS AppConfig environment \(command line\)](#)

Creating an AWS AppConfig environment (console)

Use the following procedure to create an AWS AppConfig environment by using the AWS Systems Manager console.

To create an environment

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Applications**, and then choose the name of an application to open the details page.
3. Choose the **Environments** tab, and then choose **Create environment**.
4. For **Name**, enter a name for the environment.
5. For **Description**, enter information about the environment.
6. (Optional) In the **Monitors** section, choose the **IAM role** field, and then choose an IAM role with permission to call `cloudwatch:DescribeAlarms` on the metrics you want to monitor for alarms.
7. In the **CloudWatch alarms** list, enter the Amazon Resource Names (ARNs) one or more metrics to monitor. AWS AppConfig rolls back your configuration deployment if one of these metrics goes into an ALARM state. For information about recommended metrics, see [Monitoring deployments for automatic rollback](#)

8. (Optional) In the **Associate extensions** section, choose an extension from the list. For more information, see [Understanding AWS AppConfig extensions](#).
9. (Optional) In the **Tags** section, enter a key and an optional value. You can specify a maximum of 50 tags for a resource.
10. Choose **Create environment**.

AWS AppConfig creates the environment and then displays the **Environment details** page. Proceed to [Creating a configuration profile in AWS AppConfig](#).

Creating an AWS AppConfig environment (command line)

The following procedure describes how to use the AWS CLI (on Linux or Windows) or AWS Tools for PowerShell to create an AWS AppConfig environment.

To create an environment step by step

1. Open the AWS CLI.
2. Run the following command to create an environment.

Linux

```
aws appconfig create-environment \
  --application-id The_application_ID \
  --name A_name_for_the_environment \
  --description A_description_of_the_environment \
  --monitors
  "AlarmArn=ARN_of_the_Amazon_CloudWatch_alarm,AlarmArnRole=ARN_of_the_IAM
role_for_AWS AppConfig_to_monitor_AlarmArn" \
  --tags User_defined_key_value_pair_metadata_of_the_environment
```

Windows

```
aws appconfig create-environment ^
  --application-id The_application_ID ^
  --name A_name_for_the_environment ^
  --description A_description_of_the_environment ^
  --monitors
  "AlarmArn=ARN_of_the_Amazon_CloudWatch_alarm,AlarmArnRole=ARN_of_the_IAM
role_for_AWS AppConfig_to_monitor_AlarmArn" ^
```

```
--tags User_defined_key_value_pair_metadata_of_the_environment
```

PowerShell

```
New-APPCEnvironment `
-Name Name_for_the_environment `
-ApplicationId The_application_ID
-Description Description_of_the_environment `
-Monitors
@{"AlarmArn=ARN_of_the_Amazon_CloudWatch_alarm,AlarmArnRole=ARN_of_the_IAM
role_for_AWS_AppConfig_to_monitor_AlarmArn"} `
-Tag Hashtable_type_user_defined_key_value_pair_metadata_of_the_environment
```

The system returns information like the following.

Linux

```
{
  "ApplicationId": "The application ID",
  "Id": "The_environment ID",
  "Name": "Name of the environment",
  "State": "The state of the environment",
  "Description": "Description of the environment",

  "Monitors": [
    {
      "AlarmArn": "ARN of the Amazon CloudWatch alarm",
      "AlarmRoleArn": "ARN of the IAM role for AppConfig to monitor AlarmArn"
    }
  ]
}
```

Windows

```
{
  "ApplicationId": "The application ID",
  "Id": "The environment ID",
  "Name": "Name of the environment",
  "State": "The state of the environment"
  "Description": "Description of the environment",
```

```
"Monitors": [  
  {  
    "AlarmArn": "ARN of the Amazon CloudWatch alarm",  
    "AlarmRoleArn": "ARN of the IAM role for AppConfig to monitor AlarmArn"  
  }  
]  
}
```

PowerShell

```
ApplicationId      : The application ID  
ContentLength      : Runtime of the command  
Description        : Description of the environment  
HttpStatusCode     : HTTP Status of the runtime  
Id                : The environment ID  
Monitors           : {ARN of the Amazon CloudWatch alarm, ARN of the IAM role for  
  AppConfig to monitor AlarmArn}  
Name               : Name of the environment  
Response Metadata  : Runtime Metadata  
State              : State of the environment
```

Proceed to [Creating a configuration profile in AWS AppConfig](#).

Creating a configuration profile in AWS AppConfig

Configuration data is a collection of settings that influence the behavior of your application. A *configuration profile* includes, among other things, a URI that enables AWS AppConfig to locate your configuration data in its stored location and a configure type. AWS AppConfig supports the following types of configuration profiles:

- **Feature flags:** You can use feature flags to enable or disable features within your applications or to configure different characteristics of your application features using flag attributes. AWS AppConfig stores feature flag configurations in the AWS AppConfig hosted configuration store in a feature flag format that contains data and metadata about your flags and the flag attributes. The URI for feature flag configurations is simply hosted.
- **Freeform configurations:** A freeform configuration can store data in any of the following AWS services and Systems Manager tools:
 - AWS AppConfig hosted configuration store

- Amazon Simple Storage Service
- AWS CodePipeline
- AWS Secrets Manager
- AWS Systems Manager (SSM) Parameter Store
- SSM Document Store

Note

If possible, we recommend hosting your configuration data in the AWS AppConfig hosted configuration store as it offers the most features and enhancements.

Here are some configuration data samples to help you better understand different types of configuration data and how they can be used in either a feature flag or free from configuration profile.

Feature flag configuration data

The following feature flag configuration data enables or disables mobile payments and default payments on a per-region basis.

JSON

```
{
  "allow_mobile_payments": {
    "enabled": false
  },
  "default_payments_per_region": {
    "enabled": true
  }
}
```

YAML

```
---
allow_mobile_payments:
  enabled: false
default_payments_per_region:
```

```
enabled: true
```

Operational configuration data

The following freeform configuration data enforces limits on how an application processes requests.

JSON

```
{
  "throttle-limits": {
    "enabled": "true",
    "throttles": [
      {
        "simultaneous_connections": 12
      },
      {
        "tps_maximum": 5000
      }
    ],
    "limit-background-tasks": [
      true
    ]
  }
}
```

YAML

```
---
throttle-limits:
  enabled: 'true'
  throttles:
  - simultaneous_connections: 12
  - tps_maximum: 5000
  limit-background-tasks:
  - true
```

Access control list configuration data

The following access control list freeform configuration data specifies which users or groups can access an application.

JSON

```
{
  "allow-list": {
    "enabled": "true",
    "cohorts": [
      {
        "internal_employees": true
      },
      {
        "beta_group": false
      },
      {
        "recent_new_customers": false
      },
      {
        "user_name": "Jane_Doe"
      },
      {
        "user_name": "John_Doe"
      }
    ]
  }
}
```

YAML

```
---
allow-list:
  enabled: 'true'
  cohorts:
    - internal_employees: true
    - beta_group: false
    - recent_new_customers: false
    - user_name: Jane_Doe
    - user_name: Ashok_Kumar
```

Topics

- [Creating a feature flag configuration profile in AWS AppConfig](#)
- [Creating a free form configuration profile in AWS AppConfig](#)
- [Creating a configuration profile for non-native data sources](#)

Creating a feature flag configuration profile in AWS AppConfig

You can use feature flags to enable or disable features within your applications or to configure different characteristics of your application features using flag attributes. AWS AppConfig stores feature flag configurations in the AWS AppConfig hosted configuration store in a feature flag format that contains data and metadata about your flags and the flag attributes.

Note

When you create a feature flag configuration profile, you can create a basic feature flag as part of the configuration profile workflow. AWS AppConfig also supports multi-variant feature flags. *Multi-variant feature flags* enable you to define a set of possible flag values to return for a request. When requesting a flag configured with variants, your application provides context that AWS AppConfig evaluates against a set of user-defined rules.

Depending on the context specified in the request and the rules defined for the variant, AWS AppConfig returns different flag values to the application.

To create multi-variant feature flags, create a configuration profile first, and then edit any flags within the configuration profile to add variants. For more information, see [Creating multi-variant feature flags](#).

Topics

- [Understanding feature flag attributes](#)
- [Creating a feature flag configuration profile \(console\)](#)
- [Creating a feature flag configuration profile \(command line\)](#)
- [Creating multi-variant feature flags](#)
- [Understanding the type reference for `AWS.AppConfig.FeatureFlags`](#)
- [Saving a previous feature flag version to a new version](#)

Understanding feature flag attributes

When you create a feature flag configuration profile—or create a new flag within an existing configuration profile—you can specify attributes and corresponding constraints for the flag. An attribute is a field that you associate with your feature flag to express properties related to your feature flag. Attributes are delivered to your application with your flag key and the enable or disable value of the flag.

Constraints ensure that any unexpected attribute values are not deployed to your application. The following image shows an example.

Define attributes

Key: currency Type: String Constraint: CAD,USD,MXN Remove

☐ Required ☐ Regular expression ☒ Enum

Add new attribute

Attribute Values

Key: currency Key: CAD

Cancel Apply

Note

Note the following information about flag attributes.

- For attribute names, the word "enabled" is reserved. You can't create a feature flag attribute called "enabled". There are no other reserved words.
- The attributes of a feature flag are only included in the `GetLatestConfiguration` response if that flag is enabled.
- Flag attribute keys for a given flag must be unique.

AWS AppConfig supports the following types of flag attributes and their corresponding constraints.

Type	Constraint	Description
String	Regular Expression	Regex pattern for the string
	Enum	List of acceptable values for the string

Type	Constraint	Description
Number	Minimum	Minimum numeric value for the attribute
	Maximum	Maximum numeric value for the attribute
Boolean	None	None
String array	Regular Expression	Regex pattern for the elements of the array
	Enum	List of acceptable values for the elements of the array
Number array	Minimum	Minimum numeric value for the elements of the array
	Maximum	Maximum numeric value for the elements of the array

Creating a feature flag configuration profile (console)

Use the following procedure to create an AWS AppConfig feature flag configuration profile by using the AWS AppConfig console. At the time you create the configuration profile, you can also create a basic feature flag.

To create a configuration profile

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Applications**, and then choose an application you created in [Creating a namespace for your application in AWS AppConfig](#).
3. On the **Configuration profiles and feature flags** tab, choose **Create configuration**.
4. In the **Configuration options** section, choose **Feature flag**.
5. In the **Configuration profile** section, for **Configuration profile name**, enter a name.
6. (Optional) Expand **Description** and enter a description.

7. (Optional) Expand **Additional options** and complete the following, as necessary.
 - a. In the **Encryption** list, choose an AWS Key Management Service (AWS KMS) key from the list. This customer managed key enables you to encrypt new configuration data versions in the AWS AppConfig hosted configuration store. For more information about this key, see **AWS AppConfig supports customer manager keys** in [Security in AWS AppConfig](#).
 - b. In the **Tags** section, choose **Add new tag**, and then specify a key and optional value.
8. Choose **Next**.
9. In the **Feature flag definition** section, for **Flag name**, enter a name.
10. For **Flag key** enter a flag identifier to distinguish flags within the same configuration profile. Flags within the same configuration profile can't have the same key. After the flag is created, you can edit the flag name, but not the flag key.
11. (Optional) Expand **Description** and enter information about this flag.
12. Select **This is a short-term flag** and optionally choose a date for when the flag should be disabled or deleted. AWS AppConfig does *not* disable the flag on the deprecation date.
13. (Optional) In the **Feature flag attributes** section, choose **Define attribute**. Attributes enable you to provide additional values within your flag. For more information about attributes and constraints, see [Understanding feature flag attributes](#).
 - a. For **Key**, specify a flag key and choose its type from the **Type** list. For information about the supported options for the **Value** and **Constraints** fields, see the previously referenced section about attributes.
 - b. Select **Required value** to specify whether an attribute value is required.
 - c. Choose **Define attribute** to add additional attributes.
14. In the **Feature flag value** section, choose **Enabled** to enable the flag. Use this same toggle to disable a flag when it reaches a specified deprecation date, if applicable.
15. Choose **Next**.
16. On the **Review and save** page, verify the details of the flag and then **Save and continue to deploy**.

Proceed to [Deploying feature flags and configuration data in AWS AppConfig](#).

Creating a feature flag configuration profile (command line)

The following procedure describes how to use the AWS Command Line Interface (on Linux or Windows) or Tools for Windows PowerShell to create an AWS AppConfig feature flag configuration profile. At the time you create the configuration profile, you can also create a basic feature flag.

To create a feature flag configuration

1. Open the AWS CLI.
2. Create a feature flag configuration profile specifying its **Type** as `AWS.AppConfig.FeatureFlags`. The configuration profile must use `hosted` for the location URI.

Linux

```
aws appconfig create-configuration-profile \  
  --application-id APPLICATION_ID \  
  --name CONFIGURATION_PROFILE_NAME \  
  --location-uri hosted \  
  --type AWS.AppConfig.FeatureFlags
```

Windows

```
aws appconfig create-configuration-profile ^  
  --application-id APPLICATION_ID ^  
  --name CONFIGURATION_PROFILE_NAME ^  
  --location-uri hosted ^  
  --type AWS.AppConfig.FeatureFlags
```

PowerShell

```
New-APPCConfigurationProfile `\  
  -Name CONFIGURATION_PROFILE_NAME `\  
  -ApplicationId APPLICATION_ID `\  
  -LocationUri hosted `\  
  -Type AWS.AppConfig.FeatureFlags
```

3. Create your feature flag configuration data. Your data must be in a JSON format and conform to the `AWS.AppConfig.FeatureFlags` JSON schema. For more information about the schema, see [Understanding the type reference for AWS.AppConfig.FeatureFlags](#).
4. Use the `CreateHostedConfigurationVersion` API to save your feature flag configuration data to AWS AppConfig.

Linux

```
aws appconfig create-hosted-configuration-version \  
  --application-id APPLICATION_ID \  
  --configuration-profile-id CONFIGURATION_PROFILE_ID \  
  --content-type "application/json" \  
  --content file://path/to/feature_flag_configuration_data.json \  
  --cli-binary-format raw-in-base64-out
```

Windows

```
aws appconfig create-hosted-configuration-version ^  
  --application-id APPLICATION_ID ^  
  --configuration-profile-id CONFIGURATION_PROFILE_ID ^  
  --content-type "application/json" ^  
  --content file://path/to/feature_flag_configuration_data.json ^  
  --cli-binary-format raw-in-base64-out
```

PowerShell

```
New-APPCHostedConfigurationVersion `   
  -ApplicationId APPLICATION_ID `   
  -ConfigurationProfileId CONFIGURATION_PROFILE_ID `   
  -ContentType "application/json" `   
  -Content file://path/to/feature_flag_configuration_data.json
```

The command loads the content specified for the `Content` parameter from disk. The content must be similar to the following example.

```
{  
  "flags": {  
    "ui_refresh": {  
      "name": "UI Refresh"    }  
  }  
}
```

```
    }
  },
  "values": {
    "ui_refresh": {
      "enabled": false,
      "attributeValues": {
        "dark_mode_support": true
      }
    }
  },
  "version": "1"
}
```

The system returns information like the following.

Linux

```
{
  "ApplicationId"      : "ui_refresh",
  "ConfigurationProfileId" : "UI Refresh",
  "VersionNumber"      : "1",
  "ContentType"        : "application/json"
}
```

Windows

```
{
  "ApplicationId"      : "ui_refresh",
  "ConfigurationProfileId" : "UI Refresh",
  "VersionNumber"      : "1",
  "ContentType"        : "application/json"
}
```

PowerShell

```
ApplicationId      : ui_refresh
ConfigurationProfileId : UI Refresh
VersionNumber      : 1
ContentType        : application/json
```

The `service_returned_content_file` contains your configuration data that includes some AWS AppConfig generated metadata.

Note

When you create the hosted configuration version, AWS AppConfig verifies that your data conforms to the `AWS.AppConfig.FeatureFlags` JSON schema. AWS AppConfig additionally validates that each feature flag attribute in your data satisfies the constraints you defined for those attributes.

Creating multi-variant feature flags

Feature flag variants enable you to define a set of possible flag values to return for a request. You can also configure different statuses (enabled or disabled) for multi-variant flags. When requesting a flag configured with variants, your application provides context that AWS AppConfig evaluates against a set of user-defined rules. Depending on the context specified in the request and the rules defined for the variant, AWS AppConfig returns different flag values to the application.

The following screenshot shows an example of a feature flag with three user-defined variants and the default variant.

Feature flag variants Info					Reorder variant up	Reorder variant down	Edit	Create variant
	Name	Enabled value	Attribute values	Rule				
☐	beta testers	☒ ON	-	(or (eq \$userId "Alice") (eq \$userId "123456789012"))				
☐	EU demographic	☒ ON	-	(and (ends_with \$email "@example.com") (eq \$continent "EU"))				
☐	QA testing	☒ ON	-	(and (matches pattern: ".*@example\\.com" in::\$email) (contains \$roles "Engineer") (gt \$tenure 5))				
☐	default	☒ ON	-	-				

Variant order is used for evaluation logic

Variants are evaluated as an ordered list based on the order shown and any specified rules. The variant at the top of the list is evaluated first. If no rules match the supplied context, AWS AppConfig returns the default variant.

Topics

- [Understanding multi-variant feature flag concepts and common use cases](#)
- [Understanding multi-variant feature flag rules](#)

- [Creating a multi-variant feature flag](#)

Understanding multi-variant feature flag concepts and common use cases

To help you better understand feature flag variants, this section explains flag variant concepts and common use cases.

Concepts

- **Feature flag:** An AWS AppConfig configuration type used to control the behavior of a feature in an application. A flag has a status (enabled or disabled) and an optional set of attributes containing arbitrary string, numeric, boolean, or array values.
- **Feature flag variant:** A specific combination of status and attribute values belonging to a feature flag. A feature flag may have multiple variants.
- **Variant rule:** A user-defined expression used to select a feature flag variant. Each variant has its own rule that AWS AppConfig evaluates to determine whether to return it or not.
- **Default variant:** A special variant that is returned when no other variant is selected. All multi-variant feature flags have a default variant.

Note, the default variant must be last in your ordering of variants, and it can't have rules associated with it. If it's not defined last, AWS AppConfig returns a `BadRequestException` when you try to create the multi-variant flag.

- **Context:** User-defined keys and values passed to AWS AppConfig at configuration retrieval time. Context values are used during rule evaluation to select the feature flag variant to return.

Note

AWS AppConfig agent evaluates variant rules and determines which rule applies to the request based on the provided context. For more information about retrieving multi-variant feature flags, see [Retrieving basic and multi-variant feature flags](#).

Common use cases

This section describes two common use cases for feature flag variants.

User segmentation

User segmentation is the process of dividing users based on certain attributes. As an example, you could use flag variants to expose a feature to some users but not others based on their user ID, geographic location, device type, or purchase frequency.

Using the example of purchase frequency, suppose your commerce application supports a feature to increase customer loyalty. You can use flag variants to configure different incentive types to be shown to a user based on when they last purchased something. A new user might be offered a small discount to encourage them to become a customer, whereas a repeat customer might be given a larger discount if they purchase something from a new category.

Traffic splitting

Traffic splitting is the process of selecting a random, but consistent, flag variant based on a context value you define. For example, you may want to perform an experiment where a small percentage of your users (identified by their user ID) sees a particular variant. Or, you may want to execute a gradual feature rollout where a feature is first exposed to 5% of your users, then 15%, then 40%, then 100%, while maintaining a consistent user experience throughout the rollout.

Using the experimentation example, you could use flag variants to test a new button style for the primary action on your application homepage to see if it drives more clicks. For your experiment, you could create a flag variant with a traffic splitting rule that selects 5% of users to see the new style, while the default variant indicates the users that should continue to see the existing style. If the experiment is successful, you can increase the percentage value, or even turn that variant into the default.

Understanding multi-variant feature flag rules

When you create a feature flag variant, you specify a rule for it. Rules are expressions that take context values as input and produce a boolean result as output. For example, you could define a rule to select a flag variant for beta users, identified by their account ID, testing a user interface refresh. For this scenario, you do the following:

1. Create a new feature flag configuration profile called *UI Refresh*.
2. Create a new feature flag called *ui_refresh*.
3. Edit the feature flag after you create it to add variants.
4. Create and enable a new variant called *BetaUsers*.
5. Define a rule for *BetaUsers* that selects the variant if the account ID from the request context is in a list of account IDs approved to view the new beta experience.
6. Confirm that the default variant's status is set to **Disabled**.

Note

Variants are evaluated as an ordered list based on the order they are defined in the console. The variant at the top of the list is evaluated first. If no rules match the supplied context, AWS AppConfig returns the Default variant.

When AWS AppConfig processes the feature flag request, it compares the supplied context, which includes the AccountID (for this example) to the BetaUsers variant first. If the context matches the rule for BetaUsers, AWS AppConfig returns the configuration data for the beta experience. If the context doesn't include an account ID or if the account ID ends in anything other than 123, AWS AppConfig returns configuration data for the Default rule, which means the user views the current experience in production.

Note

For information about retrieving multi-variant feature flags, see [Retrieving basic and multi-variant feature flags](#).

Defining rules for multi-variant feature flags

A variant rule is an expression comprised of one or more operands and an operator. An operand is a specific value used during the evaluation of a rule. Operand values can be either static, such as a literal number or string, or variable, such as the value found in a context or the result of another expression. An operator, such as "greater than", is a test or action applied to its operands that produces a value. A variant rule expression must produce either a "true" or "false" to be valid.

Operands

Type	Description	Example
String	A sequence of UTF-8 characters, enclosed in double-quotes.	"apple", "###ë# ##š##"
Integer	A 64-bit integer value.	-7, 42

Type	Description	Example
Float	A 64-bit IEEE-754 floating-point value.	3.14, 1.234e-5
Timestamp	A specific moment in time as described by the W3C note on date and time formats .	2012-03-04T05:06:07-08:00, 2024-01
Boolean	A true or false value.	true, false
Context value	A parameterized value in the form of <code>\$key</code> that is retrieved from the context during rule evaluation.	\$country, \$userId

Comparison operators

Operator	Description	Example
eq	Determines whether a context value is equal to a given value.	(eq \$state "Virginia")
gt	Determines whether a context value is greater than a given value.	(gt \$age 65)
gte	Determines whether a context value is greater than or equal to a given value.	(gte \$age 65)
lt	Determines whether a context value is less than a given value.	(lt \$age 65)

Operator	Description	Example
lte	Determines whether a context value is less than or equal to a given value.	<pre>(lte \$age 65)</pre>

Logical operators

Operator	Description	Example
and	Determines if both operands are true.	<pre>(and (eq \$state "Virginia") (gt \$age 65))</pre>
or	Determines if at least one of the operands is true.	<pre>(or (eq \$state "Virginia") (gt \$age 65))</pre>
not	Reverses the value of an expression.	<pre>(not (eq \$state "Virginia"))</pre>

Custom operators

Operator	Description	Example
begins_with	Determines whether a context value begins with a given prefix.	<pre>(begins_with \$state "A")</pre>
ends_with	Determines whether a context value ends with a given prefix.	<pre>(ends_with \$email "amazon.com")</pre>

Operator	Description	Example
contains	Determines whether a context value contains a given substring.	<pre>(contains \$promoCode "WIN")</pre>
in	Determines whether a context value is contained within a list of constants.	<pre>(in \$userId ["123", "456"])</pre>
matches	Determines whether a context value matches a given regex pattern.	<pre>(matches in::\$greeting pattern::"h.*y")</pre>
exists	Determines whether any value was provided for a context key.	<pre>(exists key::"country")</pre>

Operator	Description	Example
split	Evaluates to true for a given percentage of traffic based on a consistent hash of the provided context value(s). For a detailed explanation of how <code>split</code> works, see the next section in this topic, Understanding the split operator. Note that <code>seed</code> is an optional property. If you don't specify <code>seed</code>, the hash is <i>locally</i> consistent, meaning traffic will be split consistently for that flag, but other flags receiving the same context value may split traffic differently. If <code>seed</code> is provided, each unique value is guaranteed to split traffic consistently across feature flags, configuration profiles, and AWS accounts.	<pre>(split pct::10 by::\$userId seed::"abc")</pre>

Understanding the split operator

The following section describes how the `split` operator behaves when used in different scenarios. As a reminder, `split` evaluates to true for a given percentage of traffic based on a consistent hash of the provided context value. To understand this better, consider the following baseline scenario that uses `split` with two variants:

```
A: (split by::$uniqueId pct::20)
C: <no rule>
```

As expected, providing a random set of `uniqueId` values produces a distribution that's approximately:

```
A: 20%  
C: 80%
```

If you add a third variant, but use the same split percentage like so:

```
A: (split by::$uniqueId pct::20)  
B: (split by::$uniqueId pct::20)  
C: <default>
```

You end up with the following distribution:

```
A: 20%  
B: 0%  
C: 80%
```

This potentially unexpected distribution happens because each variant rule is evaluated in order and the first match determines the returned variant. When rule A is evaluated, 20% of `uniqueId` values match it, so the first variant is returned. Next, rule B is evaluated. However, all of the `uniqueId` values that would have matched the second split statement were already matched by variant rule A, so no values match B. The default variant is returned instead.

Now consider a third example.

```
A: (split by::$uniqueId pct::20)  
B: (split by::$uniqueId pct::25)  
C: <default>
```

As with the previous example, the first 20% of `uniqueId` values match rule A. For variant rule B, 25% of all `uniqueId` values would match, but most of those previously matched rule A. That leaves 5% of the total for variant B, with the remainder receiving variant C. The distribution would look like the following:

```
A: 20%
```

```
B: 5%  
C: 75%
```

Using the seed property

You can use the seed property to ensure traffic is split consistently for a given context value irrespective of where the split operator is used. If you don't specify seed, the hash is *locally* consistent, meaning traffic will be split consistently for that flag, but other flags receiving the same context value may split traffic differently. If seed is provided, each unique value is guaranteed to split traffic consistently across feature flags, configuration profiles, and AWS accounts.

Typically, customers use the same seed value across variants within a flag when splitting traffic on the same context property. However, it may occasionally make sense to use a different seed value. Here is an example that uses different seeds for rules A and B:

```
A: (split by::$uniqueId pct::20 seed::"seed_one")  
B: (split by::$uniqueId pct::25 seed::"seed_two")  
C: <default>
```

As before, 20% of the matching uniqueId values match rule A. That means 80% of values fall through and are tested against variant rule B. Because the seed is different, there is no correlation between the values that matched A and the values that match B. There are, however, only 80% as many uniqueId values to split with 25% of that number matching rule B and 75% not. That works out to the following distribution:

```
A: 20%  
B: 20% (25% of what falls through from A, or 25% of 80%)  
C: 60%
```

Creating a multi-variant feature flag

Use the procedures in this section to create variants of a feature flag.

Before you begin

Note the following important information.

- You can create variants of existing feature flags by editing them. You can't create variants of a new feature flag *when you create a new configuration profile*. You must complete the workflow

of creating the new configuration profile first. After you create the configuration profile, you can add variants to any flag within the configuration profile. For information about how to create a new configuration profile, see [Creating a feature flag configuration profile in AWS AppConfig](#).

- To retrieve feature flag variant data for Amazon EC2, Amazon ECS, and Amazon EKS compute platforms, you must use AWS AppConfig Agent version 2.0.4416 or later.
- For performance reasons, AWS CLI and SDK calls to AWS AppConfig don't retrieve variant data. For more information about AWS AppConfig Agent, see [How to use AWS AppConfig Agent to retrieve configuration data](#).
- When you create a feature flag variant, you specify a rule for it. Rules are expressions that take request context as input and produce a boolean result as output. Before you create variants, review the supported operands and operators for flag variant rules. You can create rules before you create variants. For more information, see [Understanding multi-variant feature flag rules](#).

Topics

- [Creating a multi-variant feature flag \(console\)](#)
- [Creating a multi-variant feature flag \(command line\)](#)

Creating a multi-variant feature flag (console)

The following procedure describes how to create a multi-variant feature flag for an existing configuration profile by using the AWS AppConfig console. You can also edit existing feature flags to create variants.

To create a multi-variant feature flag

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Applications**, and then choose an application.
3. On the **Configuration profiles and feature flags** tab, choose an existing feature flag configuration profile.
4. In the **Flags** section, choose **Add new flag**.
5. In the **Feature flag definition** section, for **Flag name**, enter a name.
6. For **Flag key** enter a flag identifier to distinguish flags within the same configuration profile. Flags within the same configuration profile can't have the same key. After the flag is created, you can edit the flag name, but not the flag key.

7. (Optional) In the **Description** field, enter information about this flag.
8. In the **Variants** section, choose **Multi-variant flag**.
9. (Optional) In the **Feature flag attributes** section, choose **Define attribute**. Attributes enable you to provide additional values within your flag. For more information about attributes and constraints, see [Understanding feature flag attributes](#).
 - a. For **Key**, specify a flag key and choose its type from the **Type** list. For information about the supported options for the **Value** and **Constraints** fields, see the previously referenced section about attributes.
 - b. Select **Required value** to specify whether an attribute value is required.
 - c. Choose **Define attribute** to add additional attributes.
 - d. Choose **Apply** to save attribute changes.
10. In the **Feature flag variants** section, choose **Create variant**.
 - a. For **Variant name**, enter a name.
 - b. Use the **Enabled value** toggle to enable the variant.
 - c. In the **Rule** text box, enter a rule.
 - d. Use the **Create variant > Create variant above** or **Create variant below** options to create additional variants for this flag.
 - e. In the **Default variant** section, use the **Enabled value** toggle to enable the default variant. Optionally, provide values for attributes defined in step 10.
 - f. Choose **Apply**.
11. Verify the details of the flag and its variants and choose **Create flag**.

For information about deploying your new feature flag with variants, see [Deploying feature flags and configuration data in AWS AppConfig](#).

Creating a multi-variant feature flag (command line)

The following procedure describes how to use the AWS Command Line Interface (on Linux or Windows) or Tools for Windows PowerShell to create a multi-variant feature flag for an existing configuration profile. You can also edit existing feature flags to create variants.

Before you begin

Complete the following tasks before you create a multi-variant feature flag by using the AWS CLI.

- Create a feature flag configuration profile. For more information, see [Creating a feature flag configuration profile in AWS AppConfig](#).
- Update to the latest version of the AWS CLI. For more information, see [Install or update to the latest version of the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

To create a multi-variant feature flag

1. Create a configuration file on your local machine that specifies the details of the multi-variant flag you want to create. Save the file with a `.json` file extension. The file must adhere to the [AWS.AppConfig.FeatureFlags](#) JSON schema. The schema contents of your configuration file will be similar to the following.

```
{
  "flags": {
    "FLAG_NAME": {
      "attributes": {
        "ATTRIBUTE_NAME": {
          "constraints": {
            "type": "CONSTRAINT_TYPE"
          }
        }
      },
      "description": "FLAG_DESCRIPTION",
      "name": "VARIANT_NAME"
    }
  },
  "values": {
    "VARIANT_VALUE_NAME": {
      "_variants": [
        {
          "attributeValues": {
            "ATTRIBUTE_NAME": BOOLEAN
          },
          "enabled": BOOLEAN,
          "name": "VARIANT_NAME",
          "rule": "VARIANT_RULE"
        },
        {
          "attributeValues": {
            "ATTRIBUTE_NAME": BOOLEAN
          }
        }
      ]
    }
  }
}
```

```

        "enabled": BOOLEAN,
        "name": "VARIANT_NAME",
        "rule": "VARIANT_RULE"
    },
    {
        "attributeValues": {
            "ATTRIBUTE_NAME": BOOLEAN
        },
        "enabled": BOOLEAN,
        "name": "VARIANT_NAME",
        "rule": "VARIANT_RULE"
    },
    {
        "attributeValues": {
            "ATTRIBUTE_NAME": BOOLEAN
        },
        "enabled": BOOLEAN,
        "name": "VARIANT_NAME",
        "rule": "VARIANT_RULE"
    }
]
}
},
"version": "VERSION_NUMBER"
}

```

Here is an example with three variants and the default variant.

```

{
  "flags": {
    "ui_refresh": {
      "attributes": {
        "dark_mode_support": {
          "constraints": {
            "type": "boolean"
          }
        }
      }
    },
    "description": "A release flag used to release a new UI",
    "name": "UI Refresh"
  }
},
"values": {

```

```

    "ui_refresh": {
      "_variants": [
        {
          "attributeValues": {
            "dark_mode_support": true
          },
          "enabled": true,
          "name": "QA",
          "rule": "(ends_with $email \"qa-testers.mycompany.com\")"
        },
        {
          "attributeValues": {
            "dark_mode_support": true
          },
          "enabled": true,
          "name": "Beta Testers",
          "rule": "(exists key::\"opted_in_to_beta\")"
        },
        {
          "attributeValues": {
            "dark_mode_support": false
          },
          "enabled": true,
          "name": "Sample Population",
          "rule": "(split pct::10 by::$email)"
        },
        {
          "attributeValues": {
            "dark_mode_support": false
          },
          "enabled": false,
          "name": "Default Variant"
        }
      ]
    },
    "version": "1"
  }
}

```

2. Use the `CreateHostedConfigurationVersion` API to save your feature flag configuration data to AWS AppConfig.

Linux

```
aws appconfig create-hosted-configuration-version \  
  --application-id APPLICATION_ID \  
  --configuration-profile-id CONFIGURATION_PROFILE_ID \  
  --content-type "application/json" \  
  --content file://path/to/feature_flag_configuration_data.json \  
  --cli-binary-format raw-in-base64-out \  
  outfile
```

Windows

```
aws appconfig create-hosted-configuration-version ^  
  --application-id APPLICATION_ID ^  
  --configuration-profile-id CONFIGURATION_PROFILE_ID ^  
  --content-type "application/json" ^  
  --content file://path/to/feature_flag_configuration_data.json ^  
  --cli-binary-format raw-in-base64-out ^  
  outfile
```

PowerShell

```
New-APPCHostedConfigurationVersion `   
  -ApplicationId APPLICATION_ID `   
  -ConfigurationProfileId CONFIGURATION_PROFILE_ID `   
  -ContentType "application/json" `   
  -Content file://path/to/feature_flag_configuration_data.json `   
  -Raw
```

The `service_returned_content_file` contains your configuration data that includes some AWS AppConfig generated metadata.

Note

When you create the hosted configuration version, AWS AppConfig verifies that your data conforms to the [AWS.AppConfig.FeatureFlags](#) JSON schema. AWS AppConfig additionally validates that each feature flag attribute in your data satisfies the constraints you defined for those attributes.

Understanding the type reference for AWS.AppConfig.FeatureFlags

Use the `AWS.AppConfig.FeatureFlags` JSON schema as a reference to create your feature flag configuration data.

```
{
  "$schema": "http://json-schema.org/draft-07/schema#",
  "definitions": {
    "flagSetDefinition": {
      "type": "object",
      "properties": {
        "version": {
          "$ref": "#/definitions/flagSchemaVersions"
        },
        "flags": {
          "$ref": "#/definitions/flagDefinitions"
        },
        "values": {
          "$ref": "#/definitions/flagValues"
        }
      },
      "required": ["version"],
      "additionalProperties": false
    },
    "flagDefinitions": {
      "type": "object",
      "patternProperties": {
        "^[a-z][a-zA-Z\\d_-]{0,63}$": {
          "$ref": "#/definitions/flagDefinition"
        }
      },
      "additionalProperties": false
    },
    "flagDefinition": {
      "type": "object",
      "properties": {
        "name": {
          "$ref": "#/definitions/customerDefinedName"
        },
        "description": {
          "$ref": "#/definitions/customerDefinedDescription"
        },
        "_createdAt": {
```

```

        "type": "string"
    },
    "_updatedAt": {
        "type": "string"
    },
    "_deprecation": {
        "type": "object",
        "properties": {
            "status": {
                "type": "string",
                "enum": ["planned"]
            },
            "date": {
                "type": "string",
                "format": "date"
            }
        },
        "additionalProperties": false
    },
    "attributes": {
        "$ref": "#/definitions/attributeDefinitions"
    },
    "additionalProperties": false
},
"attributeDefinitions": {
    "type": "object",
    "patternProperties": {
        "^[a-z][a-zA-Z\\d_-]{0,63}$": {
            "$ref": "#/definitions/attributeDefinition"
        }
    },
    "maxProperties": 25,
    "additionalProperties": false
},
"attributeDefinition": {
    "type": "object",
    "properties": {
        "description": {
            "$ref": "#/definitions/customerDefinedDescription"
        },
        "constraints": {
            "oneOf": [
                { "$ref": "#/definitions/numberConstraints" },

```

```

        { "$ref": "#/definitions/stringConstraints" },
        { "$ref": "#/definitions/arrayConstraints" },
        { "$ref": "#/definitions/boolConstraints" }
    ]
}
},
"additionalProperties": false
},
"flagValues": {
    "type": "object",
    "patternProperties": {
        "^[a-z][a-zA-Z\\d_-]{0,63}$": {
            "$ref": "#/definitions/flagValue"
        }
    },
    "additionalProperties": false
},
"flagValue": {
    "type": "object",
    "properties": {
        "enabled": {
            "type": "boolean"
        },
        "_createdAt": {
            "type": "string"
        },
        "_updatedAt": {
            "type": "string"
        },
        "_variants": {
            "type": "array",
            "maxLength": 32,
            "items": {
                "$ref": "#/definitions/variant"
            }
        }
    },
    "patternProperties": {
        "^[a-z][a-zA-Z\\d_-]{0,63}$": {
            "$ref": "#/definitions/attributeValue",
            "maxProperties": 25
        }
    },
    "additionalProperties": false

```

```

},
"attributeValue": {
  "oneOf": [
    { "type": "string", "maxLength": 1024 },
    { "type": "number" },
    { "type": "boolean" },
    {
      "type": "array",
      "oneOf": [
        {
          "items": {
            "type": "string",
            "maxLength": 1024
          }
        },
        {
          "items": {
            "type": "number"
          }
        }
      ]
    }
  ],
  "additionalProperties": false
},
"stringConstraints": {
  "type": "object",
  "properties": {
    "type": {
      "type": "string",
      "enum": ["string"]
    },
    "required": {
      "type": "boolean"
    },
    "pattern": {
      "type": "string",
      "maxLength": 1024
    },
    "enum": {
      "type": "array",
      "maxLength": 100,
      "items": {
        "oneOf": [

```

```

        {
            "type": "string",
            "maxLength": 1024
        },
        {
            "type": "integer"
        }
    ]
}
},
"required": ["type"],
"not": {
    "required": ["pattern", "enum"]
},
"additionalProperties": false
},
"numberConstraints": {
    "type": "object",
    "properties": {
        "type": {
            "type": "string",
            "enum": ["number"]
        },
        "required": {
            "type": "boolean"
        },
        "minimum": {
            "type": "integer"
        },
        "maximum": {
            "type": "integer"
        }
    },
    "required": ["type"],
    "additionalProperties": false
},
"arrayConstraints": {
    "type": "object",
    "properties": {
        "type": {
            "type": "string",
            "enum": ["array"]
        },
    },

```

```

    "required": {
      "type": "boolean"
    },
    "elements": {
      "$ref": "#/definitions/elementConstraints"
    }
  },
  "required": ["type"],
  "additionalProperties": false
},
"boolConstraints": {
  "type": "object",
  "properties": {
    "type": {
      "type": "string",
      "enum": ["boolean"]
    },
    "required": {
      "type": "boolean"
    }
  },
  "required": ["type"],
  "additionalProperties": false
},
"elementConstraints": {
  "oneOf": [
    { "$ref": "#/definitions/numberConstraints" },
    { "$ref": "#/definitions/stringConstraints" }
  ]
},
"variant": {
  "type": "object",
  "properties": {
    "enabled": {
      "type": "boolean"
    },
    "name": {
      "$ref": "#/definitions/customerDefinedName"
    },
    "rule": {
      "type": "string",
      "maxLength": 16384
    },
    "attributeValues": {

```

```

    "type": "object",
    "patternProperties": {
      "^[a-z][a-zA-Z\\d_-]{0,63}$": {
        "$ref": "#/definitions/attributeValue"
      }
    },
    "maxProperties": 25,
    "additionalProperties": false
  },
  "required": ["name", "enabled"],
  "additionalProperties": false
},
"customerDefinedName": {
  "type": "string",
  "pattern": "^[^\\n]{1,64}$"
},
"customerDefinedDescription": {
  "type": "string",
  "maxLength": 1024
},
"flagSchemaVersions": {
  "type": "string",
  "enum": ["1"]
}
},
"type": "object",
"$ref": "#/definitions/flagSetDefinition",
"additionalProperties": false
}

```

Important

To retrieve feature flag configuration data, your application must call the `GetLatestConfiguration` API. You can't retrieve feature flag configuration data by calling `GetConfiguration`, which is deprecated. For more information, see [GetLatestConfiguration](#) in the *AWS AppConfig API Reference*.

When your application calls [GetLatestConfiguration](#) and receives a newly deployed configuration, the information that defines your feature flags and attributes is removed. The simplified JSON contains a map of keys that match each of the flag keys you specified. The simplified JSON also

contains mapped values of true or false for the enabled attribute. If a flag sets enabled to true, any attributes of the flag will be present as well. The following JSON schema describes the format of the JSON output.

```
{
  "$schema": "http://json-schema.org/draft-07/schema#",
  "type": "object",
  "patternProperties": {
    "^[a-z][a-zA-Z\\d_-]{0,63}$": {
      "$ref": "#/definitions/attributeValuesMap"
    }
  },
  "additionalProperties": false,
  "definitions": {
    "attributeValuesMap": {
      "type": "object",
      "properties": {
        "enabled": {
          "type": "boolean"
        }
      }
    },
    "required": ["enabled"],
    "patternProperties": {
      "^[a-z][a-zA-Z\\d_-]{0,63}$": {
        "$ref": "#/definitions/attributeValue"
      }
    }
  },
  "maxProperties": 25,
  "additionalProperties": false
},
"attributeValue": {
  "oneOf": [
    { "type": "string", "maxLength": 1024 },
    { "type": "number" },
    { "type": "boolean" },
    {
      "type": "array",
      "oneOf": [
        {
          "items": {
            "oneOf": [
              {
                "type": "string",
```

```
        "maxLength": 1024
      }
    ]
  },
  {
    "items": {
      "oneOf": [
        {
          "type": "number"
        }
      ]
    }
  }
],
"additionalProperties": false
}
}
```

Saving a previous feature flag version to a new version

When you update a feature flag, AWS AppConfig automatically saves your changes to a new version. If you want to use a previous feature flag version, you must copy it to a draft version and then save it. You can't edit and save changes to a previous flag version without saving it to a new version.

To edit a previous feature flag version and save it to a new version

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Applications**, and then choose the application with the feature flag you want to edit and save to a new version.
3. On the **Configuration profiles and feature flags** tab, choose the configuration profile with the feature flag you want to edit and save to a new version.
4. On the **Feature flags** tab, use the **Version** list to choose the version you want to edit and save to a new version.
5. Choose **Copy to draft version**.

6. In the **Version label** field, enter a new label (optional, but recommended).
7. In the **Version description** field, enter a new description (optional, but recommended).
8. Choose **Save version**.
9. Choose **Start deployment** to deploy the new version.

Creating a free form configuration profile in AWS AppConfig

Configuration data is a collection of settings that influence the behavior of your application. A *configuration profile* includes, among other things, a URI that enables AWS AppConfig to locate your configuration data in its stored location and a configure type. With freeform configuration profiles, you can store your data in the AWS AppConfig hosted configuration store or any of the following AWS services and Systems Manager tools:

Location	Supported file types
AWS AppConfig hosted configuration store	YAML, JSON, and text if added using the AWS Management Console. Any file type if added using the AWS AppConfig CreateHostedConfigurationVersion API action.
Amazon Simple Storage Service (Amazon S3)	Any
AWS CodePipeline	Pipeline (as defined by the service)
AWS Secrets Manager	Secret (as defined by the service)
AWS Systems Manager Parameter Store	Standard and secure string parameters (as defined by Parameter Store)
AWS Systems Manager document store (SSM documents)	YAML, JSON, text

A configuration profile can also include optional validators to ensure your configuration data is syntactically and semantically correct. AWS AppConfig performs a check using the validators when you start a deployment. If any errors are detected, the deployment stops before making any changes to the targets of the configuration.

Note

If possible, we recommend hosting your configuration data in the AWS AppConfig hosted configuration store as it offers the most features and enhancements.

For freeform configurations stored in the AWS AppConfig hosted configuration store or SSM documents, you can create the freeform configuration by using the Systems Manager console at the time you create a configuration profile. The process is described later in this topic.

For freeform configurations stored in Parameter Store, Secrets Manager, or Amazon S3, you must create the parameter, secret, or object first and store it in the relevant configuration store. After you store the configuration data, use the procedure in this topic to create the configuration profile.

Topics

- [Understanding validators](#)
- [Understanding configuration store quotas and limitations](#)
- [Understanding the AWS AppConfig hosted configuration store](#)
- [Understanding configurations stored in Amazon S3](#)
- [Creating an AWS AppConfig freeform configuration profile \(console\)](#)
- [Creating an AWS AppConfig freeform configuration profile \(command line\)](#)

Understanding validators

When you create a configuration profile, you have the option to specify up to two validators. A validator ensures that your configuration data is syntactically and semantically correct. If you plan to use a validator, you must create it before you create the configuration profile. AWS AppConfig supports the following types of validators:

- **AWS Lambda functions:** Supported for feature flags and free form configurations.
- **JSON Schema:** Supported for free form configurations. (AWS AppConfig automatically validates feature flags against a JSON Schema.)

Topics

- [AWS Lambda function validators](#)
- [JSON Schema validators](#)

AWS Lambda function validators

Lambda function validators must be configured with the following event schema. AWS AppConfig uses this schema to invoke the Lambda function. The content is a base64-encoded string, and the URI is a string.

```
{
  "applicationId": "The application ID of the configuration profile being validated",
  "configurationProfileId": "The ID of the configuration profile being validated",
  "configurationVersion": "The version of the configuration profile being validated",
  "content": "Base64EncodedByteString",
  "uri": "The configuration uri"
}
```

AWS AppConfig verifies that the Lambda X-Amz-Function-Error header is set in the response. Lambda sets this header if the function throws an exception. For more information about X-Amz-Function-Error, see [Error Handling and Automatic Retries in AWS Lambda](#) in the *AWS Lambda Developer Guide*.

Here is a simple example of a Lambda response code for a successful validation.

```
import json

def handler(event, context):
    #Add your validation logic here
    print("We passed!")
```

Here is a simple example of a Lambda response code for an unsuccessful validation.

```
def handler(event, context):
    #Add your validation logic here
    raise Exception("Failure!")
```

Here is another example that validates only if the configuration parameter is a prime number.

```
function isPrime(value) {
  if (value < 2) {
    return false;
  }
}
```

```

    for (i = 2; i < value; i++) {
        if (value % i === 0) {
            return false;
        }
    }

    return true;
}

exports.handler = async function(event, context) {
    console.log('EVENT: ' + JSON.stringify(event, null, 2));
    const input = parseInt(Buffer.from(event.content, 'base64').toString('ascii'));
    const prime = isPrime(input);
    console.log('RESULT: ' + input + (prime ? ' is' : ' is not') + ' prime');
    if (!prime) {
        throw input + "is not prime";
    }
}

```

AWS AppConfig calls your validation Lambda when calling the `StartDeployment` and `ValidateConfigurationActivity` API operations. You must provide `appconfig.amazonaws.com` permissions to invoke your Lambda. For more information, see [Granting Function Access to AWS Services](#). AWS AppConfig limits the validation Lambda run time to 15 seconds, including start-up latency.

JSON Schema validators

If you create a configuration in an SSM document, then you must specify or create a JSON Schema for that configuration. A JSON Schema defines the allowable properties for each application configuration setting. The JSON Schema functions like a set of rules to ensure that new or updated configuration settings conform to the best practices required by your application. Here is an example.

```

{
    "$schema": "http://json-schema.org/draft-04/schema#",
    "title": "$id$",
    "description": "BasicFeatureToggle-1",
    "type": "object",
    "additionalProperties": false,
    "patternProperties": {
        "^[^\\s]+$": {
            "type": "boolean"
        }
    }
}

```

```
    }  
  },  
  "minProperties": 1  
}
```

When you create a configuration from an SSM document, the system automatically verifies that the configuration conforms to the schema requirements. If it doesn't, AWS AppConfig returns a validation error.

Important

Note the following important information about JSON Schema validators:

- Configuration data stored in SSM documents must validate against an associated JSON Schema before you can add the configuration to the system. SSM parameters do not require a validation method, but we recommend that you create a validation check for new or updated SSM parameter configurations by using AWS Lambda.
- A configuration in an SSM document uses the `ApplicationConfiguration` document type. The corresponding JSON Schema, uses the `ApplicationConfigurationSchema` document type.
- AWS AppConfig supports JSON Schema version 4.X for inline schema. If your application configuration requires a different version of JSON Schema, then you must create a Lambda validator.

Understanding configuration store quotas and limitations

Configuration stores supported by AWS AppConfig have the following quotas and limitations.

	AWS AppConfig hosted configuration store	Amazon S3	Systems Manager Parameter Store	AWS Secrets Manager	Systems Manager Document store	AWS CodePipeline
Configuration size limit	2 MB default,	2 MB	4 KB (free tier) / 8 KB (advanced	64 KB	64 KB	2 MB

	AWS AppConfig hosted configuration store	Amazon S3	Systems Manager Parameter Store	AWS Secrets Manager	Systems Manager Document store	AWS CodePipeline
	4 MB maximum	Enforced by AWS AppConfig, not S3	parameter s)			Enforced by AWS AppConfig , not CodePipeline
Resource storage limit	1 GB	Unlimited	10,000 parameter s (free tier) / 100,000 parameter s (advanced parameter s)	500,000	500 documents	Limited by the number of configuration profiles per application (100 profiles per application)
Server-side encryption	Yes	SSE-S3 , SSE-KMS	Yes	Yes	No	Yes
CloudFormation support	Yes	Not for creating or updating data	Yes	Yes	No	Yes

	AWS AppConfig hosted configura tion store	Amazon S3	Systems Manager Parameter Store	AWS Secrets Manager	Systems Manager Document store	AWS CodePipel ine
Pricing	Free	See Amazon S3 pricing	See AWS Systems Manager pricing	See AWS Secrets Manager pricing	Free	See AWS CodePipeline pricing

Understanding the AWS AppConfig hosted configuration store

AWS AppConfig includes an internal or hosted configuration store. Configurations must be 2 MB or smaller. The AWS AppConfig hosted configuration store provides the following benefits over other configuration store options.

- You don't need to set up and configure other services such as Amazon Simple Storage Service (Amazon S3) or Parameter Store.
- You don't need to configure AWS Identity and Access Management (IAM) permissions to use the configuration store.
- You can store configurations in YAML, JSON, or as text documents.
- There is no cost to use the store.
- You can create a configuration and add it to the store when you create a configuration profile.

Understanding configurations stored in Amazon S3

You can store configurations in an Amazon Simple Storage Service (Amazon S3) bucket. When you create the configuration profile, you specify the URI to a single S3 object in a bucket. You also specify the Amazon Resource Name (ARN) of an AWS Identity and Access Management (IAM) role that gives AWS AppConfig permission to get the object. Before you create a configuration profile for an Amazon S3 object, be aware of the following restrictions.

Restriction	Details
Size	Configurations stored as S3 objects can be a maximum of 1 MB in size.
Object encryption	A configuration profile can target SSE-S3 and SSE-KMS encrypted objects.
Storage classes	AWS AppConfig supports the following S3 storage classes: STANDARD, INTELLIGENT_TIERING , REDUCED_REDUNDANCY , STANDARD_IA , and ONEZONE_IA . The following classes are not supported: All S3 Glacier classes (GLACIER and DEEP_ARCHIVE).
Versioning	AWS AppConfig requires that the S3 object use versioning.

Configuring permissions for a configuration stored as an Amazon S3 object

When you create a configuration profile for a configuration stored as an S3 object, you must specify an ARN for an IAM role that gives AWS AppConfig permission to get the object. The role must include the following permissions.

Permissions to access the S3 object

- `s3:GetObject`
- `s3:GetObjectVersion`

Permissions to list S3 buckets

`s3:ListAllMyBuckets`

Permissions to access the S3 bucket where the object is stored

- `s3:GetBucketLocation`
- `s3:GetBucketVersioning`

- s3:ListBucket
- s3:ListBucketVersions

Complete the following procedure to create a role that enables AWS AppConfig to get a configuration stored in an S3 object.

Creating the IAM Policy for Accessing an S3 Object

Use the following procedure to create an IAM policy that enables AWS AppConfig to get a configuration stored in an S3 object.

To create an IAM policy for accessing an S3 object

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Policies**, and then choose **Create policy**.
3. On the **Create policy** page, choose the **JSON** tab.
4. Update the following sample policy with information about your S3 bucket and configuration object. Then paste the policy into the text field on the **JSON** tab. Replace the *placeholder values* with your own information.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:GetObjectVersion"
      ],
      "Resource": "arn:aws:s3:::amzn-s3-demo-bucket/my-configurations/my-configuration.json"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetBucketLocation",
        "s3:GetBucketVersioning",
        "s3:ListBucketVersions",
```

```
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::amzn-s3-demo-bucket"
      ]
    },
    {
      "Effect": "Allow",
      "Action": "s3:ListAllMyBuckets",
      "Resource": "*"
    }
  ]
}
```

5. Choose **Review policy**.
6. On the **Review policy** page, type a name in the **Name** box, and then type a description.
7. Choose **Create policy**. The system returns you to the **Roles** page.

Creating the IAM Role for Accessing an S3 Object

Use the following procedure to create an IAM role that enables AWS AppConfig to get a configuration stored in an S3 object.

To create an IAM role for accessing an Amazon S3 object

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Roles**, and then choose **Create role**.
3. On the **Select type of trusted entity** section, choose **AWS service**.
4. In the **Choose a use case** section, under **Common use cases**, choose **EC2**, and then choose **Next: Permissions**.
5. On the **Attach permissions policy** page, in the search box, enter the name of the policy you created in the previous procedure.
6. Choose the policy and then choose **Next: Tags**.
7. On the **Add tags (optional)** page, enter a key and an optional value, and then choose **Next: Review**.
8. On the **Review** page, type a name in the **Role name** field, and then type a description.
9. Choose **Create role**. The system returns you to the **Roles** page.

10. On the **Roles** page, choose the role you just created to open the **Summary** page. Note the **Role Name** and **Role ARN**. You will specify the role ARN when you create the configuration profile later in this topic.

Creating a Trust Relationship

Use the following procedure to configure the role you just created to trust AWS AppConfig.

To add a trust relationship

1. In the **Summary** page for the role you just created, choose the **Trust Relationships** tab, and then choose **Edit Trust Relationship**.
2. Delete "ec2.amazonaws.com" and add "appconfig.amazonaws.com", as shown in the following example.

JSON

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

3. Choose **Update Trust Policy**.

Creating an AWS AppConfig freeform configuration profile (console)

Use the following procedure to create an AWS AppConfig freeform configuration profile and (optionally) a freeform-configuration by using the AWS Systems Manager console.

To create a freeform configuration profile

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Applications**, and then choose an application you created in [Creating a namespace for your application in AWS AppConfig](#).
3. Choose the **Configuration profiles and feature flags** tab, and then choose **Create configuration**.
4. In the **Configuration options** section, choose **Freeform configuration**.
5. For **Configuration profile name**, enter a name for the configuration profile.
6. (Optional) Expand **Description** and enter a description.
7. (Optional) Expand **Additional options** and complete the following, as necessary.
 - a. In the **Associate extensions** section, choose an extension from the list.
 - b. In the **Tags** section, choose **Add new tag**, and then specify a key and optional value.
8. Choose **Next**.
9. On the **Specify configuration data** page, in the **Configuration definition** section, choose an option.
10. Complete the fields for the option you selected, as described in the following table.

Option selected	Details
AWS AppConfig hosted configuration	Choose either Text , JSON , or YAML , and enter your configuration in the field. Go to Step 12 in this procedure.
Amazon S3 object	Enter the object URI in the S3 object source field and go to Step 11 in this procedure.
AWS CodePipeline	Choose Next and go to Step 12 in this procedure.
Secrets Manager secret	Choose the secret from the list go to Step 11 in this procedure.

Option selected	Details
AWS Systems Manager parameter	Choose the parameter from the list and go to Step 11 in this procedure.
AWS Systems Manager document	1. Choose a document from the list or choose Create new document. 2. If you choose Create new document, for Document name, enter a name. Optionally, expand Version name and enter a name for the document version. 3. For Application configuration schema, either choose the JSON schema from the list or choose Create schema. If you choose Create schema, Systems Manager opens the Create schema page. Enter the schema details, and then choose Create application configuration schema. 4. In the Content section, choose either YAML or JSON and then enter the configuration data in the field.

11. In the **Service role** section, choose **New service role** to have AWS AppConfig create the IAM role that provides access to the configuration data. AWS AppConfig automatically populates the **Role name** field based on the name you entered earlier. Or, choose **Existing service role**. Choose the role by using the **Role ARN** list.
12. Optionally, on the **Add validators** page, choose either **JSON Schema** or **AWS Lambda**. If you choose **JSON Schema**, enter the JSON Schema in the field. If you choose **AWS Lambda**, choose the function Amazon Resource Name (ARN) and the version from the list.

Important

Configuration data stored in SSM documents must validate against an associated JSON Schema before you can add the configuration to the system. SSM parameters do not require a validation method, but we recommend that you create a validation check for new or updated SSM parameter configurations by using AWS Lambda.

13. Choose **Next**.
14. On the **Review and save** page, choose **Save and continue to deploy**.

Important

If you created a configuration profile for AWS CodePipeline, then you must create a pipeline in CodePipeline that specifies AWS AppConfig as the *deploy provider*. You don't need to perform [Deploying feature flags and configuration data in AWS AppConfig](#). However, you must configure a client to receive application configuration updates as described in [Retrieving configuration data without AWS AppConfig Agent](#). For information about creating a pipeline that specifies AWS AppConfig as the deploy provider, see [Tutorial: Create a Pipeline that Uses AWS AppConfig as a Deployment Provider](#) in the *AWS CodePipeline User Guide*.

Proceed to [Deploying feature flags and configuration data in AWS AppConfig](#).

Creating an AWS AppConfig freeform configuration profile (command line)

The following procedure describes how to use the AWS CLI (on Linux or Windows) or AWS Tools for PowerShell to create an AWS AppConfig freeform configuration profile. If you prefer, you can use AWS CloudShell to run the commands listed below. For more information, see [What is AWS CloudShell?](#) in the *AWS CloudShell User Guide*.

Note

For freeform configurations hosted in the AWS AppConfig hosted configuration store, you specify hosted for the location URI.

To create a configuration profile by using the AWS CLI

1. Open the AWS CLI.
2. Run the following command to create a freeform configuration profile.

Linux

```
aws appconfig create-configuration-profile \
```

```
--application-id APPLICATION_ID \
--name NAME \
--description CONFIGURATION_PROFILE_DESCRIPTION \
--location-uri CONFIGURATION_URI or hosted \
--retrieval-role-arn IAM_ROLE_ARN \
--tags TAGS \
--validators "Content=SCHEMA_CONTENT or LAMBDA_FUNCTION_ARN,Type=JSON_SCHEMA
or LAMBDA"
```

Windows

```
aws appconfig create-configuration-profile ^
--application-id APPLICATION_ID ^
--name NAME ^
--description CONFIGURATION_PROFILE_DESCRIPTION ^
--location-uri CONFIGURATION_URI or hosted ^
--retrieval-role-arn IAM_ROLE_ARN ^
--tags TAGS ^
--validators "Content=SCHEMA_CONTENT or LAMBDA_FUNCTION_ARN,Type=JSON_SCHEMA
or LAMBDA"
```

PowerShell

```
New-APPCConfigurationProfile `
-Name NAME `
-ApplicationId APPLICATION_ID `
-Description CONFIGURATION_PROFILE_DESCRIPTION `
-LocationUri CONFIGURATION_URI or hosted `
-RetrievalRoleArn IAM_ROLE_ARN `
-Tag TAGS `
-Validators "Content=SCHEMA_CONTENT or LAMBDA_FUNCTION_ARN,Type=JSON_SCHEMA
or LAMBDA"
```

Important

Note the following important information.

- If you created a configuration profile for AWS CodePipeline, then you must create a pipeline in CodePipeline that specifies AWS AppConfig as the *deploy provider*. You don't need to perform [Deploying feature flags and configuration data in AWS AppConfig](#).

However, you must configure a client to receive application configuration updates as described in [Retrieving configuration data without AWS AppConfig Agent](#). For information about creating a pipeline that specifies AWS AppConfig as the deploy provider, see [Tutorial: Create a Pipeline that Uses AWS AppConfig as a Deployment Provider](#) in the *AWS CodePipeline User Guide*.

- If you created a configuration in the AWS AppConfig hosted configuration store, you can create new versions of the configuration by using the [CreateHostedConfigurationVersion](#) API operations. To view AWS CLI details and sample commands for this API operation, see [create-hosted-configuration-version](#) in the *AWS CLI Command Reference*.

Proceed to [Deploying feature flags and configuration data in AWS AppConfig](#).

Creating a configuration profile for non-native data sources

AWS AppConfig supports deploying configuration data from most any data store. Natively, AWS AppConfig supports deploying configuration data stored in the following services:

- The AWS AppConfig hosted configuration store
- Amazon S3
- AWS Secrets Manager
- AWS Systems Manager Parameter Store
- Systems Manager Document Store
- AWS CodePipeline

If your configuration data is stored in a location not natively supported by AWS AppConfig, you can create an [AWS AppConfig extension](#) to retrieve your data from its source. For example, by using an AWS AppConfig extension, you can retrieve configuration data stored in Amazon Relational Database Service (Amazon RDS), Amazon DynamoDB (DynamoDB), GitHub, GitLab, or a local repo, to name a few. By implementing an extension, you can take advantage of AWS AppConfig security and DevOps enhancements for your applications and computing environment. You can also use this method as you migrate configuration data from legacy systems into AWS AppConfig.

Creating a configuration profile for data sources not natively supported in AWS AppConfig involves the following processes or actions:

1. Create an [AWS Lambda function](#) that fetches data from your data source. As long as a Lambda function can access the data source, your AWS AppConfig extension will be able to retrieve the data.
2. Create a custom AWS AppConfig extension that invokes your Lambda function. For more information, see [Walkthrough: Creating custom AWS AppConfig extensions](#).
3. Create an AWS AppConfig free-form configuration profile. Specifically, create a configuration profile that uses the **AWS AppConfig hosted configuration** definition. The configuration profile functions as a temporary data store after your Lambda function retrieves your configuration from your source. Your application will retrieve the configuration data from the AWS AppConfig hosted configuration store. For more information, see [Creating a free form configuration profile in AWS AppConfig](#).
4. Create an extension association that triggers using the `PRE_CREATE_HOSTED_CONFIGURATION_VERSION` action point. For more information, see [Step 4: Create an extension association for a custom AWS AppConfig extension](#).

Once configured, when your application requests a new version of the configuration data, the Lambda fetches your configuration data and pulls it into the configuration profile. AWS AppConfig then saves the configuration profile and your third-party data.

When you're ready, you can deploy the configuration profile to your applications, just like any other type of configuration data.

Note

You can choose to insert third-party data in line with existing configuration data, or have the entire contents of the configuration data contain only the third-party data. If you want to have the data in line with other existing data, that logic should be part of the Lambda function that imports the data from the third-party source.

Migrating to AWS AppConfig from legacy and home-grown configuration services

If you've started using AWS AppConfig and still have legacy configuration data or feature flags in another system, you can use the process described earlier in this topic to migrate off of your legacy system and onto AWS AppConfig. You can build an extension that pulls data out of your legacy system and deploys it through AWS AppConfig. Using AWS AppConfig in this way provides you with all of the safety guardrail controls and benefits while still using your legacy data stores.

Deploying feature flags and configuration data in AWS AppConfig

After you [create required artifacts](#) for working with feature flags and freeform configuration data, you can create a new deployment. When you create a new deployment, you specify the following information:

- An application ID
- A configuration profile ID
- A configuration version
- An environment ID where you want to deploy the configuration data
- A deployment strategy ID that defines how fast you want the changes to take effect
- An AWS Key Management Service (AWS KMS) key ID to encrypt the data using a customer managed key.

When you call the [StartDeployment](#) API action, AWS AppConfig performs the following tasks:

1. Retrieves the configuration data from the underlying data store by using the location URI in the configuration profile.
2. Verifies the configuration data is syntactically and semantically correct by using the validators you specified when you created your configuration profile.
3. Caches a copy of the data so it is ready to be retrieved by your application. This cached copy is called the *deployed data*.

You can mitigate situations where deploying configuration data causes errors in your application by using a combination of AWS AppConfig deployment strategies and automatic rollbacks based on Amazon CloudWatch alarms. A deployment strategy enables you to slowly release changes to all targets or specific segments over minutes or hours—either session-based or along your own target dimension by leveraging entity-based deployments. After you configure CloudWatch, if one or more alarms go into the alarm state during a deployment, AWS AppConfig automatically rolls back your configuration data to the previous version. For more information about deployment strategies, see [Working with deployment strategies](#). For more information about automatic rollbacks, see [Monitoring deployments for automatic rollback](#).

Note

AWS AppConfig Agent (version 2.0.136060 or later) supports deploying feature flag or free-form configuration data to specific segments or individual users during a gradual rollout. Entity-based gradual deployments ensure that once a user or segment receives a configuration version, they continue to receive that same version throughout the deployment period, regardless of which compute resource serves their requests. For more information, see [Using AWS AppConfig Agent for user- or entity-based gradual deployments](#).

Topics

- [Working with deployment strategies](#)
- [Deploying a configuration](#)
- [Deploying AWS AppConfig configurations using CodePipeline](#)
- [Reverting a configuration](#)

Working with deployment strategies

A deployment strategy enables you to slowly release changes to all targets or specific segments over minutes or hours—either session based or along your own target dimension by leveraging entity-based deployments.

Note

AWS AppConfig Agent (version 2.0.136060 or later) supports deploying feature flag or free-form configuration data to specific segments or individual users during a gradual rollout. Entity-based gradual deployments ensure that once a user or segment receives a configuration version, they continue to receive that same version throughout the deployment period, regardless of which compute resource serves their requests. For more information, see [Using AWS AppConfig Agent for user- or entity-based gradual deployments](#).

An AWS AppConfig deployment strategy defines the following important aspects of a configuration deployment.

Setting	Description														
Deployment type	Deployment type defines how the configuration deploys or <i>rolls out</i>. AWS AppConfig supports Linear and Exponential deployment types. • Linear: For this type, AWS AppConfig processes the deployment by increments of the growth factor evenly distributed over the deployment. Here's an example timeline for a 10 hour deployment that uses 20% linear growth: <table><tr><th>Elapsed time</th><th>Deployment progress</th></tr><tr><td>0 hour</td><td>0%</td></tr><tr><td>2 hour</td><td>20%</td></tr><tr><td>4 hour</td><td>40%</td></tr><tr><td>6 hour</td><td>60%</td></tr><tr><td>8 hour</td><td>80%</td></tr><tr><td>10 hour</td><td>100%</td></tr></table> • Exponential: For this type, AWS AppConfig processes the deployment exponentially using the following formula: $G * (2^N)$. In this formula, G is the step percentage specified by the user and N is the number of steps until the configuration is deployed to all targets. For example, if you specify a growth factor of 2, then the system rolls out the configuration as follows:	Elapsed time	Deployment progress	0 hour	0%	2 hour	20%	4 hour	40%	6 hour	60%	8 hour	80%	10 hour	100%
Elapsed time	Deployment progress														
0 hour	0%														
2 hour	20%														
4 hour	40%														
6 hour	60%														
8 hour	80%														
10 hour	100%														

Setting	Description
	$2*(2^0)$ $2*(2^1)$ $2*(2^2)$ Expressed numerically, the deployment rolls out as follows: 2% of the targets, 4% of the targets, 8% of the targets, and continues until the configuration has been deployed to all targets.
Step percentage (growth factor)	This setting specifies the percentage of callers to target during each step of the deployment. Note In the SDK and the AWS AppConfig API Reference, step percentage is called growth factor.
Deployment time	This setting specifies an amount of time during which AWS AppConfig deploys to hosts. This is not a timeout value. It is a window of time during which the deployment is processed in intervals.

Setting	Description
Bake time	This setting specifies the amount of time AWS AppConfig monitors for Amazon CloudWatch alarms after the configuration has been deployed to 100% of its targets, before considering the deployment to be complete. If an alarm is triggered during this time, AWS AppConfig rolls back the deployment. You must configure permissions for AWS AppConfig to roll back based on CloudWatch alarms. For more information, see Configure permissions for automatic rollback .

You can choose a predefined strategy included with AWS AppConfig or create your own.

 Note

AWS AppConfig Agent (version 2.0.136060 or later) supports deploying feature flag or free-form configuration data to specific segments or individual users during a gradual rollout. Entity-based gradual deployments ensure that once a user or segment receives a configuration version, they continue to receive that same version throughout the deployment period, regardless of which compute resource serves their requests. For more information, see [Using AWS AppConfig Agent for user- or entity-based gradual deployments](#).

Topics

- [Using predefined deployment strategies](#)
- [Create a deployment strategy](#)

Using predefined deployment strategies

AWS AppConfig includes predefined deployment strategies to help you quickly deploy a configuration. Instead of creating your own strategies, you can choose one of the following when you deploy a configuration.

Deployment strategy	Description
AppConfig.Linear20PercentEvery6Minutes	AWS recommended: This strategy deploys the configuration to 20% of all targets every six minutes for a 30 minute deployment. The system monitors for Amazon CloudWatch alarms for 30 minutes. If no alarms are received in this time, the deployment is complete. If an alarm is triggered during this time, AWS AppConfig rolls back the deployment. We recommend using this strategy for production deployments because it aligns with AWS best practices and includes additional emphasis on deployment safety due to its long duration and bake time.
AppConfig.Canary10Percent20Minutes	AWS recommended: This strategy processes the deployment exponentially using a 10% growth factor over 20 minutes. The system monitors for CloudWatch alarms for 10 minutes. If no alarms are received in this time, the deployment is complete. If an alarm is triggered during this time, AWS AppConfig rolls back the deployment. We recommend using this strategy for production deployments because it aligns with AWS best practices for configuration deployments.
AppConfig.AllAtOnce	Quick:

Deployment strategy	Description
	This strategy deploys the configuration to all targets immediately. The system monitors for CloudWatch alarms for 10 minutes. If no alarms are received in this time, the deployment is complete. If an alarm is triggered during this time, AWS AppConfig rolls back the deployment.
AppConfig.Linear50PercentEvery30Seconds	Testing/demonstration: This strategy deploys the configuration to half of all targets every 30 seconds for a one-minute deployment. The system monitors for Amazon CloudWatch alarms for 1 minute. If no alarms are received in this time, the deployment is complete. If an alarm is triggered during this time, AWS AppConfig rolls back the deployment. We recommend using this strategy only for testing or demonstration purposes because it has a short duration and bake time.

Create a deployment strategy

If you don't want to use one of the predefined deployment strategies, you can create your own. You can create a maximum of 20 deployment strategies. When you deploy a configuration, you can choose the deployment strategy that works best for the application and the environment.

Creating an AWS AppConfig deployment strategy (console)

Use the following procedure to create an AWS AppConfig deployment strategy by using the AWS Systems Manager console.

To create a deployment strategy

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Deployment strategies**, and then choose **Create deployment strategy**.
3. For **Name**, enter a name for the deployment strategy.
4. For **Description**, enter information about the deployment strategy.
5. For **Deployment type**, choose a type.
6. For **Step percentage**, choose the percentage of callers to target during each step of the deployment.
7. For **Deployment time**, enter the total duration for the deployment in minutes or hours.
8. For **Bake time**, enter the total time, in minutes or hours, to monitor for Amazon CloudWatch alarms before proceeding to the next step of a deployment or before considering the deployment to be complete.
9. In the **Tags** section, enter a key and an optional value. You can specify a maximum of 50 tags for a resource.
10. Choose **Create deployment strategy**.

Important

If you created a configuration profile for AWS CodePipeline, then you must create a pipeline in CodePipeline that specifies AWS AppConfig as the *deploy provider*. You don't need to perform [Deploying a configuration](#). However, you must configure a client to receive application configuration updates as described in [Retrieving configuration data without AWS AppConfig Agent](#). For information about creating a pipeline that specifies AWS AppConfig as the deploy provider, see [Tutorial: Create a Pipeline that Uses AWS AppConfig as a Deployment Provider](#) in the *AWS CodePipeline User Guide*.

Proceed to [Deploying a configuration](#).

Creating an AWS AppConfig deployment strategy (command line)

The following procedure describes how to use the AWS CLI (on Linux or Windows) or AWS Tools for PowerShell to create an AWS AppConfig deployment strategy.

To create a deployment strategy step by step

1. Open the AWS CLI.
2. Run the following command to create a deployment strategy.

Linux

```
aws appconfig create-deployment-strategy \
  --name A_name_for_the_deployment_strategy \
  --description A_description_of_the_deployment_strategy \
  --deployment-duration-in-minutes Total_amount_of_time_for_a_deployment_to_last \
  --final-bake-time-in-minutes Amount_of_time_AWS
AppConfig_monitors_for_alarms_before_considering_the_deployment_to_be_complete \
  --growth-
factor The_percentage_of_targets_to_receive_a_deployed_configuration_during_each_interval \
  --growth-
type The_linear_or_exponential_algorithm_used_to_define_how_percentage_grows_over_time \
  --replicate-
to To_save_the_deployment_strategy_to_a_Systems_Manager(SSM)_document \
  --tags User_defined_key_value_pair_metadata_of_the_deployment_strategy
```

Windows

```
aws appconfig create-deployment-strategy ^
  --name A_name_for_the_deployment_strategy ^
  --description A_description_of_the_deployment_strategy ^
  --deployment-duration-in-minutes Total_amount_of_time_for_a_deployment_to_last ^
  --final-bake-time-in-minutes Amount_of_time_AWS
AppConfig_monitors_for_alarms_before_considering_the_deployment_to_be_complete ^
  --growth-
factor The_percentage_of_targets_to_receive_a_deployed_configuration_during_each_interval ^
  --growth-
type The_linear_or_exponential_algorithm_used_to_define_how_percentage_grows_over_time ^
  --name A_name_for_the_deployment_strategy ^
```

```
--replicate-
to To_save_the_deployment_strategy_to_a_Systems_Manager_(SSM)_document ^
--tags User_defined_key_value_pair_metadata_of_the_deployment_strategy
```

PowerShell

```
New-APPCDeploymentStrategy `
--Name A_name_for_the_deployment_strategy `
--Description A_description_of_the_deployment_strategy `
--DeploymentDurationInMinutes Total_amount_of_time_for_a_deployment_to_last `
--FinalBakeTimeInMinutes Amount_of_time_AWS
AppConfig_monitors_for_alarms_before_considering_the_deployment_to_be_complete
`
--
GrowthFactor The_percentage_of_targets_to_receive_a_deployed_configuration_during_each_i
`
--
GrowthType The_linear_or_exponential_algorithm_used_to_define_how_percentage_grows_over
`
--
ReplicateTo To_save_the_deployment_strategy_to_a_Systems_Manager_(SSM)_document
`
--
Tag Hashtable_type_User_defined_key_value_pair_metadata_of_the_deployment_strategy
```

The system returns information like the following.

Linux

```
{
  "Id": "Id of the deployment strategy",
  "Name": "Name of the deployment strategy",
  "Description": "Description of the deployment strategy",
  "DeploymentDurationInMinutes": "Total amount of time the deployment lasted",
  "GrowthType": "The linear or exponential algorithm used to define how
percentage grew over time",
  "GrowthFactor": "The percentage of targets that received a deployed
configuration during each interval",
  "FinalBakeTimeInMinutes": "The amount of time AWS AppConfig monitored for
alarms before considering the deployment to be complete",
  "ReplicateTo": "The Systems Manager (SSM) document where the deployment
strategy is saved"
```

```
}
```

Windows

```
{
  "Id": "Id of the deployment strategy",
  "Name": "Name of the deployment strategy",
  "Description": "Description of the deployment strategy",
  "DeploymentDurationInMinutes": "Total amount of time the deployment lasted",
  "GrowthType": "The linear or exponential algorithm used to define how
percentage grew over time",
  "GrowthFactor": "The percentage of targets that received a deployed
configuration during each interval",
  "FinalBakeTimeInMinutes": "The amount of time AWS AppConfig monitored for
alarms before considering the deployment to be complete",
  "ReplicateTo": "The Systems Manager (SSM) document where the deployment
strategy is saved"
}
```

PowerShell

```
ContentLength           : Runtime of the command
DeploymentDurationInMinutes : Total amount of time the deployment lasted
Description              : Description of the deployment strategy
FinalBakeTimeInMinutes   : The amount of time AWS AppConfig monitored for
alarms before considering the deployment to be complete
GrowthFactor             : The percentage of targets that received a deployed
configuration during each interval
GrowthType               : The linear or exponential algorithm used to define
how percentage grew over time
HttpStatusCode           : HTTP Status of the runtime
Id                       : The deployment strategy ID
Name                     : Name of the deployment strategy
ReplicateTo              : The Systems Manager (SSM) document where the
deployment strategy is saved
ResponseMetadata         : Runtime Metadata
```

Deploying a configuration

After you [create required artifacts](#) for working with feature flags and freeform configuration data, you can create a new deployment by using the AWS Management Console, the AWS CLI, or the

SDK. Starting a deployment in AWS AppConfig calls the [StartDeployment](#) API operation. This call includes the IDs of the AWS AppConfig application, the environment, the configuration profile, and (optionally) the configuration data version to deploy. The call also includes the ID of the deployment strategy to use, which determines how the configuration data is deployed.

If you deploy secrets stored in AWS Secrets Manager, Amazon Simple Storage Service (Amazon S3) objects encrypted with a customer managed key, or secure string parameters stored in AWS Systems Manager Parameter Store encrypted with a customer managed key, you must specify a value for the `KmsKeyIdIdentifier` parameter. If your configuration is not encrypted or is encrypted with an AWS managed key, specifying a value for the `KmsKeyIdIdentifier` parameter is not required.

Note

The value you specify for `KmsKeyIdIdentifier` must be a customer managed key. This doesn't have to be the same key you used to encrypt your configuration.

When you start a deployment with a `KmsKeyIdIdentifier`, the permission policy attached to your AWS Identity and Access Management (IAM) principal must allow the `kms:GenerateDataKey` operation.

AWS AppConfig monitors the distribution to all hosts and reports status. If a distribution fails, then AWS AppConfig rolls back the configuration.

Note

You can only deploy one configuration at a time to an environment. However, you can deploy one configuration each to different environments at the same time.

Deploy a configuration (console)

Use the following procedure to deploy an AWS AppConfig configuration by using the AWS Systems Manager console.

To deploy a configuration by using the console

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.

2. In the navigation pane, choose **Applications**, and then choose an application you created in [Creating a namespace for your application in AWS AppConfig](#).
3. On the **Environments** tab, fill the radio button for an environment, and then choose **View details**.
4. Choose **Start deployment**.
5. For **Configuration**, choose a configuration from the list.
6. Depending on the source of your configuration, use the version list to choose the version you want to deploy.
7. For **Deployment strategy**, choose a strategy from the list.
8. (Optional) For **Deployment description**, enter a description.
9. For **Additional encryption options**, choose a AWS Key Management Service key from the list.
10. (Optional) In the **Tags** section, choose **Add new tag** and enter a key and an optional value. You can specify a maximum of 50 tags for a resource.
11. Choose **Start deployment**.

Deploy a configuration (commandline)

The following procedure describes how to use the AWS CLI (on Linux or Windows) or AWS Tools for PowerShell to deploy an AWS AppConfig configuration.

To deploy a configuration step by step

1. Open the AWS CLI.
2. Run the following command to deploy a configuration.

Linux

```
aws appconfig start-deployment \  
  --application-id The_application_ID \  
  --environment-id The_environment_ID \  
  --deployment-strategy-id The_deployment_strategy_ID \  
  --configuration-profile-id The_configuration_profile_ID \  
  --configuration-version The_configuration_version_to_deploy \  
  --description A_description_of_the_deployment \  
  --tags User_defined_key_value_pair_metadata_of_the_deployment
```

Windows

```
aws appconfig start-deployment ^  
  --application-id The_application_ID ^  
  --environment-id The_environment_ID ^  
  --deployment-strategy-id The_deployment_strategy_ID ^  
  --configuration-profile-id The_configuration_profile_ID ^  
  --configuration-version The_configuration_version_to_deploy ^  
  --description A_description_of_the_deployment ^  
  --tags User_defined_key_value_pair_metadata_of_the_deployment
```

PowerShell

```
Start-APPCDeployment `   
  -ApplicationId The_application_ID `   
  -ConfigurationProfileId The_configuration_profile_ID `   
  -ConfigurationVersion The_configuration_version_to_deploy `   
  -DeploymentStrategyId The_deployment_strategy_ID `   
  -Description A_description_of_the_deployment `   
  -EnvironmentId The_environment_ID `   
  -Tag Hashtable_type_user_defined_key_value_pair_metadata_of_the_deployment
```

The system returns information like the following.

Linux

```
{  
  "ApplicationId": "The ID of the application that was deployed",  
  "EnvironmentId" : "The ID of the environment",  
  "DeploymentStrategyId": "The ID of the deployment strategy that was  
deployed",  
  "ConfigurationProfileId": "The ID of the configuration profile that was  
deployed",  
  "DeploymentNumber": The sequence number of the deployment,  
  "ConfigurationName": "The name of the configuration",  
  "ConfigurationLocationUri": "Information about the source location of the  
configuration",  
  "ConfigurationVersion": "The configuration version that was deployed",  
  "Description": "The description of the deployment",  
  "DeploymentDurationInMinutes": Total amount of time the deployment lasted,
```

```

    "GrowthType": "The linear or exponential algorithm used to define how
percentage grew over time",
    "GrowthFactor": "The percentage of targets to receive a deployed configuration
during each interval",
    "FinalBakeTimeInMinutes": "Time AWS AppConfig monitored for alarms before
considering the deployment to be complete",
    "State": "The state of the deployment",

    "EventLog": [
        {
            "Description": "A description of the deployment event",
            "EventType": "The type of deployment event",
            "OccurredAt": "The date and time the event occurred",
            "TriggeredBy": "The entity that triggered the deployment event"
        }
    ],

    "PercentageComplete": "The percentage of targets for which the deployment is
available",
    "StartedAt": "The time the deployment started",
    "CompletedAt": "The time the deployment completed"
}

```

Windows

```

{
    "ApplicationId": "The ID of the application that was deployed",
    "EnvironmentId": "The ID of the environment",
    "DeploymentStrategyId": "The ID of the deployment strategy that was
deployed",
    "ConfigurationProfileId": "The ID of the configuration profile that was
deployed",
    "DeploymentNumber": "The sequence number of the deployment",
    "ConfigurationName": "The name of the configuration",
    "ConfigurationLocationUri": "Information about the source location of the
configuration",
    "ConfigurationVersion": "The configuration version that was deployed",
    "Description": "The description of the deployment",
    "DeploymentDurationInMinutes": "Total amount of time the deployment lasted",
    "GrowthType": "The linear or exponential algorithm used to define how
percentage grew over time",
    "GrowthFactor": "The percentage of targets to receive a deployed configuration
during each interval",
}

```

```

    "FinalBakeTimeInMinutes": Time AWS AppConfig monitored for alarms before
    considering the deployment to be complete,
    "State": "The state of the deployment",

    "EventLog": [
        {
            "Description": "A description of the deployment event",
            "EventType": "The type of deployment event",
            "OccurredAt": The date and time the event occurred,
            "TriggeredBy": "The entity that triggered the deployment event"
        }
    ],

    "PercentageComplete": The percentage of targets for which the deployment is
    available,
    "StartedAt": The time the deployment started,
    "CompletedAt": The time the deployment completed
}

```

PowerShell

```

ApplicationId           : The ID of the application that was deployed
CompletedAt             : The time the deployment completed
ConfigurationLocationUri : Information about the source location of the
    configuration
ConfigurationName       : The name of the configuration
ConfigurationProfileId   : The ID of the configuration profile that was
    deployed
ConfigurationVersion     : The configuration version that was deployed
ContentLength           : Runtime of the deployment
DeploymentDurationInMinutes : Total amount of time the deployment lasted
DeploymentNumber         : The sequence number of the deployment
DeploymentStrategyId     : The ID of the deployment strategy that was
    deployed
Description             : The description of the deployment
EnvironmentId           : The ID of the environment that was deployed
EventLog                : {Description : A description of the deployment
    event, EventType : The type of deployment event, OccurredAt : The date and time
    the event occurred,
        TriggeredBy : The entity that triggered the deployment event}
FinalBakeTimeInMinutes   : Time AWS AppConfig monitored for alarms before
    considering the deployment to be complete

```

GrowthFactor	: The percentage of targets to receive a deployed configuration during each interval
GrowthType	: The linear or exponential algorithm used to define how percentage grew over time
HttpStatusCode	: HTTP Status of the runtime
PercentageComplete	: The percentage of targets for which the deployment is available
ResponseMetadata	: Runtime Metadata
StartedAt	: The time the deployment started
State	: The state of the deployment

Deploying AWS AppConfig configurations using CodePipeline

AWS AppConfig is an integrated deploy action for AWS CodePipeline (CodePipeline). CodePipeline is a fully managed continuous delivery service that helps you automate your release pipelines for fast and reliable application and infrastructure updates. CodePipeline automates the build, test, and deploy phases of your release process every time there is a code change, based on the release model you define. For more information, see [What is AWS CodePipeline?](#)

The integration of AWS AppConfig with CodePipeline offers the following benefits:

- Customers who use CodePipeline to manage orchestration now have a lightweight means of deploying configuration changes to their applications without having to deploy their entire code base.
- Customers who want to use AWS AppConfig to manage configuration deployments but are limited because AWS AppConfig does not support their current code or configuration store, now have additional options. CodePipeline supports AWS CodeCommit, GitHub, and BitBucket (to name a few).

Note

AWS AppConfig integration with CodePipeline is only supported in AWS Regions where CodePipeline is [available](#).

How integration works

You start by setting up and configuring CodePipeline. This includes adding your configuration to a CodePipeline-supported code store. Next, you set up your AWS AppConfig environment by performing the following tasks:

- [Create a namespace and a configuration profile](#)
- [Choose a predefined deployment strategy or create your own](#)

After you complete these tasks, you create a pipeline in CodePipeline that specifies AWS AppConfig as the *deploy provider*. You can then make a change to your configuration and upload it to your CodePipeline code store. Uploading the new configuration automatically starts a new deployment in CodePipeline. After the deployment completes, you can verify your changes. For information about creating a pipeline that specifies AWS AppConfig as the deploy provider, see [Tutorial: Create a Pipeline That Uses AWS AppConfig as a Deployment Provider](#) in the *AWS CodePipeline User Guide*.

Reverting a configuration

During a deployment, you can mitigate situations where malformed or incorrect configuration data causes errors in your application by using automatic rollbacks (if an alarm triggers during a deployment) or by reverting the configuration data to the previous version (if a deployment successfully completed).

For automatic rollbacks, you can use a combination of AWS AppConfig [deployment strategies](#) and Amazon CloudWatch alarms. Once configured, if one or more CloudWatch alarms go into the ALARM state during a deployment, AWS AppConfig automatically rolls back your configuration data to the previous version, thereby preventing application outages or errors. To get started, see [Configure permissions for automatic rollback](#).

Note

You can also roll back a configuration by calling the [StopDeployment](#) API operation while a deployment is still in progress.

For deployments that successfully complete, AWS AppConfig also supports reverting configuration data to a previous version by using the AllowRevert parameter with the [StopDeployment](#) API

operation. For some customers, reverting to a previous configuration after a successful deployment guarantees the data will be the same as it was before the deployment. Reverting also ignores alarm monitors, which may prevent a roll forward from progressing during an application emergency.

 Important

If you call `StopDeployment` with the `AllowRevert` parameter enabled, AWS AppConfig will revert the deployment only if the deployment succeeded within the last 72 hours. After 72 hours, the deployment can no longer be reverted. You must create a new deployment.

Here's a breakdown of the `StopDeployment` functionality based on different situations.

1. If `StopDeployment` is called on an in-progress deployment, the resulting deployment state will be `ROLLED_BACK`.
2. If `StopDeployment` (with `AllowRevert`) is called on an in-progress deployment, the resulting deployment state will be `ROLLED_BACK`.
3. If `StopDeployment` is called on a completed deployment, a `BadRequestException` will be thrown.
4. If `StopDeployment` (with `AllowRevert`) is called on a completed deployment, the resulting deployment state will be `REVERTED`.
5. If `StopDeployment` (with `AllowRevert`) is called on a completed deployment after 72 hours, a `BadRequestException` will be thrown.

You can use the AWS CLI to call the [StopDeployment](#) operation with the `AllowRevert` parameter. Here's an example AWS CLI command that includes the `AllowRevert` parameter.

```
aws appconfig stop-deployment \  
  --application-id 339ohji \  
  --environment-id 54j1r29 \  
  --deployment-number 2 \  
  --allow-revert
```

Retrieving feature flags and configuration data in AWS AppConfig

Your application retrieves feature flags and free form configuration data by establishing a configuration session using the AWS AppConfig Data service. We recommend you use AWS AppConfig Agent to retrieve configuration data. The agent (or the AWS AppConfig Agent Lambda extension for Lambda compute environments) manages a series of API calls and session tokens on your behalf. From a high level, the process works as follows:

1. You configure AWS AppConfig Agent as a local host and have the agent poll AWS AppConfig for configuration updates.
2. The agent calls the [StartConfigurationSession](#) and [GetLatestConfiguration](#) API actions and caches your configuration data locally.
3. To retrieve the data, your application makes an HTTP call to the localhost server. AWS AppConfig Agent supports several use cases, as described in [How to use AWS AppConfig Agent to retrieve configuration data](#).

If you prefer, you can manually call these API actions to retrieve a configuration. The API process works as follows:

1. Your application establishes a configuration session using the `StartConfigurationSession` API action. Your session's client then makes periodic calls to `GetLatestConfiguration` to check for and retrieve the latest data available.
2. When calling `StartConfigurationSession`, your code sends identifiers (ID or name) of an AWS AppConfig application, environment, and configuration profile that the session tracks.
3. In response, AWS AppConfig provides an `InitialConfigurationToken` to be given to the session's client and used the first time it calls `GetLatestConfiguration` for that session.
4. When calling `GetLatestConfiguration`, your client code sends the most recent `ConfigurationToken` value it has and receives in response:
 - `NextPollConfigurationToken`: the `ConfigurationToken` value to use on the next call to `GetLatestConfiguration`.
 - The configuration: the latest data intended for the session. This may be empty if the client already has the latest version of the configuration.

Note

Retrieving configuration data from a separate AWS account isn't supported.

Contents

- [What is AWS AppConfig Agent?](#)
- [How to use AWS AppConfig Agent to retrieve configuration data](#)

What is AWS AppConfig Agent?


AWS AppConfig Agent is an Amazon-developed and managed process for retrieving configuration data from AWS AppConfig. With the agent, you can cache configuration data locally and asynchronously poll the AWS AppConfig data plane service for updates. This caching/polling process ensures that your configuration data is always available for your application while minimizing latency and cost. The agent is not the only way to retrieve configuration data from AWS AppConfig, but it is the recommended way. The agent enhances application processing and management in the following ways:

- The agent calls AWS AppConfig on your behalf by using an AWS Identity and Access Management (IAM) principal and managing a local cache of configuration data. By retrieving configuration data from the local cache, your application requires fewer code updates to manage configuration data, retrieves configuration data in milliseconds, and isn't affected by network issues that can disrupt calls for such data.
- The agent offers a native experience for retrieving and resolving AWS AppConfig feature flags.
- Out of the box, the agent provides best practices for caching strategies, polling intervals, and availability of local configuration data while tracking the configuration tokens needed for subsequent service calls.
- While running in the background, the agent periodically polls the AWS AppConfig data plane service for configuration data updates. Your application can retrieve the data by connecting to localhost on port 2772 (a customizable default port value) and calling HTTP GET to retrieve the data.

Note

AWS AppConfig Agent caches data the first time the service retrieves your configuration data. For this reason, the first call to retrieve data is slower than subsequent calls.

The following diagram shows how AWS AppConfig Agent works.



1. Your application requests configuration data from the agent.
2. The agent returns data from an in-memory cache.
3. The agent asynchronously polls the AWS AppConfig service for the latest configuration data on a predefined cadence. The latest configuration data is always stored in a cache in memory.

How to use AWS AppConfig Agent to retrieve configuration data

The AWS AppConfig Agent is the recommended method for retrieving AWS AppConfig feature flags or free form configuration data. The agent is supported on all forms of AWS Compute including Amazon EC2, Amazon ECS, Amazon EKS, and Lambda. After you complete the initial agent set up, using the agent to retrieve configuration data is simpler than directly calling AWS AppConfig APIs. The agent automatically implements best practices and may lower your cost of using AWS AppConfig as a result of fewer API calls to retrieve configurations.

Note

Retrieving configuration data from a separate AWS account isn't supported.

Using AWS AppConfig Agent for user- or entity-based gradual deployments

AWS AppConfig Agent supports deploying feature flag or free-form configuration data to specific segments or individual users during a gradual rollout. Entity-based gradual deployments ensure that once a user or segment receives a configuration version, they continue to receive that same version throughout the deployment period, regardless of which compute resource serves their requests.

With entity-based gradual deployments, AWS AppConfig Agent evaluates a unique identifier (Entity-Id) supplied with each HTTP request. Based on this identifier, the agent consistently serves either the new or previous configuration version throughout the deployment period. This process ensures that once a user receives the updated configuration, they continue to receive it. It also ensures that rollback alarms have sufficient time and data to detect issues.

Consider entity-based gradual deployments if a configuration directly changes user-facing behavior and if the change in blast radius (impacting specific users fully rather than all users partially) is acceptable for your application.

Important

Note the following important information about entity-based gradual deployments:

- Entity-based gradual deployments require AWS AppConfig Agent version 2.0.136060 or later.
- Entity-Id accepts a maximum string size of 2 KB.
- Unique identifiers must not be hard-coded or low cardinality.
- Identifiers aren't sent to the AWS AppConfig service. AWS AppConfig Agent evaluates unique identifiers client-side.
- During a deployment, AWS AppConfig Agent keeps track of entities and deployed configurations. This tracking is maintained only during a deployment period. The tracking ends when the deployment completes.
- Each instance of AWS AppConfig polls for deployment state independently, so agents may briefly serve different versions of the same entity as the deployment percentage changes. This window is determined by the poll interval and a short synchronization period. You can reduce it by setting `RequiredMinimumPollIntervalInSeconds` in the [StartConfigurationSession](#) API action. However, shorter intervals increase API call rates, which can raise costs and risk throttling. Choose a poll interval that meets your requirements. For more information, see [StartConfigurationSession](#) in the AWS CLI Reference.

Enabling entity-based gradual deployments

To enable entity-based gradual deployments:

1. Update to AWS AppConfig Agent version 2.0.136060 or later.
2. Provide a unique identifier in the Entity-Id HTTP header when retrieving configuration data.

The following example request uses an email address for Entity-Id

```
GET /applications/myapp/environments/prod/configurations/featureflags HTTP/1.1
Host: localhost:2772
Entity-Id: example@AWS-example-email.com
```

You can include the Entity-Id header in requests to:

`/applications/{Application}/environments/{Environment}/configurations/{Configuration}`

To view code samples with Entity-Id, see [Using AWS AppConfig Agent to read a freeform configuration profile](#).

The Entity-Id value can be any string that uniquely identifies an entity in your system, such as:

- Customer ID
- Email address
- Account ID
- Backend job ID
- Session-scoped identifier (if appropriate)

Using AWS AppConfig Agent with AWS Lambda

An AWS Lambda extension is a companion process that augments the capabilities of a Lambda function. An extension can start before a function is invoked, run in parallel with a function, and continue to run after a function invocation is processed. In essence, a Lambda extension is like a client that runs in parallel to a Lambda invocation. This parallel client can interface with your function at any point during its lifecycle.

If you use AWS AppConfig feature flags or other dynamic configuration data in a Lambda function, then we recommend that you add the AWS AppConfig Agent Lambda extension as a layer to your Lambda function. This makes calling feature flags simpler, and the extension itself includes best practices that simplify using AWS AppConfig while reducing costs. Reduced costs result from fewer API calls to the AWS AppConfig service and shorter Lambda function processing times. For more information about Lambda extensions, see [Lambda extensions](#) in the *AWS Lambda Developer Guide*.

Note

AWS AppConfig is a capability of AWS Systems Manager. AWS AppConfig [pricing](#) is based on the number of times a configuration is called and received. Your costs increase if your Lambda performs multiple cold starts and retrieves new configuration data frequently.

Topics


- [Understanding how the AWS AppConfig Agent Lambda extension works](#)
- [Adding the AWS AppConfig Agent Lambda extension](#)
- [Configuring the AWS AppConfig Agent Lambda extension](#)

- [Understanding available versions of the AWS AppConfig Agent Lambda extension](#)

Understanding how the AWS AppConfig Agent Lambda extension works

If you use AWS AppConfig to manage configurations for a Lambda function *without* Lambda extensions, then you must configure your Lambda function to receive configuration updates by integrating with the [StartConfigurationSession](#) and [GetLatestConfiguration](#) API actions.

Integrating the AWS AppConfig Agent Lambda extension with your Lambda function simplifies this process. The extension takes care of calling the AWS AppConfig service, managing a local cache of retrieved data, tracking the configuration tokens needed for the next service calls, and periodically checking for configuration updates in the background. The following diagram shows how it works.



1. You configure the AWS AppConfig Agent Lambda extension as a layer of your Lambda function.

2. To access its configuration data, your function calls the AWS AppConfig extension at an HTTP endpoint running on `localhost:2772`.
3. The extension maintains a local cache of the configuration data. If the data isn't in the cache, the extension calls AWS AppConfig to get the configuration data.
4. Upon receiving the configuration from the service, the extension stores it in the local cache and passes it to the Lambda function.
5. AWS AppConfig Agent Lambda extension periodically checks for updates to your configuration data in the background. Each time your Lambda function is invoked, the extension checks the elapsed time since it retrieved a configuration. If the elapsed time is greater than the configured poll interval, the extension calls AWS AppConfig to check for newly deployed data, updates the local cache if there has been a change, and resets the elapsed time.

Note

- Lambda instantiates separate instances corresponding to the concurrency level that your function requires. Each instance is isolated and maintains its own local cache of your configuration data. For more information about Lambda instances and concurrency, see [Managing concurrency for a Lambda function](#).
- The amount of time it takes for a configuration change to appear in a Lambda function, after you deploy an updated configuration from AWS AppConfig, depends on the deployment strategy you used for the deployment and the polling interval you configured for the extension.

Adding the AWS AppConfig Agent Lambda extension

To use the AWS AppConfig Agent Lambda extension, you need to add the extension to your Lambda. This can be done by adding the AWS AppConfig Agent Lambda extension to your Lambda function as a layer or by enabling the extension on a Lambda function as a container image.

Note

The AWS AppConfig extension is runtime agnostic and supports all runtimes.

Before you begin

Before you enable the AWS AppConfig Agent Lambda extension, do the following:

- Organize the configurations in your Lambda function so that you can externalize them into AWS AppConfig.
- Create AWS AppConfig artifacts and configuration data, including feature flags or freeform configuration data. For more information, see [Creating feature flags and free form configuration data in AWS AppConfig](#).
- Add `appconfig:StartConfigurationSession` and `appconfig:GetLatestConfiguration` to the AWS Identity and Access Management (IAM) policy used by the Lambda function execution role. For more information, see [AWS Lambda execution role](#) in the *AWS Lambda Developer Guide*. For more information about AWS AppConfig permissions, see [Actions, resources, and condition keys for AWS AppConfig](#) in the *Service Authorization Reference*.

Adding the AWS AppConfig Agent Lambda extension by using a layer and an ARN

To use the AWS AppConfig Agent Lambda extension, you add the extension to your Lambda function as a layer. For information about how to add a layer to your function, see [Configuring extensions](#) in the *AWS Lambda Developer Guide*. The name of the extension in the AWS Lambda console is **AWS-AppConfig-Extension**. Also note that when you add the extension as a layer to your Lambda, you must specify an Amazon Resource Name (ARN). Choose an ARN from one of the following lists that corresponds with the platform and AWS Region where you created the Lambda.

- [x86-64 platform](#)
- [ARM64 platform](#)

If you want to test the extension before you add it to your function, you can verify that it works by using the following code example.

```
import urllib.request

def lambda_handler(event, context):
    url = f'http://localhost:2772/applications/application_name/environments/environment_name/configurations/configuration_name'
    config = urllib.request.urlopen(url).read()
    return config
```

To test it, create a new Lambda function for Python, add the extension, and then run the Lambda function. After you run the Lambda function, the AWS AppConfig Lambda function returns the configuration you specified for the `http://localhost:2772` path. For information about creating a Lambda function, see [Create a Lambda function with the console](#) in the *AWS Lambda Developer Guide*.

Important

You can view log data for the AWS AppConfig Agent Lambda extension in the AWS Lambda logs. Log entries are prefaced with `appconfig agent`. Here's an example.

```
[appconfig agent] 2024/05/07 04:19:01 ERROR retrieve failure for
'SourceEventConfig:SourceEventConfigEnvironment:SourceEventConfigProfile':
StartConfigurationSession: api error AccessDenied: User:
arn:aws:sts::0123456789:assumed-role/us-east-1-LambdaRole/
extension1 is not authorized to perform: sts:AssumeRole on resource:
arn:aws:iam::0123456789:role/test1 (retry in 60s)
```

Configuring the AWS AppConfig Agent Lambda extension

You can configure the extension by changing the following AWS Lambda environment variables. For more information, see [Using AWS Lambda environment variables](#) in the *AWS Lambda Developer Guide*.

Prefetching configuration data

The environment variable `AWS_APPCONFIG_EXTENSION_PREFETCH_LIST` can improve the start-up time of your function. When the AWS AppConfig Agent Lambda extension is initialized, it retrieves the specified configuration from AWS AppConfig before Lambda starts to initialize your function and invoke your handler. In some cases, the configuration data is already available in the local cache before your function requests it.

To use the prefetch capability, set the value of the environment variable to the path corresponding to your configuration data. For example, if your configuration corresponds to an application, environment, and configuration profile respectively named "my_application", "my_environment", and "my_configuration_data", the path would be `/applications/my_application/environments/my_environment/configurations/my_configuration_data`. You can

specify multiple configuration items by listing them as a comma-separated list (If you have a resource name that includes a comma, use the resource's ID value instead of its name).

Accessing configuration data from another account

The AWS AppConfig Agent Lambda extension can retrieve configuration data from another account by specifying an IAM role that grants [permissions](#) to the data. To set this up, follow these steps:

1. In the account where AWS AppConfig is used to manage the configuration data, create a role with a trust policy that grants the account running the Lambda function access to the `appconfig:StartConfigurationSession` and `appconfig:GetLatestConfiguration` actions, along with the partial or complete ARNs corresponding to the AWS AppConfig configuration resources.
2. In the account running the Lambda function, add the `AWS_APPCONFIG_EXTENSION_ROLE_ARN` environment variable to the Lambda function with the ARN of the role created in step 1.
3. (Optional) If needed, an [external ID](#) can be specified using the `AWS_APPCONFIG_EXTENSION_ROLE_EXTERNAL_ID` environment variable. Similarly, a session name can be configured using the `AWS_APPCONFIG_EXTENSION_ROLE_SESSION_NAME` environment variable.

Note

Note the following information.

- The AWS AppConfig Agent Lambda extension can only retrieve data from one account. If you specify an IAM role, the extension will not be able to retrieve configuration data from the account in which the Lambda function is running.
- AWS Lambda logs information about the AWS AppConfig Agent Lambda extension and the Lambda function by using Amazon CloudWatch Logs.
- The following table includes a **Sample values** column. Depending on your monitor resolution, you might need to scroll to the bottom of the table and then scroll to the right to view the column.

Environment variable	Details	Default value	Sample values
AWS_APPCONFIG_EXTENSION_ASSIGNMENT_LOG_DESTINATION	This environment variable enables AWS AppConfig Agent to log experiment assignment records, which capture each occurrence of an experiment treatment being assigned to a given entity at a given time. Possible values: <code>file:<path>/</code> indicates a base directory under which assignment records are written to disk; <code>stderr</code> indicates that assignment records are emitted to stderr. This feature requires AWS AppConfig Agent Lambda extension version 2.0.20159 or later.	None	<code>file:/tmp/aws-appconfig/assignments/</code> <code>stderr</code>
AWS_APPCONFIG_EXTENSION_HTTP_PORT	This environment variable specifies the port on which the local HTTP server that hosts the extension runs.	2772	2772

Environment variable	Details	Default value	Sample values
AWS_APPCONFIG_EXTENSION_LOG_LEVEL	This environment variable specifies the level of detail that the agent logs. Each level includes the current level and all higher levels. The value is case insensitive. From most to least detailed, the log levels are: trace, debug, info, warn, error, fatal, and none. The trace log includes detailed information, including timing information, about the agent.	info	trace debug info warn error fatal none
AWS_APPCONFIG_EXTENSION_MAX_CONNECTIONS	This environment variable configures the maximum number of connections the extension uses to retrieve configurations from AWS AppConfig.	3	3

Environment variable	Details	Default value	Sample values
AWS_APPCONFIG_EXTENSION_POLL_INTERVAL_SECONDS	This environment variable controls how often the agent polls AWS AppConfig for updated configuration data. You can specify a number of seconds for the interval. You can also specify a number with a time unit: s for seconds, m for minutes, and h for hours. If a unit isn't specified, the agent defaults to seconds. For example, 60, 60s, and 1m result in the same poll interval.	45	45 45s 5m 1h

Environment variable	Details	Default value	Sample values
AWS_APPCONFIG_EXTENSION_POLL_TIMEOUT_MILLIS	This environment variable controls the maximum amount of time, in milliseconds, the extension waits for a response from AWS AppConfig when refreshing data in the cache. If AWS AppConfig does not respond in the specified amount of time, the extension skips this poll interval and returns the previously updated cached data.	3000ms	3000 300ms 5s

Environment variable	Details	Default value	Sample values
AWS_APPCONFIG_EXTENSION_PREFETCH_LIST	This environment variable specifies the configuration data the agent requests from AWS AppConfig as soon as it starts. Multiple configuration identifiers may be provided in a comma-separated list. Prefetching configuration data from AWS AppConfig can reduce the cold start time of your function significantly.	None	MyApp:MyEnv:MyConfig abcd123:efgh456:ijkl789 MyApp:MyEnv:Config1,MyApp:MyEnv:Config2
AWS_APPCONFIG_PROXY_HEADERS	This environment variable specifies headers required by the proxy referenced in the AWS_APPCONFIG_EXTENSION_PROXY_URL environment variable. The value is a comma-separated list of headers.	None	header: value h1: v1, h2: v2

Environment variable	Details	Default value	Sample values
AWS_APPCONFIG_EXTENSION_PROXY_URL	This environment variable specifies the proxy URL to use for connections from the AWS AppConfig extension to AWS services. HTTPS and HTTP URLs are supported.	None	http://localhost:7474 https://my-proxy.example.com
AWS_APPCONFIG_EXTENSION_ROLE_ARN	This environment variable specifies the IAM role ARN corresponding to a role that should be assumed by the AWS AppConfig extension to retrieve configuration.	None	arn:aws:iam::123456789012:role/MyRole
AWS_APPCONFIG_EXTENSION_ROLE_EXTERNAL_ID	This environment variable specifies the external id to use in conjunction with the assumed role ARN.	None	MyExternalId
AWS_APPCONFIG_EXTENSION_ROLE_SESSION_NAME	This environment variable specifies the session name to be associated with the credentials for the assumed IAM role.	None	AWSAppConfigAgentSession

Environment variable	Details	Default value	Sample values
AWS_APPCONFIG_EXTENSION_SERVICE_REGION	This environment variable specifies an alternative Region the extension should use to call the AWS AppConfig service. When undefined, the extension uses the endpoint in the current Region.	None	us-east-1 eu-west-1
AWS_APPCONFIG_EXTENSION_MANIFEST	This environment variable configures AWS AppConfig Agent to take advantage of additional per-configuration features like multi-account retrievals and save configuration to disk. For more information about these features, see Using a manifest to enable additional retrieval features .	None	When using AWS AppConfig configuration as manifest: MyApp:MyEnvironment:MyManifestConfig . When loading manifest from disk: file:/path/to/manifest.json
AWS_APPCONFIG_EXTENSION_WAIT_ON_MANIFEST	This environment variable configures AWS AppConfig Agent to wait until the manifest is processed before completing startup.	true	true false

Understanding available versions of the AWS AppConfig Agent Lambda extension

This topic includes information about AWS AppConfig Agent Lambda extension versions. The AWS AppConfig Agent Lambda extension supports Lambda functions developed for the x86-64 and ARM64 (Graviton2) platforms. To work properly, your Lambda function must be configured to use the specific Amazon Resource Name (ARN) for the AWS Region where it is currently hosted. You can view AWS Region and ARN details later in this section.

Important

Note the following important details about the AWS AppConfig Agent Lambda extension.

- The `GetConfiguration` API action was deprecated on January 28, 2022. Calls to receive configuration data should use the `StartConfigurationSession` and `GetLatestConfiguration` APIs instead. If you are using a version of the AWS AppConfig Agent Lambda extension created after January 28, 2022, you need to configure permissions to the new APIs. For more information, see [Retrieving configuration data without AWS AppConfig Agent](#).
- AWS AppConfig supports all of the versions listed in [Older extension versions](#). We recommend that you periodically update to the latest version to take advantage of extension enhancements.

Topics

- [AWS AppConfig Agent Lambda Extension release notes](#)
- [Finding your Lambda extension version number](#)
- [Retrieving the latest extension ARN from a public parameter](#)
- [x86-64 platform](#)
- [ARM64 platform](#)
- [Older extension versions](#)

AWS AppConfig Agent Lambda Extension release notes

The following table describes changes made to recent versions of the AWS AppConfig Lambda extension.

Version	Launch date	Notes
2.0.22365	08/24/2026	Added support for experiment assignment logging to stderr. Minor enhancements and bug fixes.
2.0.20159	07/13/2026	Added support for AWS AppConfig experimentation.
2.0.18836	06/22/2026	Minor enhancements and bug fixes.
2.0.17054	05/14/2026	Minor enhancements and bug fixes.
2.0.14126	03/25/2026	Minor enhancements and bug fixes.
2.0.11962	02/20/2026	Improved environment support, minor enhancements, and bug fixes.
2.0.8693	11/20/2025	Improved environment support, minor enhancements, and bug fixes. Added support for the following AWS Regions • Asia Pacific (Taipei), ap-east-2 • Asia Pacific (New Zealand), ap-southeast-6 • Asia Pacific (Thailand), ap-southeast-7 • Mexico (Central), mx-central-1

Version	Launch date	Notes
2.0.2037	05/12/2025	Added /ping path, which exposes a simple health check that returns that agent's version. Also includes minor enhancements and bug fixes.
2.0.1079	12/12/2024	Minor enhancements and bug fixes.
2.0.719	08/08/2024	Minor enhancements and bug fixes.
2.0.678	07/23/2024	Enhancements to support feature flag targets, variants, and splits. For more information, see Creating multi-variant feature flags .
2.0.501	07/01/2024	Minor enhancements and bug fixes.

Version	Launch date	Notes
2.0.358	12/01/2023	Added support for the following retrieval features: • <i>Multi-account retrieval</i>: Use AWS AppConfig Agent from a primary or <i>retrieval</i> AWS account to retrieve configuration data from multiple vendor accounts. • <i>Write configuration copy to disk</i>: Use AWS AppConfig Agent to write configuration data to disk. This feature enables customers with applications that read configuration data from disk to integrate with AWS AppConfig.
2.0.181	08/14/2023	Added support for the Israel (Tel Aviv) il-central-1 AWS Region.

Version	Launch date	Notes
2.0.165	02/21/2023	Minor bug fixes. No longer restricting extension use to specific runtime versions via the AWS Lambda console. Added support for the following AWS Regions: • Middle East (UAE), me-central-1 • Asia Pacific (Hyderabad), ap-south-2 • Asia Pacific (Melbourne), ap-southeast-4 • Europe (Spain), eu-south-2 • Europe (Zurich), eu-central-2
2.0.122	08/23/2022	Added support for a tunneling proxy, which can be configured with the <code>AWS_APPCONFIG_EXTENSION_PROXY_URL</code> and <code>AWS_APPCONFIG_EXTENSION_PROXY_HEADER</code> environment variables. Added .NET 6 as a runtime. For more information about environment variables, see Configuring the AWS AppConfig Agent Lambda extension.

Version	Launch date	Notes
2.0.58	05/03/2022	Improved support for Graviton2 (ARM64) processors in Lambda.
2.0.45	03/15/2022	Added support for calling a single feature flag. Previously, customers called feature flags grouped into a configuration profile and had to parse the response client-side. With this release, customers can use a <code>flag=<flag-name></code> parameter when calling the HTTP localhost endpoint to get the value of a single flag. Also added initial support for Graviton2 (ARM64) processors.

Finding your Lambda extension version number

Use the following procedure to locate the version number of your currently configured AWS AppConfig Agent Lambda extension. To work properly, your Lambda function must be configured to use the specific Amazon Resource Name (ARN) for the AWS Region where it is currently hosted.

1. Sign in to the AWS Management Console and open the AWS Lambda console at <https://console.aws.amazon.com/lambda/>.
2. Choose the Lambda function where you want to add the AWS-AppConfig-Extension layer.
3. In the **Layers** section, choose **Add a layer**.
4. In the **Choose a layer** section, choose **AWS-AppConfig-Extension** from the **AWS layers** list.
5. Use the **Version** list to choose a version number.
6. Choose **Add**.
7. Use the **Test** tab to test the function.

8. After the test completes, view the log output. Locate the AWS AppConfig Agent Lambda extension version in the **Details of the Execution** section. This version must match the required URLs for that version.

Retrieving the latest extension ARN from a public parameter

The latest AWS AppConfig Agent Lambda extension layer ARN is available as a public parameter in AWS Systems Manager Parameter Store in every commercial AWS Region. You can reference these parameters programmatically to always deploy your Lambda function with the latest extension version without manual ARN updates.

The public parameter paths are:

- **x86_64**: `/aws/service/aws-appconfig/lambda-extension/x86/latest`
- **arm64**: `/aws/service/aws-appconfig/lambda-extension/arm64/latest`

Note

The x86 parameter path uses x86 (not x86_64), even though the Lambda architecture is x86_64.

You don't need special IAM permissions beyond `ssm:GetParameter`. The parameter is public and readable from any AWS account.

AWS CLI

To retrieve the latest ARN versions, run the following commands.

x86_64

```
aws ssm get-parameter \
  --name "/aws/service/aws-appconfig/lambda-extension/x86/latest" \
  --query "Parameter.Value" \
  --output text
```

arm64

```
aws ssm get-parameter \
  --name "/aws/service/aws-appconfig/lambda-extension/arm64/latest" \
```

```
--query "Parameter.Value" \  
--output text
```

CloudFormation

When you deploy Lambda functions with CloudFormation, you can use the `{{resolve:ssm:...}}` dynamic reference syntax to retrieve the latest extension ARN during stack operations. This ensures that your function always uses the latest extension version without manual updates.

x86_64

```
Resources:  
  MyFunction:  
    Type: AWS::Lambda::Function  
    Properties:  
      FunctionName: my-function  
      Runtime: python3.13  
      Handler: index.handler  
      Architectures:  
        - x86_64  
      Code:  
        ZipFile: |  
          def handler(event, context):  
              return {'statusCode': 200}  
      Layers:  
        - '{{resolve:ssm:/aws/service/aws-appconfig/lambda-extension/x86/latest}}'  
      Role: !GetAtt MyFunctionRole.Arn
```

arm64

```
Resources:  
  MyFunction:  
    Type: AWS::Lambda::Function  
    Properties:  
      FunctionName: my-function  
      Runtime: python3.13  
      Handler: index.handler  
      Architectures:  
        - arm64  
      Code:  
        ZipFile: |  
          def handler(event, context):
```

```

        return {'statusCode': 200}
    Layers:
    - '{{resolve:ssm:/aws/service/aws-appconfig/lambda-extension/arm64/latest}}'
    Role: !GetAtt MyFunctionRole.Arn

```

Note

The `{{resolve:ssm:parameter-name}}` syntax automatically retrieves the parameter value during stack operations. This ensures that you always deploy with the current ARN.

x86-64 platform

When you add the extension as a layer to your Lambda, you must specify an ARN. Choose an ARN from the following table that corresponds with the AWS Region where you created the Lambda. These ARNs are for Lambda functions developed for the x86-64 platform.

Version 2.0.22365

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:332
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:287
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:408
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:381
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:268

Region	ARN
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:178
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:301
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:228
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:309
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:252
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:279
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:373
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:259
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:221

Region	ARN
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:249
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:247
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:263
Asia Pacific (Taipei)	arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension:153
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:278
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:282
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:280
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:264
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:319

Region	ARN
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:262
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:198
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension:170
Asia Pacific (New Zealand)	arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension:103
Asia Pacific (Thailand)	arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension:140
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:297
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:223
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:338
Mexico (Central)	arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension:148

Region	ARN
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:266
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:194
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:227
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:226
AWS European Sovereign Cloud (Germany)	arn:aws-eusc:lambda:eusc-de-east-1:426221636601:layer:AWS-AppConfig-Extension:73

ARM64 platform

When you add the extension as a layer to your Lambda, you must specify an ARN. Choose an ARN from the following table that corresponds with the AWS Region where you created the Lambda. These ARNs are for Lambda functions developed for the ARM64 platform.

Version 2.0.22365

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:265

Region	ARN
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:239
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:285
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:283
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:188
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension-Arm64:168
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:244
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:186
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:247
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:204

Region	ARN
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:198
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:232
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:187
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:182
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension-Arm64:171
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension-Arm64:169
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:191
Asia Pacific (Taipei)	arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension-Arm64:127
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:231

Region	ARN
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:190
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:195
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:216
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:262
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:199
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:183
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension-Arm64:145
Asia Pacific (New Zealand)	arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension-Arm64:93
Asia Pacific (Thailand)	arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension-Arm64:139

Region	ARN
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:239
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:181
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:226
Mexico (Central)	arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension-Arm64:147
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:194
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:177
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension-Arm64:173
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension-Arm64:172
AWS European Sovereign Cloud (Germany)	arn:aws-eusc:lambda:eusc-de-east-1:426221636601:layer:AWS-AppConfig-Extension-Arm64:48

Older extension versions

This section lists the ARNs and AWS Regions for older versions of the AWS AppConfig Lambda extension. This list doesn't contain information for all previous versions of the AWS AppConfig Agent Lambda extension, but it will be updated when new versions are released.

Topics

- [Older extension versions \(x86-64 platform\)](#)
- [Older extension versions \(ARM64 platform\)](#)

Older extension versions (x86-64 platform)

The following tables list ARNs and the AWS Regions for older versions of the AWS AppConfig Agent Lambda extension developed for the x86-64 platform.

Date replaced by newer extension: 08/24/2026

Version 2.0.20159

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:329
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:280
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:404
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:378

Region	ARN
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:265
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:175
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:298
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:223
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:306
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:249
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:276
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:370
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:252

Region	ARN
Europe (Spain)	<code>arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:217</code>
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:246</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:244</code>
Asia Pacific (Hong Kong)	<code>arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:259</code>
Asia Pacific (Taipei)	<code>arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension:150</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:275</code>
Asia Pacific (Seoul)	<code>arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:279</code>
Asia Pacific (Osaka)	<code>arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:277</code>
Asia Pacific (Singapore)	<code>arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:261</code>

Region	ARN
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:316
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:259
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:195
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension:164
Asia Pacific (New Zealand)	arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension:100
Asia Pacific (Thailand)	arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension:137
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:293
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:220
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:335

Region	ARN
Mexico (Central)	arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension:145
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:263
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:191
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:222
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:223
AWS European Sovereign Cloud (Germany)	arn:aws-eusc:lambda:eusc-de-east-1:426221636601:layer:AWS-AppConfig-Extension:73

Date replaced by newer extension: 07/13/2026

Version 2.0.18836

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:328

Region	ARN
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:279
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:403
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:375
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:264
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:174
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:297
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:222
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:305
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:248

Region	ARN
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:275
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:369
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:251
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:216
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:245
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:243
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:258
Asia Pacific (Taipei)	arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension:149
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:274

Region	ARN
Asia Pacific (Seoul)	<code>arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:278</code>
Asia Pacific (Osaka)	<code>arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:276</code>
Asia Pacific (Singapore)	<code>arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:260</code>
Asia Pacific (Sydney)	<code>arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:315</code>
Asia Pacific (Jakarta)	<code>arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:258</code>
Asia Pacific (Melbourne)	<code>arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:194</code>
Asia Pacific (Malaysia)	<code>arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension:163</code>
Asia Pacific (New Zealand)	<code>arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension:99</code>
Asia Pacific (Thailand)	<code>arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension:136</code>

Region	ARN
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:292
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:219
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:334
Mexico (Central)	arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension:144
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:262
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:190
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension:212
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:254
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:221

Region	ARN
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:222
AWS European Sovereign Cloud (Germany)	arn:aws-eusc:lambda:eusc-de-east-1:426221636601:layer:AWS-AppConfig-Extension:74

Date replaced by newer extension: 06/22/2026

Version 2.0.17054

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:321
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:277
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:395
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:373
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:262

Region	ARN
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:172
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:295
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:220
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:303
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:242
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:273
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:367
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:249
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:214

Region	ARN
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:243
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:241
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:254
Asia Pacific (Taipei)	arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension:147
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:272
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:273
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:272
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:258
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:313

Region	ARN
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:256
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:188
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension:161
Asia Pacific (New Zealand)	arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension:97
Asia Pacific (Thailand)	arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension:134
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:290
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:217
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:327
Mexico (Central)	arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension:142

Region	ARN
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:260
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:188
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension:212
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:254
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:219
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:220

Date replaced by newer extension: 05/14/2026

Version 2.0.14126

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:317

Region	ARN
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:273
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:383
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:369
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:258
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:168
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:291
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:216
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:299
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:238

Region	ARN
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:269
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:363
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:245
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:210
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:249
Asia Pacific (Taipei)	arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension:139
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:286
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:213
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:268

Region	ARN
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:269
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:268
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:254
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:309
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:252
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:184
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension:157
Asia Pacific (New Zealand)	arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension:87
Asia Pacific (Thailand)	arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension:130

Region	ARN
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:256
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:184
Mexico (Central)	arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension:138
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:323
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:236
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:237
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:215
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:214

Date replaced by newer extension: 03/25/2026

Version 2.0.11962

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:296
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:252
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:359
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:348
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:239
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:147
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:270
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:195
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:278

Region	ARN
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:217
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:248
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:342
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:226
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:189
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:219
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:221
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:228
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:245

Region	ARN
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:248
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:247
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:233
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:288
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:231
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:163
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension:136
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:264
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:192

Region	ARN
Asia Pacific (New Zealand)	arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension:58
Asia Pacific (Thailand)	arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension:109
Asia Pacific (Taipei)	arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension:118
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:302
Mexico (Central)	arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension:115
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:234
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:168
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension:206
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:244

Region	ARN
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:184
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:182

Date replaced by newer extension: 02/17/2026

Version 2.0.8693

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:279
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:235
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:348
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:335
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:228

Region	ARN
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:130
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:261
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:178
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:261
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:207
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:235
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:333
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:215
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:176

Region	ARN
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:205</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:203</code>
Asia Pacific (Hong Kong)	<code>arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:217</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:228</code>
Asia Pacific (Seoul)	<code>arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:239</code>
Asia Pacific (Osaka)	<code>arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:234</code>
Asia Pacific (Singapore)	<code>arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:224</code>
Asia Pacific (Sydney)	<code>arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:272</code>
Asia Pacific (Jakarta)	<code>arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:222</code>

Region	ARN
Asia Pacific (Melbourne)	<code>arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:152</code>
Asia Pacific (Malaysia)	<code>arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension:127</code>
Asia Pacific (Mumbai)	<code>arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:248</code>
Asia Pacific (Hyderabad)	<code>arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:179</code>
Asia Pacific (New Zealand)	<code>arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension:41</code>
Asia Pacific (Thailand)	<code>arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension:98</code>
Asia Pacific (Taipei)	<code>arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension:100</code>
South America (São Paulo)	<code>arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:288</code>
Mexico (Central)	<code>arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension:98</code>

Region	ARN
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:225
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:155
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension:195
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:227
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:184
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:182

Date replaced by newer extension: 11/20/2025

Version 2.0.2037

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:207

Region	ARN
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:162
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:258
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:262
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:152
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:57
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:189
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:106
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:189
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:133

Region	ARN
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:162
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:259
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:140
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:102
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:133
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:131
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:142
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:155
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:165

Region	ARN
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:159
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:156
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:199
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:150
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:78
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension:55
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:175
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:104
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:215

Region	ARN
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:152
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:81
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension:120
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:154
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:110
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:109

Date replaced by newer extension: 05/20/2025

Version 2.0.1079

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:174

Region	ARN
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:133
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:223
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:230
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:123
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:27
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:159
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:77
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:160
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:121

Region	ARN
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:133
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:225
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:111
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:74
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:106
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:104
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:113
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:126
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:136

Region	ARN
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:130
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:134
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:165
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:121
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:49
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension:26
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:146
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:75
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:179

Region	ARN
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:123
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:52
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension:91
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:125
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:80
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:80

Date replaced by newer extension: 12/12/2024

Version 2.0.719

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:173

Region	ARN
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:132
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:221
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:229
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:121
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:27
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:158
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:75
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:159
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:120

Region	ARN
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:132
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:224
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:110
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:72
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:104
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:102
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:112
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:125
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:135

Region	ARN
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:129
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:132
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:164
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:120
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:48
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension:25
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:145
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:74
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:178

Region	ARN
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:122
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:50
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension:90
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:124
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:79
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:79

Date replaced by newer extension: 08/08/2024

Version 2.0.678

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:167

Region	ARN
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:126
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:213
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:223
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:116
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:21
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:152
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:70
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:153
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:114

Region	ARN
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:126
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:218
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:104
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:67
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:99
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:97
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:106
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:119
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:129

Region	ARN
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:123
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:127
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:158
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:114
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:42
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:139
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:68
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:172
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:116

Region	ARN
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:45
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension:84
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:118
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:73
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:73

Date replaced by newer extension: 07/23/2024

Version 2.0.501

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:153
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:112

Region	ARN
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:195
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:210
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:101
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:136
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:53
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:144
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:99
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:111
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:201

Region	ARN
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:89</code>
Europe (Spain)	<code>arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:50</code>
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:85</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:83</code>
Asia Pacific (Hong Kong)	<code>arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:91</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:104</code>
Asia Pacific (Seoul)	<code>arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:114</code>
Asia Pacific (Osaka)	<code>arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:107</code>
Asia Pacific (Singapore)	<code>arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:112</code>

Region	ARN
Asia Pacific (Sydney)	<code>arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:142</code>
Asia Pacific (Jakarta)	<code>arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:98</code>
Asia Pacific (Melbourne)	<code>arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:26</code>
Asia Pacific (Mumbai)	<code>arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:125</code>
Asia Pacific (Hyderabad)	<code>arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:53</code>
South America (São Paulo)	<code>arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:155</code>
Africa (Cape Town)	<code>arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:102</code>
Israel (Tel Aviv)	<code>arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:28</code>
Middle East (UAE)	<code>arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension:68</code>

Region	ARN
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:103
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:59
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:59

Date replaced by newer extension: 07/01/2024

Version 2.0.358

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:128
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:93
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:141
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:161

Region	ARN
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:93
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:106
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:47
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:125
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:93
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:98
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:159
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:83
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:44

Region	ARN
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:76</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:76</code>
Asia Pacific (Hong Kong)	<code>arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:83</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:98</code>
Asia Pacific (Seoul)	<code>arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:108</code>
Asia Pacific (Osaka)	<code>arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:101</code>
Asia Pacific (Singapore)	<code>arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:106</code>
Asia Pacific (Sydney)	<code>arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:106</code>
Asia Pacific (Jakarta)	<code>arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:79</code>

Region	ARN
Asia Pacific (Melbourne)	<code>arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:20</code>
Asia Pacific (Mumbai)	<code>arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:107</code>
Asia Pacific (Hyderabad)	<code>arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:47</code>
South America (São Paulo)	<code>arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:128</code>
Africa (Cape Town)	<code>arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:83</code>
Israel (Tel Aviv)	<code>arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:22</code>
Middle East (UAE)	<code>arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension:49</code>
Middle East (Bahrain)	<code>arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:85</code>
AWS GovCloud (US-East)	<code>arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:54</code>

Region	ARN
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:54

Date replaced by newer extension: 12/01/2023

Version 2.0.181

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:113
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:81
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:124
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:146
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:81
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:93

Region	ARN
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:32
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:110
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:81
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:82
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:142
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:73
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:29
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:68
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:68

Region	ARN
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:73
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:84
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:93
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:86
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:91
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:93
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:64
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:5
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:94

Region	ARN
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:32
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:113
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:73
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension:7
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension:34
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:73
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:46
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:46

Date replaced by newer extension: 08/14/2023

Version 2.0.165

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:110
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:79
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:121
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:143
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:79
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:91
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension:29
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:108
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:79

Region	ARN
Europe (Paris)	<code>arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:80</code>
Europe (Stockholm)	<code>arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:139</code>
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:71</code>
Europe (Spain)	<code>arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension:26</code>
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:66</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:66</code>
Asia Pacific (Hong Kong)	<code>arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:71</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:82</code>
Asia Pacific (Seoul)	<code>arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:91</code>

Region	ARN
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:84
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:89
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:91
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:60
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension:2
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:92
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension:29
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:110
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:71

Region	ARN
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension:31
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:71
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:44
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:44

Date replaced by newer extension: 02/21/2023

Version 2.0.122

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:82
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:59
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:93

Region	ARN
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:114
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:59
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:70
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:82
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:59
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:60
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:111
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:54
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:52

Region	ARN
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:52
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:54
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:62
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:70
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:59
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:64
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:70
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:37
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:71

Region	ARN
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:82
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:54
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:54
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:29
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:29

Date replaced by newer extension: 08/23/2022

Version 2.0.58

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:69
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:50

Region	ARN
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:78
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:101
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:50
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:59
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:69
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:50
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:51
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:98
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:47

Region	ARN
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:46</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:46</code>
Asia Pacific (Hong Kong)	<code>arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:47</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:49</code>
Asia Pacific (Seoul)	<code>arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:59</code>
Asia Pacific (Osaka)	<code>arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:46</code>
Asia Pacific (Singapore)	<code>arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:51</code>
Asia Pacific (Sydney)	<code>arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:59</code>
Asia Pacific (Jakarta)	<code>arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:24</code>

Region	ARN
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:60
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:69
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:47
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:47
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:23
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:23

Date replaced by newer extension: 04/21/2022

Version 2.0.45

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:68

Region	ARN
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:49
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:77
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:100
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:49
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:58
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:68
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:49
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:50
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:97

Region	ARN
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:46
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:45
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:45
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:46
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:48
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:58
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:45
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:50
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:58

Region	ARN
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:23
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:59
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:68
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:46
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:46
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:22
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:22

Date replaced by newer extension: 03/15/2022

Version 2.0.30

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension:61
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension:47
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension:61
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension:89
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension:47
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension:54
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension:59
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension:47
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension:48

Region	ARN
Europe (Stockholm)	<code>arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension:86</code>
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension:44</code>
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension:43</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension:43</code>
Asia Pacific (Hong Kong)	<code>arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension:44</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension:45</code>
Asia Pacific (Osaka)	<code>arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension:42</code>
Asia Pacific (Seoul)	<code>arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension:54</code>
Asia Pacific (Singapore)	<code>arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension:45</code>

Region	ARN
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension:54
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension:13
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension:55
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension:61
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension:44
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension:44
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension:20
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension:20

Older extension versions (ARM64 platform)

The following tables list ARNs and the AWS Regions for older versions of the AWS AppConfig Agent Lambda extension developed for the ARM64 platform.

Date replaced by newer extension: 08/24/2026

Version 2.0.20159

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:262
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:232
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:281
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:280
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:185
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension-Arm64:165
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:241
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:181

Region	ARN
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:244
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:201
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:195
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:229
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:180
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:178
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension-Arm64:168
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension-Arm64:166
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:187

Region	ARN
Asia Pacific (Taipei)	arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension-Arm64:124
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:228
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:187
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:192
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:213
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:259
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:196
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:180
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension-Arm64:139

Region	ARN
Asia Pacific (New Zealand)	arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension-Arm64:90
Asia Pacific (Thailand)	arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension-Arm64:136
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:235
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:178
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:223
Mexico (Central)	arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension-Arm64:144
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:191
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:174
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension-Arm64:168

Region	ARN
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension-Arm64:169
AWS European Sovereign Cloud (Germany)	arn:aws-eusc:lambda:eusc-de-east-1:426221636601:layer:AWS-AppConfig-Extension-Arm64:48

Date replaced by newer extension: 07/13/2026

Version 2.0.18836

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:261
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:231
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:280
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:277
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:184

Region	ARN
Canada West (Calgary)	<code>arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension-Arm64:164</code>
Europe (Frankfurt)	<code>arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:240</code>
Europe (Zurich)	<code>arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:180</code>
Europe (Ireland)	<code>arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:243</code>
Europe (London)	<code>arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:200</code>
Europe (Paris)	<code>arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:194</code>
Europe (Stockholm)	<code>arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:228</code>
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:179</code>
Europe (Spain)	<code>arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:177</code>

Region	ARN
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:186
Asia Pacific (Taipei)	arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension-Arm64:123
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:227
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:186
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:191
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:212
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:258
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:195
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:179

Region	ARN
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension-Arm64:138
Asia Pacific (New Zealand)	arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension-Arm64:89
Asia Pacific (Thailand)	arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension-Arm64:135
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:234
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:177
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:222
Mexico (Central)	arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension-Arm64:143
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:190
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension-Arm64:168

Region	ARN
Middle East (Bahrain)	<code>arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension-Arm64:182</code>
Israel (Tel Aviv)	<code>arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:173</code>
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension-Arm64:167</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension-Arm64:165</code>
AWS GovCloud (US-East)	<code>arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension-Arm64:167</code>
AWS GovCloud (US-West)	<code>arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension-Arm64:168</code>
AWS European Sovereign Cloud (Germany)	<code>arn:aws-eusc:lambda:eusc-de-east-1:426221636601:layer:AWS-AppConfig-Extension-Arm64:49</code>

Date replaced by newer extension: 06/22/2026

Version 2.0.17054

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:254
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:229
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:272
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:275
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:182
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension-Arm64:162
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:238
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:178
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:241

Region	ARN
Europe (London)	<code>arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:194</code>
Europe (Paris)	<code>arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:192</code>
Europe (Stockholm)	<code>arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:226</code>
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:177</code>
Europe (Spain)	<code>arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:175</code>
Asia Pacific (Hong Kong)	<code>arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:182</code>
Asia Pacific (Taipei)	<code>arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension-Arm64:121</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:225</code>
Asia Pacific (Seoul)	<code>arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:181</code>

Region	ARN
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:187
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:210
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:256
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:193
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:173
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension-Arm64:136
Asia Pacific (New Zealand)	arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension-Arm64:87
Asia Pacific (Thailand)	arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension-Arm64:133
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:232

Region	ARN
Asia Pacific (Hyderabad)	<code>arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:175</code>
South America (São Paulo)	<code>arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:215</code>
Mexico (Central)	<code>arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension-Arm64:141</code>
Africa (Cape Town)	<code>arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:188</code>
Middle East (UAE)	<code>arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension-Arm64:168</code>
Middle East (Bahrain)	<code>arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension-Arm64:182</code>
Israel (Tel Aviv)	<code>arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:171</code>
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension-Arm64:165</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension-Arm64:163</code>

Region	ARN
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension-Arm64:165
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension-Arm64:166

Date replaced by newer extension: 05/14/2026

Version 2.0.14126

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:250
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:225
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:260
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:271
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:178

Region	ARN
Canada West (Calgary)	<code>arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension-Arm64:158</code>
Europe (Frankfurt)	<code>arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:234</code>
Europe (Zurich)	<code>arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:174</code>
Europe (Ireland)	<code>arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:237</code>
Europe (London)	<code>arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:190</code>
Europe (Paris)	<code>arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:188</code>
Europe (Stockholm)	<code>arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:222</code>
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:173</code>
Europe (Spain)	<code>arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:171</code>

Region	ARN
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:177
Asia Pacific (Taipei)	arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension-Arm64:113
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:228
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:171
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:221
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:177
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:183
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:206
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:252

Region	ARN
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:189
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:169
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension-Arm64:132
Asia Pacific (New Zealand)	arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension-Arm64:77
Asia Pacific (Thailand)	arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension-Arm64:129
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:184
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:167
Mexico (Central)	arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension-Arm64:137
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:211

Region	ARN
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension-Arm64:158
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension-Arm64:159
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension-Arm64:161
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension-Arm64:160

Date replaced by newer extension: 03/25/2026

Version 2.0.11962

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:229
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:204
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:236

Region	ARN
US West (Oregon)	<code>arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:250</code>
Canada (Central)	<code>arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:159</code>
Canada West (Calgary)	<code>arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension-Arm64:137</code>
Europe (Frankfurt)	<code>arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:213</code>
Europe (Zurich)	<code>arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:153</code>
Europe (Ireland)	<code>arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:216</code>
Europe (London)	<code>arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:169</code>
Europe (Paris)	<code>arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:167</code>
Europe (Stockholm)	<code>arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:201</code>

Region	ARN
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:154</code>
Europe (Spain)	<code>arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:150</code>
Asia Pacific (Hong Kong)	<code>arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:156</code>
Asia Pacific (Taipei)	<code>arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension-Arm64:92</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:198</code>
Asia Pacific (Seoul)	<code>arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:156</code>
Asia Pacific (Osaka)	<code>arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:162</code>
Asia Pacific (Singapore)	<code>arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:185</code>
Asia Pacific (Sydney)	<code>arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:231</code>

Region	ARN
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:168
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:148
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension-Arm64:111
Asia Pacific (New Zealand)	arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension-Arm64:48
Asia Pacific (Thailand)	arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension-Arm64:108
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:206
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:150
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:190
Mexico (Central)	arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension-Arm64:114

Region	ARN
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:162
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension-Arm64:162
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension-Arm64:172
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:151
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension-Arm64:141
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension-Arm64:143
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension-Arm64:130
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension-Arm64:128

Date replaced by newer extension: 02/17/2026

Version 2.0.8693

Region	ARN
US East (N. Virginia)	<code>arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:212</code>
US East (Ohio)	<code>arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:187</code>
US West (N. California)	<code>arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:225</code>
US West (Oregon)	<code>arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:237</code>
Canada (Central)	<code>arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:148</code>
Canada West (Calgary)	<code>arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension-Arm64:120</code>
Europe (Frankfurt)	<code>arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:204</code>
Europe (Zurich)	<code>arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:136</code>
Europe (Ireland)	<code>arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:199</code>

Region	ARN
Europe (London)	<code>arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:159</code>
Europe (Paris)	<code>arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:154</code>
Europe (Stockholm)	<code>arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:192</code>
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:143</code>
Europe (Spain)	<code>arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:137</code>
Asia Pacific (Hong Kong)	<code>arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:145</code>
Asia Pacific (Taipei)	<code>arn:aws:lambda:ap-east-2:730335625313:layer:AWS-AppConfig-Extension-Arm64:74</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:181</code>
Asia Pacific (Seoul)	<code>arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:147</code>

Region	ARN
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:149
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:176
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:215
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:159
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:137
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension-Arm64:102
Asia Pacific (New Zealand)	arn:aws:lambda:ap-southeast-6:381491832265:layer:AWS-AppConfig-Extension-Arm64:31
Asia Pacific (Thailand)	arn:aws:lambda:ap-southeast-7:851725616657:layer:AWS-AppConfig-Extension-Arm64:97
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:190

Region	ARN
Asia Pacific (Hyderabad)	<code>arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:137</code>
South America (São Paulo)	<code>arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:176</code>
Mexico (Central)	<code>arn:aws:lambda:mx-central-1:891376990304:layer:AWS-AppConfig-Extension-Arm64:97</code>
Africa (Cape Town)	<code>arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:153</code>
Middle East (UAE)	<code>arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension-Arm64:151</code>
Middle East (Bahrain)	<code>arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension-Arm64:155</code>
Israel (Tel Aviv)	<code>arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:138</code>
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension-Arm64:127</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension-Arm64:125</code>

Region	ARN
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension-Arm64:130
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension-Arm64:128

Date replaced by newer extension: 11/20/2025

Version 2.0.2037

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:140
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:114
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:135
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:164
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:72

Region	ARN
Canada West (Calgary)	<code>arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension-Arm64:47</code>
Europe (Frankfurt)	<code>arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:132</code>
Europe (Zurich)	<code>arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:64</code>
Europe (Ireland)	<code>arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:127</code>
Europe (London)	<code>arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:85</code>
Europe (Paris)	<code>arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:81</code>
Europe (Stockholm)	<code>arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:118</code>
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:68</code>
Europe (Spain)	<code>arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:63</code>

Region	ARN
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:70
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:108
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:73
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:74
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:108
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:142
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:87
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:63
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension-Arm64:30

Region	ARN
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:117
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:62
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:103
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:80
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension-Arm64:76
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension-Arm64:82
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:64
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension-Arm64:55
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension-Arm64:53

Region	ARN
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension-Arm64:56
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension-Arm64:55

Date replaced by newer extension: 05/20/2025

Version 2.0.1079

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:107
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:85
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:100
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:132
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:43

Region	ARN
Canada West (Calgary)	<code>arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension-Arm64:18</code>
Europe (Frankfurt)	<code>arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:102</code>
Europe (Zurich)	<code>arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:35</code>
Europe (Ireland)	<code>arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:98</code>
Europe (London)	<code>arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:73</code>
Europe (Paris)	<code>arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:52</code>
Europe (Stockholm)	<code>arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:84</code>
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:39</code>
Europe (Spain)	<code>arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:35</code>

Region	ARN
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:41
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:79
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:44
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:45
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:86
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:108
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:58
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:34
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension-Arm64:1

Region	ARN
Asia Pacific (Mumbai)	<code>arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:88</code>
Asia Pacific (Hyderabad)	<code>arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:33</code>
South America (São Paulo)	<code>arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:67</code>
Africa (Cape Town)	<code>arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:51</code>
Middle East (UAE)	<code>arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension-Arm64:47</code>
Middle East (Bahrain)	<code>arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension-Arm64:53</code>
Israel (Tel Aviv)	<code>arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:35</code>
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension-Arm64:28</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension-Arm64:26</code>

Region	ARN
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension-Arm64:26
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension-Arm64:26

Date replaced by newer extension: 12/12/2024

Version 2.0.678

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:106
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:84
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:98
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:131
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:41

Region	ARN
Canada West (Calgary)	<code>arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension-Arm64:17</code>
Europe (Frankfurt)	<code>arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:101</code>
Europe (Zurich)	<code>arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:33</code>
Europe (Ireland)	<code>arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:97</code>
Europe (London)	<code>arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:72</code>
Europe (Paris)	<code>arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:51</code>
Europe (Stockholm)	<code>arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:83</code>
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:38</code>
Europe (Spain)	<code>arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:33</code>

Region	ARN
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:40
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:78
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:43
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:44
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:84
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:107
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:57
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:33
Asia Pacific (Malaysia)	arn:aws:lambda:ap-southeast-5:631746059939:layer:AWS-AppConfig-Extension-Arm64:1

Region	ARN
Asia Pacific (Mumbai)	<code>arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:87</code>
Asia Pacific (Hyderabad)	<code>arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:32</code>
South America (São Paulo)	<code>arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:66</code>
Africa (Cape Town)	<code>arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:50</code>
Middle East (UAE)	<code>arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension-Arm64:46</code>
Middle East (Bahrain)	<code>arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension-Arm64:52</code>
Israel (Tel Aviv)	<code>arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:33</code>
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension-Arm64:26</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension-Arm64:24</code>

Region	ARN
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension-Arm64:25
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension-Arm64:25

Date replaced by newer extension: 08/08/2024

Version 2.0.678

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:100
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:78
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:90
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:125
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:11

Region	ARN
Canada (Central)	<code>arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:36</code>
Europe (Frankfurt)	<code>arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:95</code>
Europe (Zurich)	<code>arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:28</code>
Europe (Ireland)	<code>arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:91</code>
Europe (London)	<code>arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:66</code>
Europe (Paris)	<code>arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:45</code>
Europe (Stockholm)	<code>arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:77</code>
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:32</code>
Europe (Spain)	<code>arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:28</code>

Region	ARN
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:34
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:72
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:37
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:38
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:79
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:101
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:51
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:27
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:81

Region	ARN
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:26
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:60
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:44
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension-Arm64:40
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension-Arm64:46
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:28
China (Beijing)	arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension-Arm64:21
China (Ningxia)	arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension-Arm64:19
AWS GovCloud (US-East)	arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension-Arm64:19

Region	ARN
AWS GovCloud (US-West)	arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension-Arm64:19

Date replaced by newer extension: 07/23/2024

Version 2.0.501

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:86
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:64
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:72
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:112
Canada West (Calgary)	arn:aws:lambda:ca-west-1:436199621743:layer:AWS-AppConfig-Extension:1
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:21

Region	ARN
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:79
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:11
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:82
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:51
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:30
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:60
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:17
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:11
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:19

Region	ARN
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:57
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:22
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:22
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:64
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:85
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:35
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:11
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:67
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:11

Region	ARN
South America (São Paulo)	<code>arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:43</code>
Africa (Cape Town)	<code>arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:30</code>
Middle East (UAE)	<code>arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension-Arm64:24</code>
Middle East (Bahrain)	<code>arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension-Arm64:31</code>
Israel (Tel Aviv)	<code>arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:11</code>
China (Beijing)	<code>arn:aws-cn:lambda:cn-north-1:615057806174:layer:AWS-AppConfig-Extension-Arm64:7</code>
China (Ningxia)	<code>arn:aws-cn:lambda:cn-northwest-1:615084187847:layer:AWS-AppConfig-Extension-Arm64:5</code>
AWS GovCloud (US-East)	<code>arn:aws-us-gov:lambda:us-gov-east-1:946561847325:layer:AWS-AppConfig-Extension-Arm64:5</code>
AWS GovCloud (US-West)	<code>arn:aws-us-gov:lambda:us-gov-west-1:946746059096:layer:AWS-AppConfig-Extension-Arm64:5</code>

Date replaced by newer extension: 07/01/2024

Version 2.0.358

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:61
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:45
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:18
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:63
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:13
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:49
Europe (Zurich)	arn:aws:lambda:eu-central-2:758369105281:layer:AWS-AppConfig-Extension-Arm64:5
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:63

Region	ARN
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:45
Europe (Paris)	arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:17
Europe (Stockholm)	arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:18
Europe (Milan)	arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:11
Europe (Spain)	arn:aws:lambda:eu-south-2:586093569114:layer:AWS-AppConfig-Extension-Arm64:5
Asia Pacific (Hong Kong)	arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:11
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:51
Asia Pacific (Seoul)	arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:16
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:16

Region	ARN
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:58
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:49
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:16
Asia Pacific (Melbourne)	arn:aws:lambda:ap-southeast-4:307021474294:layer:AWS-AppConfig-Extension-Arm64:5
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:49
Asia Pacific (Hyderabad)	arn:aws:lambda:ap-south-2:489524808438:layer:AWS-AppConfig-Extension-Arm64:5
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:16
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:11
Middle East (UAE)	arn:aws:lambda:me-central-1:662846165436:layer:AWS-AppConfig-Extension-Arm64:5

Region	ARN
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension-Arm64:13
Israel (Tel Aviv)	arn:aws:lambda:il-central-1:895787185223:layer:AWS-AppConfig-Extension-Arm64:5

Date replaced by newer extension: 12/01/2023

Version 2.0.181

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:46
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:33
US West (N. California)	arn:aws:lambda:us-west-1:958113053741:layer:AWS-AppConfig-Extension-Arm64:1
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:48
Canada (Central)	arn:aws:lambda:ca-central-1:039592058896:layer:AWS-AppConfig-Extension-Arm64:1

Region	ARN
Europe (Frankfurt)	<code>arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:36</code>
Europe (Ireland)	<code>arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:48</code>
Europe (London)	<code>arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:33</code>
Europe (Paris)	<code>arn:aws:lambda:eu-west-3:493207061005:layer:AWS-AppConfig-Extension-Arm64:1</code>
Europe (Stockholm)	<code>arn:aws:lambda:eu-north-1:646970417810:layer:AWS-AppConfig-Extension-Arm64:1</code>
Europe (Milan)	<code>arn:aws:lambda:eu-south-1:203683718741:layer:AWS-AppConfig-Extension-Arm64:1</code>
Asia Pacific (Hong Kong)	<code>arn:aws:lambda:ap-east-1:630222743974:layer:AWS-AppConfig-Extension-Arm64:1</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:37</code>
Asia Pacific (Seoul)	<code>arn:aws:lambda:ap-northeast-2:826293736237:layer:AWS-AppConfig-Extension-Arm64:1</code>

Region	ARN
Asia Pacific (Osaka)	arn:aws:lambda:ap-northeast-3:706869817123:layer:AWS-AppConfig-Extension-Arm64:1
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:43
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:36
Asia Pacific (Jakarta)	arn:aws:lambda:ap-southeast-3:418787028745:layer:AWS-AppConfig-Extension-Arm64:1
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:36
South America (São Paulo)	arn:aws:lambda:sa-east-1:000010852771:layer:AWS-AppConfig-Extension-Arm64:1
Africa (Cape Town)	arn:aws:lambda:af-south-1:574348263942:layer:AWS-AppConfig-Extension-Arm64:1
Middle East (Bahrain)	arn:aws:lambda:me-south-1:559955524753:layer:AWS-AppConfig-Extension-Arm64:1

Date replaced by newer extension: 03/30/2023

Version 2.0.165

Region	ARN
US East (N. Virginia)	<code>arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:43</code>
US East (Ohio)	<code>arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:31</code>
US West (Oregon)	<code>arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:45</code>
Europe (Frankfurt)	<code>arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:34</code>
Europe (Ireland)	<code>arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:46</code>
Europe (London)	<code>arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:31</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:35</code>
Asia Pacific (Singapore)	<code>arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:41</code>
Asia Pacific (Sydney)	<code>arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:34</code>

Region	ARN
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:34

Date replaced by newer extension: 02/21/2023

Version 2.0.122

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:15
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:11
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:16
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:13
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:20
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:11

Region	ARN
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:15
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:16
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:13
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:13

Date replaced by newer extension: 08/23/2022

Version 2.0.58

Region	ARN
US East (N. Virginia)	arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:2
US East (Ohio)	arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:2
US West (Oregon)	arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:3

Region	ARN
Europe (Frankfurt)	arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:2
Europe (Ireland)	arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:7
Europe (London)	arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:2
Asia Pacific (Tokyo)	arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:2
Asia Pacific (Singapore)	arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:3
Asia Pacific (Sydney)	arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:2
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:2

Date replaced by newer extension: 04/21/2022

Version 2.0.45

Region	ARN
US East (N. Virginia)	<code>arn:aws:lambda:us-east-1:027255383542:layer:AWS-AppConfig-Extension-Arm64:1</code>
US East (Ohio)	<code>arn:aws:lambda:us-east-2:728743619870:layer:AWS-AppConfig-Extension-Arm64:1</code>
US West (Oregon)	<code>arn:aws:lambda:us-west-2:359756378197:layer:AWS-AppConfig-Extension-Arm64:2</code>
Europe (Frankfurt)	<code>arn:aws:lambda:eu-central-1:066940009817:layer:AWS-AppConfig-Extension-Arm64:1</code>
Europe (Ireland)	<code>arn:aws:lambda:eu-west-1:434848589818:layer:AWS-AppConfig-Extension-Arm64:6</code>
Europe (London)	<code>arn:aws:lambda:eu-west-2:282860088358:layer:AWS-AppConfig-Extension-Arm64:1</code>
Asia Pacific (Tokyo)	<code>arn:aws:lambda:ap-northeast-1:980059726660:layer:AWS-AppConfig-Extension-Arm64:1</code>
Asia Pacific (Singapore)	<code>arn:aws:lambda:ap-southeast-1:421114256042:layer:AWS-AppConfig-Extension-Arm64:2</code>
Asia Pacific (Sydney)	<code>arn:aws:lambda:ap-southeast-2:080788657173:layer:AWS-AppConfig-Extension-Arm64:1</code>

Region	ARN
Asia Pacific (Mumbai)	arn:aws:lambda:ap-south-1:554480029851:layer:AWS-AppConfig-Extension-Arm64:1

Using AWS AppConfig Agent with Amazon EC2 and on-premises machines

You can integrate AWS AppConfig with applications running on your Amazon Elastic Compute Cloud (Amazon EC2) Linux instances by using AWS AppConfig Agent. The agent enhances application processing and management in the following ways:

- The agent calls AWS AppConfig on your behalf by using an AWS Identity and Access Management (IAM) role and managing a local cache of configuration data. By pulling configuration data from the local cache, your application requires fewer code updates to manage configuration data, retrieves configuration data in milliseconds, and isn't affected by network issues that can disrupt calls for such data.*
- The agent offers a native experience for retrieving and resolving AWS AppConfig feature flags.
- Out of the box, the agent provides best practices for caching strategies, polling intervals, and availability of local configuration data while tracking the configuration tokens needed for subsequent service calls.
- While running in the background, the agent periodically polls the AWS AppConfig data plane for configuration data updates. Your application can retrieve the data by connecting to localhost on port 2772 (a customizable default port value) and calling HTTP GET to retrieve the data.

*AWS AppConfig Agent caches data the first time the service retrieves your configuration data. For this reason, the first call to retrieve data is slower than subsequent calls.

Topics

- [Step 1: \(Required\) Creating resources and configuring permissions](#)
- [Step 2: \(Required\) Installing and starting AWS AppConfig Agent on Amazon EC2 instances](#)
- [Step 3: \(Optional, but recommended\) Sending log files to CloudWatch Logs](#)

- [Step 4: \(Optional\) Using environment variables to configure AWS AppConfig Agent for Amazon EC2](#)
- [Step 5: \(Required\) Retrieving configuration data](#)
- [Step 6 \(Optional, but recommended\): Automating updates to AWS AppConfig Agent](#)

Step 1: (Required) Creating resources and configuring permissions

To integrate AWS AppConfig with applications running on your Amazon EC2 instances, you must create AWS AppConfig artifacts and configuration data, including feature flags or freeform configuration data. For more information, see [Creating feature flags and free form configuration data in AWS AppConfig](#).

To retrieve configuration data hosted by AWS AppConfig, your applications must be configured with access to the AWS AppConfig data plane. To give your applications access, update the IAM permissions policy that is assigned to the Amazon EC2 instance role. Specifically, you must add the `appconfig:StartConfigurationSession` and `appconfig:GetLatestConfiguration` actions to the policy. Here is an example:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "appconfig:StartConfigurationSession",
        "appconfig:GetLatestConfiguration"
      ],
      "Resource": "*"
    }
  ]
}
```

For more information about adding permissions to a policy, see [Adding and removing IAM identity permissions](#) in the *IAM User Guide*.

Step 2: (Required) Installing and starting AWS AppConfig Agent on Amazon EC2 instances

AWS AppConfig Agent is hosted in an Amazon Simple Storage Service (Amazon S3) bucket that is managed by AWS. Use the following procedure to install the latest version of the agent on your Linux instance. If your application is distributed across multiple instances, then you must perform this procedure on each instance that hosts the application.

Note

Note the following information:

- AWS AppConfig Agent is available for Linux operating systems running kernel version 4.15 or greater. Debian-based systems, such as Ubuntu, are not supported.
- The agent supports x86_64 and ARM64 architectures.
- For distributed applications, we recommend adding the install and startup commands to the Amazon EC2 user data of your Auto Scaling group. If you do, each instance runs the commands automatically. For more information, see [Run commands on your Linux instance at launch](#) in the *Amazon EC2 User Guide*. Additionally, see [Tutorial: Configure user data to retrieve the target lifecycle state through instance metadata](#) in the *Amazon EC2 Auto Scaling User Guide*.
- The procedures throughout this topic describe how to perform actions like installing the agent by logging into the instance to run the command. You can run the commands from a local client machine and target one or more instances by using Run Command, which is a tool in AWS Systems Manager. For more information, see [AWS Systems Manager Run Command](#) in the *AWS Systems Manager User Guide*.
- AWS AppConfig Agent on Amazon EC2 Linux instances is a systemd service.

To install and start AWS AppConfig Agent on an instance

1. Log into your Linux instance.
2. Open a terminal and run one of the following commands with Administrator permissions:

x86_64

```
sudo yum install https://s3.amazonaws.com/aws-appconfig-downloads/aws-appconfig-agent/linux/x86_64/latest/aws-appconfig-agent.rpm
```

ARM64

```
sudo yum install https://s3.amazonaws.com/aws-appconfig-downloads/aws-appconfig-agent/linux/arm64/latest/aws-appconfig-agent.rpm
```

If you want to install a specific version of AWS AppConfig Agent, replace `latest` in the URL with a specific version number. Here's an example for `x86_64`:

```
sudo yum install https://s3.amazonaws.com/aws-appconfig-downloads/aws-appconfig-agent/linux/x86_64/2.0.2/aws-appconfig-agent.rpm
```

3. Run the following command to start the agent:

```
sudo systemctl start aws-appconfig-agent
```

4. Run the following command to verify the agent is running:

```
sudo systemctl status aws-appconfig-agent
```

If successful, the command returns information like the following:

```
aws-appconfig-agent.service - aws-appconfig-agent
...
Active: active (running) since Mon 2023-07-26 00:00:00 UTC; 0s ago
...
```

Note

To stop the agent, run the following command:

```
sudo systemctl stop aws-appconfig-agent
```

Step 3: (Optional, but recommended) Sending log files to CloudWatch Logs

By default, AWS AppConfig Agent publishes logs to STDERR. Systemd redirects STDOUT and STDERR for all services running on the Linux instance to the systemd journal. You can view and manage log data in the systemd journal if you're running AWS AppConfig Agent on only one or two instances. A better solution, a solution we highly recommend for distributed applications, is to write log files to disk and then use Amazon CloudWatch agent to upload the log data to the AWS cloud. Additionally, you can configure the CloudWatch agent to delete old log files from your instance, which prevents your instance from running out of disk space.

To enable logging to disk, you must set the LOG_PATH environment variable, as described in [Step 4: \(Optional\) Using environment variables to configure AWS AppConfig Agent for Amazon EC2](#).

To get started with the CloudWatch agent, see [Collect metrics and logs from Amazon EC2 instances and on-premises servers with the CloudWatch agent](#) in the *Amazon CloudWatch User Guide*. You can use Quick Setup, a tool in Systems Manager to quickly install the CloudWatch agent. For more information, see [Quick Setup Host Management](#) in the *AWS Systems Manager User Guide*.

Warning

If you choose to write log files to disk without using the CloudWatch agent, you must delete old log files. AWS AppConfig Agent automatically rotates log files every hour. If you don't delete old log files, your instance can run out of disk space.

After you install the CloudWatch agent on your instance, create a CloudWatch agent configuration file. The configuration file instructs CloudWatch agent on how to work with AWS AppConfig Agent log files. For more information about creating a CloudWatch agent configuration file, see [Create the CloudWatch agent configuration file](#).

Add the following logs section to the CloudWatch agent configuration file on the instance and save your changes:

```
"logs": {
  "logs_collected": {
    "files": {
      "collect_list": [
        {
          "file_path": "/path_you_specified_for_logging",
```

```
        "log_group_name": "${YOUR_LOG_GROUP_NAME}/aws-appconfig-agent.log",
        "auto_removal": true
    },
    ...
]
},
...
},
...
}
```

If the value of `auto_removal` is `true`, the CloudWatch agent automatically deletes rotated AWS AppConfig Agent log files.

Step 4: (Optional) Using environment variables to configure AWS AppConfig Agent for Amazon EC2

You can configure AWS AppConfig Agent for Amazon EC2 by using environment variables. To set environment variables for a systemd service, you create a drop-in unit file. The following example shows how to create drop-in unit file to set the AWS AppConfig Agent logging level to `DEBUG`.

Example of how to create a drop-in unit file for environment variables

1. Log into your Linux instance.
2. Open a terminal and run the following command with Administrator permissions. The command creates a configuration directory:

```
sudo mkdir /etc/systemd/system/aws-appconfig-agent.service.d
```

3. Run the following command to create the drop-in unit file. Replace *file_name* with a name for the file. The extension must be `.conf`:

```
sudo touch /etc/systemd/system/aws-appconfig-agent.service.d/file_name.conf
```

4. Enter information in the drop-in unit file. The following example adds a `Service` section that defines an environment variable. The example sets AWS AppConfig Agent log level to `DEBUG`.

```
[Service]
Environment=LOG_LEVEL=DEBUG
```

5. Run the following command to reload the systemd configuration:

```
sudo systemctl daemon-reload
```

6. Run the following command to restart AWS AppConfig Agent:

```
sudo systemctl restart aws-appconfig-agent
```

You can configure AWS AppConfig Agent for Amazon EC2 by specifying the following environment variables in a drop-in unit file.

 **Note**

The following table includes a **Sample values** column. Depending on your monitor resolution, you might need to scroll to the bottom of the table and then scroll to the right to view the column.

Environment variable	Details	Default value	Sample value(s)
ACCESS_TOKEN	This environment variable defines a token that must be provided when requesting configuration data from the agent HTTP server. The value of the token must be set in the HTTP request authorization header with an authorization type of Bearer. Here is an example. GET /applications/my_app/...	None	MyAccessToken

Environment variable	Details	Default value	Sample value(s)
	Host: localhost :2772 Authorization: Bearer <token value>		

Environment variable	Details	Default value	Sample value(s)
BACKUP_DIRECTORY	This environment variable enables AWS AppConfig Agent to save a backup of each configuration it retrieves to the specified directory.  Important Configurations backed up to disk are not encrypted. If your configuration contains sensitive data, AWS AppConfig recommends that you practice the principle of least privilege with your filesystem permissions. For more information, see Security in AWS AppConfig.	None	/path/to/backups

Environment variable	Details	Default value	Sample value(s)
EXPERIMENT_ASSIGNMENT_LOG_DESTINATION	This environment variable enables AWS AppConfig Agent to log experiment assignment records, which capture each occurrence of an experiment treatment being assigned to a given entity at a given time. Possible values: <code>file:<path>/</code> indicates a base directory under which assignment records are written to disk; <code>stderr</code> indicates that assignment records are emitted to stderr. This feature requires AWS AppConfig Agent version 2.0.175233 or later.	None	<code>file:/tmp/aws-appconfig/assignments/</code> <code>stderr</code>
HTTP_PORT	This environment variable specifies the port on which the HTTP server for the agent runs.	2772	2772

Environment variable	Details	Default value	Sample value(s)
HTTP_HOST	The HTTP_HOST variable controls how the AWS AppConfig Agent binds to network interfaces. The binding behavior differs based on the runtime environment to ensure optimal security and accessibility.	ECS, EKS - Default binding: All network interfaces (0.0.0.0) EC2 and on-prem - Default binding: localhost only - IPv4 address: 127.0.0.1:2772 - IPv6 address: [::1]:2772	Custom Configuration Options. You can override the default behavior using these values: - all (binds to all interfaces) - localhost (explicitly binds to localhost interfaces) - Specific IP address (e.g 192.168.1.1) - Custom hostname

Environment variable	Details	Default value	Sample value(s)
LOG_LEVEL	This environment variable specifies the level of detail that the agent logs. Each level includes the current level and all higher levels. The value is case insensitive. From most to least detailed, the log levels are: trace, debug, info, warn, error, fatal, and none. The trace log includes detailed information, including timing information, about the agent.	info	trace debug info warn error fatal none
LOG_PATH	The disk location where logs are written. If not specified, logs are written to stderr.	None	/path/to/logs/agent.log

Environment variable	Details	Default value	Sample value(s)
MANIFEST	This environment variable configures AWS AppConfig Agent to take advantage of additional per-configuration features like multi-account retrievals and save configuration to disk. For more information about these features, see Using a manifest to enable additional retrieval features .	None	When using AWS AppConfig configuration as manifest: <code>MyApp:MyEnvironment:MyManifestConfig</code> . When loading manifest from disk: <code>file:/path/to/manifest.json</code>
MAX_CONNECTIONS	This environment variable configures the maximum number of connections that the agent uses to retrieve configurations from AWS AppConfig.	3	3

Environment variable	Details	Default value	Sample value(s)
POLL_INTERVAL	This environment variable controls how often the agent polls AWS AppConfig for updated configuration data. You can specify a number of seconds for the interval. You can also specify a number with a time unit: s for seconds, m for minutes, and h for hours. If a unit isn't specified, the agent defaults to seconds. For example, 60, 60s, and 1m result in the same poll interval.	45 seconds	45 45s 5m 1h
PREFETCH_LIST	This environment variable specifies the configuration data the agent requests from AWS AppConfig as soon as it starts. Multiple configuration identifiers may be provided in a comma-separated list.	None	MyApp:MyEnv:MyConfig abcd123:efgh456:ijkl789 MyApp:MyEnv:Config1,MyApp:MyEnv:Config2

Environment variable	Details	Default value	Sample value(s)
PRELOAD_BACKUPS	If set to <code>true</code> , AWS AppConfig Agent loads configuration backups found in the <code>BACKUP_DIRECTORY</code> into memory and immediately checks to see if a newer version exists from the service. If set to <code>false</code> , AWS AppConfig Agent only loads the contents from a configuration backup if it cannot retrieve configuration data from the service, for example if there is a problem with your network.	<code>true</code>	<code>true</code> <code>false</code>
PROXY_HEADERS	This environment variable specifies headers that are required by the proxy referenced in the <code>PROXY_URL</code> environment variable. The value is a comma-separated list of headers.	None	header: value h1: v1, h2: v2

Environment variable	Details	Default value	Sample value(s)
PROXY_URL	This environment variable specifies the proxy URL to use for connections from the agent to AWS services, including AWS AppConfig. HTTPS and HTTP URLs are supported.	None	<code>http://localhost:7474</code> <code>https://my-proxy.example.com</code>

Environment variable	Details	Default value	Sample value(s)
REQUEST_TIMEOUT	This environment variable controls the amount of time the agent waits for a response from AWS AppConfig. If the service does not respond, the request fails. If the request is for the initial data retrieval, the agent returns an error to your application. If the timeout occurs during a background check for updated data, the agent logs the error and tries again after a short delay. You can specify the number of milliseconds for the timeout. You can also specify a number with a time unit: ms for milliseconds and s for seconds. If a unit isn't specified, the agent defaults	3000ms	3000 3000ms 5s

Environment variable	Details	Default value	Sample value(s)
	to milliseconds. As an example, 5000, 5000ms and 5s result in the same request timeout value.		
ROLE_ARN	This environment variable specifies the Amazon Resource Name (ARN) of an IAM role. AWS AppConfig Agent assumes this role to retrieve configuration data.	None	arn:aws:iam::123456789012:role/MyRole
ROLE_EXTERNAL_ID	This environment variable specifies the external ID to use with the assumed role ARN.	None	MyExternalId
ROLE_SESSION_NAME	This environment variable specifies the session name to be associated with the credentials for the assumed IAM role.	None	AWSAppConfigAgentSession

Environment variable	Details	Default value	Sample value(s)
SERVICE_REGION	This environment variable specifies an alternative AWS Region that AWS AppConfig Agent uses to call the AWS AppConfig service. If left undefined, the agent attempts to determine the current Region. If it can't, the agent fails to start.	None	us-east-1 eu-west-1
WAIT_ON_MANIFEST	This environment variable configures AWS AppConfig Agent to wait until the manifest is processed before completing startup.	true	true false

Step 5: (Required) Retrieving configuration data

You can retrieve configuration data from AWS AppConfig Agent by using an HTTP localhost call. The following examples use `curl` with an HTTP client. You can call the agent using any available HTTP client supported by your application language or available libraries, including an AWS SDK.

To retrieve the full content of any deployed configuration

```
$ curl "http://localhost:2772/applications/application_name/environments/environment_name/configurations/configuration_name"
```

To retrieve a single flag and its attributes from an AWS AppConfig configuration of type Feature Flag

```
$ curl "http://localhost:2772/applications/application_name/
environments/environment_name/configurations/configuration_name?flag=flag_name"
```

To access multiple flags and their attributes from an AWS AppConfig configuration of type Feature Flag

```
$ curl "http://localhost:2772/applications/application_name/
environments/environment_name/configurations/configuration_name?
flag=flag_name_one&flag=flag_name_two"
```

Step 6 (Optional, but recommended): Automating updates to AWS AppConfig Agent

AWS AppConfig Agent is updated periodically. To ensure you are running the latest version of AWS AppConfig Agent on your instances, we recommend that you add the following commands to your Amazon EC2 user data. You can add the commands to the user data on either the instance or the EC2 Auto Scaling group. The script installs and starts the latest version of the agent each time an instance starts or reboots.

```
#!/bin/bash
# install the latest version of the agent
yum install -y https://s3.amazonaws.com/aws-appconfig-downloads/aws-appconfig-agent/
linux/x86_64/latest/aws-appconfig-agent.rpm
# optional: configure the agent
mkdir /etc/systemd/system/aws-appconfig-agent.service.d
echo "${MY_AGENT_CONFIG}" > /etc/systemd/system/aws-appconfig-agent.service.d/
overrides.conf
systemctl daemon-reload
# start the agent
systemctl start aws-appconfig-agent
```

Using AWS AppConfig Agent with Amazon ECS and Amazon EKS

You can integrate AWS AppConfig with Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS) by using AWS AppConfig Agent. The agent functions as a sidecar container running alongside your Amazon ECS and Amazon EKS container

applications. The agent enhances containerized application processing and management in the following ways:

- The agent calls AWS AppConfig on your behalf by using an AWS Identity and Access Management (IAM) role and managing a local cache of configuration data. By pulling configuration data from the local cache, your application requires fewer code updates to manage configuration data, retrieves configuration data in milliseconds, and isn't affected by network issues that can disrupt calls for such data.*
- The agent offers a native experience for retrieving and resolving AWS AppConfig feature flags.
- Out of the box, the agent provides best practices for caching strategies, polling intervals, and local configuration data availability while tracking the configuration tokens needed for subsequent service calls.
- While running in the background, the agent periodically polls the AWS AppConfig data plane for configuration data updates. Your containerized application can retrieve the data by connecting to localhost on port 2772 (a customizable default port value) and calling HTTP GET to retrieve the data.
- AWS AppConfig Agent updates configuration data in your containers without having to restart or recycle those containers.

*AWS AppConfig Agent caches data the first time the service retrieves your configuration data. For this reason, the first call to retrieve data is slower than subsequent calls.

Before you begin

To integrate AWS AppConfig with your container applications, you must create AWS AppConfig artifacts and configuration data, including feature flags or freeform configuration data. For more information, see [Creating feature flags and free form configuration data in AWS AppConfig](#).

To retrieve configuration data hosted by AWS AppConfig, your container applications must be configured with access to the AWS AppConfig data plane. To give your applications access, update the IAM permissions policy that is used by your container service IAM role. Specifically, you must add the `appconfig:StartConfigurationSession` and `appconfig:GetLatestConfiguration` actions to the policy. Container service IAM roles include the following:

- The Amazon ECS task role
- The Amazon EKS node role

- The AWS Fargate pod execution role (if your Amazon EKS containers use Fargate for compute processing)

For more information about adding permissions to a policy, see [Adding and removing IAM identity permissions](#) in the *IAM User Guide*.

Topics

- [Starting the AWS AppConfig agent for Amazon ECS integration](#)
- [Starting the AWS AppConfig agent for Amazon EKS integration](#)
- [\(Optional\) Running AWS AppConfig as a DaemonSet in Amazon EKS](#)
- [\(Optional\) Using environment variables to configure AWS AppConfig Agent for Amazon ECS and Amazon EKS](#)
- [Retrieving configuration data for applications running in Amazon ECS and Amazon EKS](#)

Starting the AWS AppConfig agent for Amazon ECS integration

The AWS AppConfig Agent sidecar container is automatically available in your Amazon ECS environment. To use it, you must start it, as described in the following procedure.

To start Amazon ECS (console)

1. Open the console at <https://console.aws.amazon.com/ecs/v2>.
2. In the navigation pane, choose **Task definitions**.
3. Choose the task definition for your application, and then select the latest revision.
4. Choose **Create new revision**, **Create new revision**.
5. Choose **Add more containers**.
6. For **Name**, enter a unique name for the AWS AppConfig Agent container.
7. For **Image URI**, enter: **public.ecr.aws/aws-appconfig/aws-appconfig-agent:2.x**
8. For **Essential container**, choose **Yes**.
9. In the **Port mappings** section, choose **Add port mapping**.
10. For **Container port**, enter **2772**.

Note

AWS AppConfig Agent runs on port 2772, by default. You can specify a different port.

11. Choose **Create**. Amazon ECS creates a new container revision and displays the details.
12. In the navigation pane, choose **Clusters**, and then choose your application cluster in the list.
13. On the **Services** tab, select the service for your application.
14. Choose **Update**.
15. Under **Deployment configuration**, for **Revision**, choose the latest revision.
16. Choose **Update**. Amazon ECS deploys the latest task definition.
17. After the deployment finishes, you can verify that AWS AppConfig Agent is running on the **Configuration and tasks** tab. On the **Tasks** tab, choose the running task.
18. In the **Containers** section, verify that the AWS AppConfig Agent container is listed.
19. To verify that AWS AppConfig Agent started, choose the **Logs** tab. Locate a statement like the following for the AWS AppConfig Agent container: [appconfig agent] 1970/01/01 00:00:00 INFO serving on localhost:2772

Note

Note the following information.

- AWS AppConfig Agent is a long-running process. As a best practice for Amazon ECS containers, configure health checks for your containers, specifically setting the container dependency to the HEALTHY condition. For more information, see [ContainerDependency](#) in the *Amazon Elastic Container Service API Reference*.
- You can adjust the default behavior of AWS AppConfig Agent by entering or changing environment variables. For information about the available environment variables, see [\(Optional\) Using environment variables to configure AWS AppConfig Agent for Amazon ECS and Amazon EKS](#). For information about how to change environment variables in Amazon ECS, see [Passing environment variables to a container](#) in the *Amazon Elastic Container Service Developer Guide*.

Starting the AWS AppConfig agent for Amazon EKS integration

The AWS AppConfig Agent sidecar container is automatically available in your Amazon EKS environment. To use it, you must start it. The following procedure describes how to use the Amazon EKS `kubectl` command line tool to start the agent.

Note

Before you continue, ensure that your `kubeconfig` file is up to date. For more information about creating or editing a `kubeconfig` file, see [Creating or updating a kubeconfig file for an Amazon EKS cluster](#) in the Amazon EKS User Guide.

To start AWS AppConfig Agent (kubectl command line tool)

1. Open the manifest for your application and verify that your Amazon EKS application is running as a single-container deployment. Contents of the file should look similar to the following.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: my-app
  namespace: my-namespace
  labels:
    app: my-application-label
spec:
  replicas: 1
  selector:
    matchLabels:
      app: my-application-label
  template:
    metadata:
      labels:
        app: my-application-label
    spec:
      containers:
        - name: my-app
          image: my-repo/my-image
          imagePullPolicy: IfNotPresent
```

2. Add the AWS AppConfig Agent container definition details to your deployment manifest.

```
- name: appconfig-agent
  image: public.ecr.aws/aws-appconfig/aws-appconfig-agent:2.x
  ports:
    - name: http
      containerPort: 2772
      protocol: TCP
  env:
    - name: SERVICE_REGION
      value: AWS Region
  imagePullPolicy: IfNotPresent
```

Note

Note the following information.

- AWS AppConfig Agent runs on port 2772, by default. You can specify a different port.
- You can adjust the default behavior of AWS AppConfig Agent by entering environment variables. For more information, see [\(Optional\) Using environment variables to configure AWS AppConfig Agent for Amazon ECS and Amazon EKS](#).
- For *AWS Region*, specify the AWS Region code (for example, us-west-1) where AWS AppConfig Agent retrieves configuration data.

3. Run the following `kubectl` command to apply the changes to your cluster. Replace *my-deployment* with the name of your deployment manifest.

```
kubectl apply -f my-deployment.yaml
```

4. After the deployment finishes, verify that AWS AppConfig Agent is running. Use the following command to view the application pod log file.

```
kubectl logs -n my-namespace -c appconfig-agent my-pod
```

Locate a statement like the following for the AWS AppConfig Agent container: [appconfig agent] 1970/01/01 00:00:00 INFO serving on localhost:2772

Note

You can adjust the default behavior of AWS AppConfig Agent by entering or changing environment variables. For information about the available environment variables, see [\(Optional\) Using environment variables to configure AWS AppConfig Agent for Amazon ECS and Amazon EKS](#).

(Optional) Running AWS AppConfig as a DaemonSet in Amazon EKS

With Amazon EKS, you can run AWS AppConfig Agent as a sidecar, which results in one agent container *per application pod*. Or, if you prefer, you can run AWS AppConfig Agent as a [DaemonSet](#), which results in one agent container *per node in your cluster*.

Note

If you run AWS AppConfig Agent as a DaemonSet, the agent runs in a separate pod, which means you can't access it with calls to `localhost`. You must inject or otherwise discover the agent pod's IP address in order to call it.

To run AWS AppConfig Agent as a DaemonSet, create a manifest file with the following contents. Replace *highlighted* text with details for your application and environment. For *AWS Region*, specify an AWS Region code (for example, `us-west-1`).

```
apiVersion: apps/v1
kind: DaemonSet
metadata:
  name: aws-appconfig-agent
  namespace: my_namespace
  labels:
    app: my_application_label
spec:
  selector:
    matchLabels:
      app: my_application_label
  template:
    metadata:
      labels:
        app: my_application_label
```

```
spec:
  containers:
  - name: aws-appconfig-agent
    image: public.ecr.aws/aws-appconfig/aws-appconfig-agent:2.x
    ports:
    - name: http
      containerPort: 2772
      protocol: TCP
    env:
    - name: SERVICE_REGION
      value: AWS Region
    imagePullPolicy: IfNotPresent
  # set a high priority class to ensure the agent is running on every node
  priorityClassName: system-node-critical
```

Run the following command to apply the AWS AppConfig Agent DaemonSet to your cluster. Replace *aws_appconfig_agent_daemonset* with the name of your DaemonSet manifest.

```
kubectl apply -f aws_appconfig_agent_daemonset.yaml
```

(Optional) Using environment variables to configure AWS AppConfig Agent for Amazon ECS and Amazon EKS

You can configure AWS AppConfig Agent by changing the following environment variables for your agent container.

Note

The following table includes a **Sample values** column. Depending on your monitor resolution, you might need to scroll to the bottom of the table and then scroll to the right to view the column.

Environment variable	Details	Default value	Sample value(s)
ACCESS_TOKEN	This environment variable defines a token that must	None	MyAccessToken

Environment variable	Details	Default value	Sample value(s)
	be provided when requesting configuration data from the agent HTTP server. The value of the token must be set in the HTTP request authorization header with an authorization type of Bearer. Here is an example. <pre>GET /applications/my_app/... Host: localhost:2772 Authorization: Bearer <token value></pre>		

Environment variable	Details	Default value	Sample value(s)
BACKUP_DIRECTORY	This environment variable enables AWS AppConfig Agent to save a backup of each configuration it retrieves to the specified directory.  Important Configurations backed up to disk are not encrypted. If your configuration contains sensitive data, AWS AppConfig recommends that you practice the principle of least privilege with your filesystem permissions. For more information, see Security in AWS AppConfig.	None	/path/to/backups

Environment variable	Details	Default value	Sample value(s)
EXPERIMENT_ASSIGNMENT_LOG_DESTINATION	This environment variable enables AWS AppConfig Agent to log experiment assignment records, which capture each occurrence of an experiment treatment being assigned to a given entity at a given time. Possible values: <code>file:<path>/</code> indicates a base directory under which assignment records are written to disk; <code>stderr</code> indicates that assignment records are emitted to stderr. This feature requires AWS AppConfig Agent version 2.0.175233 or later.	None	<code>file:/tmp/aws-appconfig/assignments/</code> <code>stderr</code>
HTTP_PORT	This environment variable specifies the port on which the HTTP server for the agent runs.	2772	2772

Environment variable	Details	Default value	Sample value(s)
HTTP_HOST	The HTTP_HOST variable controls how the AWS AppConfig Agent binds to network interfaces. The binding behavior differs based on the runtime environment to ensure optimal security and accessibility.	ECS, EKS • Default binding: All network interfaces (0.0.0.0) EC2 and on-prem • Default binding: localhost only • IPv4 address: 127.0.0.1:2772 • IPv6 address: [::1]:2772	Custom Configuration Options. You can override the default behavior using these values: • all (binds to all interfaces) • localhost (explicitly binds to localhost interfaces) • Specific IP address (e.g 192.168.1.1) • Custom hostname

Environment variable	Details	Default value	Sample value(s)
LOG_LEVEL	This environment variable specifies the level of detail that the agent logs. Each level includes the current level and all higher levels. The value is case insensitive. From most to least detailed, the log levels are: trace, debug, info, warn, error, fatal, and none. The trace log includes detailed information, including timing information, about the agent.	info	trace debug info warn error fatal none
LOG_PATH	The disk location where logs are written. If not specified, logs are written to stderr.	None	/path/to/logs/agent.log

Environment variable	Details	Default value	Sample value(s)
MANIFEST	This environment variable configures AWS AppConfig Agent to take advantage of additional per-configuration features like multi-account retrievals and save configuration to disk. For more information about these features, see Using a manifest to enable additional retrieval features .	None	When using AWS AppConfig configuration as manifest: <code>MyApp:MyEnvironment:MyManifestConfig</code> . When loading manifest from disk: <code>file:/path/to/manifest.json</code>
MAX_CONNECTIONS	This environment variable configures the maximum number of connections that the agent uses to retrieve configurations from AWS AppConfig.	3	3

Environment variable	Details	Default value	Sample value(s)
POLL_INTERVAL	This environment variable controls how often the agent polls AWS AppConfig for updated configuration data. You can specify a number of seconds for the interval. You can also specify a number with a time unit: s for seconds, m for minutes, and h for hours. If a unit isn't specified, the agent defaults to seconds. For example, 60, 60s, and 1m result in the same poll interval.	45 seconds	45 45s 5m 1h
PREFETCH_LIST	This environment variable specifies the configuration data the agent requests from AWS AppConfig as soon as it starts. Multiple configuration identifiers may be provided in a comma-separated list.	None	MyApp:MyEnv:MyConfig abcd123:efgh456:ijkl789 MyApp:MyEnv:Config1,MyApp:MyEnv:Config2

Environment variable	Details	Default value	Sample value(s)
PRELOAD_BACKUPS	If set to <code>true</code> , AWS AppConfig Agent loads configuration backups found in the <code>BACKUP_DIRECTORY</code> into memory and immediately checks to see if a newer version exists from the service. If set to <code>false</code> , AWS AppConfig Agent only loads the contents from a configuration backup if it cannot retrieve configuration data from the service, for example if there is a problem with your network.	<code>true</code>	<code>true</code> <code>false</code>
PROXY_HEADERS	This environment variable specifies headers that are required by the proxy referenced in the <code>PROXY_URL</code> environment variable. The value is a comma-separated list of headers.	None	header: value h1: v1, h2: v2

Environment variable	Details	Default value	Sample value(s)
PROXY_URL	This environment variable specifies the proxy URL to use for connections from the agent to AWS services, including AWS AppConfig. HTTPS and HTTP URLs are supported.	None	<code>http://localhost:7474</code> <code>https://my-proxy.example.com</code>

Environment variable	Details	Default value	Sample value(s)
REQUEST_TIMEOUT	This environment variable controls the amount of time the agent waits for a response from AWS AppConfig. If the service does not respond, the request fails. If the request is for the initial data retrieval, the agent returns an error to your application. If the timeout occurs during a background check for updated data, the agent logs the error and tries again after a short delay. You can specify the number of milliseconds for the timeout. You can also specify a number with a time unit: ms for milliseconds and s for seconds. If a unit isn't specified, the agent defaults	3000ms	3000 3000ms 5s

Environment variable	Details	Default value	Sample value(s)
	to milliseconds. As an example, 5000, 5000ms and 5s result in the same request timeout value.		
ROLE_ARN	This environment variable specifies the Amazon Resource Name (ARN) of an IAM role. AWS AppConfig Agent assumes this role to retrieve configuration data.	None	arn:aws:iam::123456789012:role/MyRole
ROLE_EXTERNAL_ID	This environment variable specifies the external ID to use with the assumed role ARN.	None	MyExternalId
ROLE_SESSION_NAME	This environment variable specifies the session name to be associated with the credentials for the assumed IAM role.	None	AWSAppConfigAgentSession

Environment variable	Details	Default value	Sample value(s)
SERVICE_REGION	This environment variable specifies an alternative AWS Region that AWS AppConfig Agent uses to call the AWS AppConfig service. If left undefined, the agent attempts to determine the current Region. If it can't, the agent fails to start.	None	us-east-1 eu-west-1
WAIT_ON_MANIFEST	This environment variable configures AWS AppConfig Agent to wait until the manifest is processed before completing startup.	true	true false

Retrieving configuration data for applications running in Amazon ECS and Amazon EKS

You can retrieve configuration data from AWS AppConfig Agent for applications running in Amazon ECS and Amazon EKS by using an HTTP localhost call. The following examples use `curl` with an HTTP client. You can call the agent using any available HTTP client supported by your application language or available libraries.

Note

To retrieve configuration data if your application uses a forward slash, for example "test-backend/test-service", you will need to use URL encoding.

To retrieve the full content of any deployed configuration

```
$ curl "http://localhost:2772/applications/application_name/
environments/environment_name/configurations/configuration_name"
```

To retrieve a single flag and its attributes from an AWS AppConfig configuration of type Feature Flag

```
$ curl "http://localhost:2772/applications/application_name/
environments/environment_name/configurations/configuration_name?flag=flag_name"
```

To access multiple flags and their attributes from an AWS AppConfig configuration of type Feature Flag

```
$ curl "http://localhost:2772/applications/application_name/
environments/environment_name/configurations/configuration_name?
flag=flag_name_one&flag=flag_name_two"
```

The call returns configuration metadata in HTTP headers, including the configuration version, content type, and configuration version label (if applicable). The body of the agent response contains the configuration content. Here is an example:

```
HTTP/1.1 200 OK
Configuration-Version: 1
Content-Type: application/json
Date: Tue, 18 Feb 2025 20:20:16 GMT
Content-Length: 31

My test config
```

Retrieving basic and multi-variant feature flags

For feature flag configurations (configurations of type `AWS.AppConfig.FeatureFlags`), the AWS AppConfig Agent enables you to retrieve a single flag or a subset of flags in a configuration. Retrieving one or two flags is useful if your use case only needs to use a few flags from the configuration profile. The following examples use cURL.

Note

The ability to call a single feature flag or a subset of flags in a configuration is only available in the AWS AppConfig Agent version 2.0.45 and higher.

You can retrieve AWS AppConfig configuration data from a local HTTP endpoint. To access a specific flag or a list of flags, use the `?flag=FLAG_KEY` query parameter for an AWS AppConfig configuration profile.

To retrieve a single flag and its attributes

```
curl "http://localhost:2772/applications/APPLICATION_NAME/
environments/ENVIRONMENT_NAME/configurations/CONFIGURATION_NAME?flag=FLAG_KEY"
```

To retrieve multiple flags and their attributes

```
curl "http://localhost:2772/applications/APPLICATION_NAME/
environments/ENVIRONMENT_NAME/configurations/CONFIGURATION_NAME?
flag=FLAG_KEY_ONE&flag=FLAG_KEY_TWO"
```

To retrieve feature flag variants based on caller context

The following cURL examples show how to retrieve feature flag variants based on caller context. To best illustrate how to make these calls, this section uses sample calls based on a scenario where a customer created variants similar to the following:

Feature flag variants Info				
Reorder variant up Reorder variant down Edit Create variant				
Name	Enabled value	Attribute values	Rule	
☐ beta testers	☒ ON	-	(or (eq \$userId "Alice") (eq \$userId "123456789012"))	
☐ EU demographic	☒ ON	-	(and (ends_with \$email "@example.com") (eq \$continent "EU"))	
☐ QA testing	☒ ON	-	(and (matches pattern:".*@example\\.com" in:\$email) (contains \$roles "Engineer") (gt \$tenure 5))	
☐ default	☒ ON	-	-	

Variant order is used for evaluation logic.
Variants are evaluated as an ordered list based on the order shown and any specified rules. The variant at the top of the list is evaluated first. If no rules match the supplied context, AWS AppConfig returns the default variant.

Note

To retrieve flag variants, you must use the latest version of AWS AppConfig Agent in your compute environment. For more information, see the following topics that describe how to update, install, or add the agent for each of the following compute environments:

- For Lambda compute environments: [Adding the AWS AppConfig Agent Lambda extension](#)
- For Amazon EC2 compute environments: [Step 2: \(Required\) Installing and starting AWS AppConfig Agent on Amazon EC2 instances](#)
- For Amazon ECS compute environments: [Starting the AWS AppConfig agent for Amazon ECS integration](#)
- For Amazon EKS compute environments: [Starting the AWS AppConfig agent for Amazon EKS integration](#)

To retrieve flag data using the caller context of jane_doe@example.org (who has not opted into the beta program):

```
curl http://localhost:2772/applications/UIRefresh/environments/Production/
configurations/Features \
-H "Context: email=jane_doe@example.org" \
-H "Context: opted_in_to_beta=false"
{
  "ui_refresh": {"_variant":"QA","dark_mode_support":true,"enabled":true}
}
```

To retrieve flag data using the caller context of jane_doe@example.org (who *has* opted into the beta program):

```
curl http://localhost:2772/applications/UIRefresh/environments/Production/
configurations/Features \
-H "Context: email=jane_doe@example.org" \
-H "Context: opted_in_to_beta=true"
{
  "ui_refresh": {"_variant":"QA","dark_mode_support":true,"enabled":true}
}
```

To retrieve flag data using the caller context of jane_doe@qa-testers.example.org (who is a quality assurance tester at Example Organization):

```
curl http://localhost:2772/applications/UIRefresh/environments/Production/
configurations/Features \
-H "Context: email=jane_doe@qa-testers.example.org"
{
  "ui_refresh": {"_variant":"QA","dark_mode_support":true,"enabled":true}
}
```

To retrieve flag data without caller context (which returns the Default variant)

```
curl http://localhost:2772/applications/UIRefresh/environments/Production/
configurations/Features
{
  "ui_refresh": {"_variant":"Default Variant","enabled":false}
}
```

To retrieve flag data for a traffic splitting scenario to determine if 1 out of 10 random callers receive the 'sample population' variant

```
for i in {0..9} do ; \
curl http://localhost:2772/applications/UIRefresh/environments/Production/
configurations/Features \
-H "Context: email=$i@example.org"
{
  "ui_refresh": {"_variant":"Default Variant","enabled":false}
}
{
  "ui_refresh": {"_variant":"Default Variant","enabled":false}
}
```

```

{
  "ui_refresh": {"_variant": "Default Variant", "enabled": false}
}
{
  "ui_refresh": {"_variant": "Default Variant", "enabled": false}
}
{
  "ui_refresh": {"_variant": "Sample
Population", "dark_mode_support": false, "enabled": true}
}
{
  "ui_refresh": {"_variant": "Default Variant", "enabled": false}
}
{
  "ui_refresh": {"_variant": "Default Variant", "enabled": false}
}
{
  "ui_refresh": {"_variant": "Default Variant", "enabled": false}
}
{
  "ui_refresh": {"_variant": "Default Variant", "enabled": false}
}
{
  "ui_refresh": {"_variant": "Default Variant", "enabled": false}
}
{
  "ui_refresh": {"_variant": "Default Variant", "enabled": false}
}
{
  "ui_refresh": {"_variant": "Default Variant", "enabled": false}
}
}

```

Retrieving experiment treatments

When an experiment runs on a feature flag, you retrieve the assigned treatment the same way you retrieve any feature flag from AWS AppConfig Agent. This requires two additions: you identify the entity, and you provide the caller context that your experiment audience rule evaluates.

1. Retrieve the flag from AWS AppConfig Agent.
2. Specify the exact flag or flags for which you want treatments by using the `?flag=FLAG_KEY` query parameter.
3. Provide an entity identifier, such as an end-user ID, in the `Entity-Id` header. The agent returns the same treatment for the same entity for the life of the run.
4. Pass the request context that your experiment audience rule evaluates in the `Context` header.

For example, to discover which treatment `user123` is assigned to for an experiment running on the `inference-settings` flag, make this request:

```
curl "http://localhost:2772/applications/MyApp/environments/MyEnv/configurations/MyFlags?flag=inference-settings" \
-H "Entity-Id: user123"
```

If `user123` is assigned to treatment `t1`, the agent returns a response like the following:

```
{
  "_variant": "__t1__",
  "enabled": true
}
```

The `_variant` field identifies the flag variant assigned to the entity. When you start an experiment, AWS AppConfig maps each treatment to a flag variant that is named for the assigned treatment (`__c__` for the control and `__t1__`, `__t2__`, and so on for treatments). If the treatment defines attribute values, they also appear in the response.

Note

If the entity is not part of the experiment audience, or is part of the audience but is not currently exposed to the experiment (based on the experiment exposure percentage), the agent does not return an experiment treatment. Instead, it falls through to the flag value currently deployed to the environment.

Using a manifest to enable additional retrieval features

AWS AppConfig Agent offers the following additional features to help you retrieve configurations for your applications.

- [Configuring AWS AppConfig Agent to retrieve configurations from multiple accounts](#): Use AWS AppConfig Agent from a primary or *retrieval* AWS account to retrieve configuration data from multiple vendor accounts.
- [Configuring AWS AppConfig Agent to write configuration copies to disk](#): Use AWS AppConfig Agent to write configuration data to disk. This feature enables customers with applications that read configuration data from disk to integrate with AWS AppConfig.

Understanding agent manifests

To enable these AWS AppConfig Agent features, you create a manifest. A manifest is a set of configuration data that you provide to control actions the agent can perform. A manifest is written in JSON. It contains a set of top-level keys that correspond to different configurations you've deployed using AWS AppConfig.

A manifest can include multiple configurations. Furthermore, each configuration in the manifest can identify one or more agent features to use for the specified configuration. The content of the manifest uses the following format:

```
{
  "application_name:environment_name:configuration_name": {
    "agent_feature_to_enable_1": {
      "feature-setting-key": "feature-setting-value"
    },
    "agent_feature_to_enable_2": {
      "feature-setting-key": "feature-setting-value"
    }
  }
}
```

Here is example JSON for a manifest with two configurations. The first configuration (*MyApp*) doesn't use any AWS AppConfig Agent features. The second configuration (*My2ndApp*) uses the *write configuration copy to disk* and the *multi-account retrieval* features:

```
{
  "MyApp:Test:MyAllowListConfiguration": {},
  "My2ndApp:Beta:MyEnableMobilePaymentsFeatureFlagConfiguration": {
    "credentials": {
      "roleArn": "arn:aws:us-west-1:iam::123456789012:role/MyTestRole",
      "roleExternalId": "00b148e2-4ea4-46a1-ab0f-c422b54d0aac",
      "roleSessionName": "AwsAppConfigAgent",
      "credentialsDuration": "2h"
    },
    "writeTo": {
      "path": "/tmp/aws-appconfig/my-2nd-app/beta/my-enable-payments-feature-flag-configuration.json"
    }
  }
}
```

```
}
```

How to supply an agent manifest

You can store the manifest as a file in a location where AWS AppConfig Agent can read it. Or, you can store the manifest as an AWS AppConfig configuration and point the agent to it. To supply an agent manifest, you must set a MANIFEST environment variable with one of the following values:

Manifest location	Environment variable value	Use case
File	file:/path/to/agent-manifest.json	Use this method if your manifest won't change often.
AWS AppConfig configuration	<i>application-name:environment-name:configuration-name</i>	Use this method for dynamic updates. You can update and deploy a manifest stored in AWS AppConfig as a configuration in the same ways you store other AWS AppConfig configurations.
Environment variable	Manifest content (JSON)	Use this method if your manifest won't change often. This method is useful in container environments where it's easier to set an environment variable than it is to expose a file.

For more information about setting variables for AWS AppConfig Agent, see the relevant topic for your use case:

- [Configuring the AWS AppConfig Agent Lambda extension](#)
- [Using AWS AppConfig Agent with Amazon EC2](#)
- [Using AWS AppConfig Agent with Amazon ECS and Amazon EKS](#)

Configuring AWS AppConfig Agent to retrieve configurations from multiple accounts

You can configure AWS AppConfig Agent to retrieve configurations from multiple AWS accounts by entering credential overrides in the AWS AppConfig Agent manifest. *Credential overrides* include the Amazon Resource Name (ARN) of an AWS Identity and Access Management (IAM) role, a role ID, a session name, and a duration for how long the agent can assume the role.

You enter these details in a "credentials" section in the manifest. The "credentials" section uses the following format:


```
{
  "application_name:environment_name:configuration_name": {
    "credentials": {
      "roleArn": "arn:partition:iam::account_ID:role/roleName",
      "roleExternalId": "string",
      "roleSessionName": "string",
      "credentialsDuration": "time_in_hours"
    }
  }
}
```

Here is an example:

```
{
  "My2ndApp:Beta:MyEnableMobilePaymentsFeatureFlagConfiguration": {
    "credentials": {
      "roleArn": "arn:aws:us-west-1:iam::123456789012:role/MyTestRole",
      "roleExternalId": "00b148e2-4ea4-46a1-ab0f-c422b54d0aac",
      "roleSessionName": "AWSAppConfigAgent",
      "credentialsDuration": "2h"
    }
  }
}
```

Before retrieving a configuration, the agent reads the credential details for the configuration from the manifest and then assumes the IAM role specified for that configuration. You can specify a different set of credential overrides for different configurations in a single manifest. The following diagram shows how AWS AppConfig Agent, while running in Account A (the retrieval account), assumes separate roles specified for Accounts B and C (the vendor accounts) and then calls the

[GetLatestConfiguration](#) API operation to retrieve configuration data from AWS AppConfig running in those accounts:



Configure permissions to retrieve configuration data from vendor accounts

AWS AppConfig Agent running in the retrieval account needs permission to retrieve configuration data from the vendor accounts. You give the agent permission by creating an AWS Identity and Access Management (IAM) role in each of the vendor accounts. AWS AppConfig Agent in the retrieval account assumes this role to get data from vendor accounts. Complete the procedures in this section to create an IAM permissions policy, an IAM role, and add agent overrides to the manifest.

Before you begin

Collect the following information before you create a permission policy and a role in IAM.

- The IDs for each AWS account. The *retrieval* account is the account that will call other accounts for configuration data. The *vendor* accounts are the accounts that will vend configuration data to the retrieval account.
- The name of the IAM role used by AWS AppConfig in the retrieval account. Here's a list of the roles used by AWS AppConfig, by default:

- For Amazon Elastic Compute Cloud (Amazon EC2), AWS AppConfig uses the instance role.
- For AWS Lambda, AWS AppConfig uses the Lambda execution role.
- For Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS), AWS AppConfig uses the container role.

If you configured AWS AppConfig Agent to use a different IAM role by specifying the `ROLE_ARN` environment variable, make a note of that name.

Create the permissions policy

Use the following procedure to create a permissions policy using the IAM console. Complete the procedure in each AWS account that will vend configuration data for the retrieval account.

To create an IAM policy

1. Sign in to the AWS Management Console in a vendor account.
2. Open the IAM console at <https://console.aws.amazon.com/iam/>.
3. In the navigation pane, choose **Policies**, and then choose **Create policy**.
4. Choose the **JSON** option.
5. In the **Policy editor**, replace the default JSON with the following policy statement. Update each *example resource placeholder* with vendor account details.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "appconfig:StartConfigurationSession",
        "appconfig:GetLatestConfiguration"
      ],
      "Resource": "arn:aws:appconfig:us-  
east-1:111122223333:application/vendor_application_ID/  
environment/vendor_environment_ID/configuration/vendor_configuration_ID"
    }
  ]
}
```

```
}
```

Here's an example:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "appconfig:StartConfigurationSession",
      "appconfig:GetLatestConfiguration"
    ],
    "Resource": "arn:aws:appconfig:us-east-2:111122223333:application/
abc123/environment/def456/configuration/hij789"
  ]
}
```

6. Choose **Next**.
7. In the **Policy name** field, enter a name.
8. (Optional) For **Add tags**, add one or more tag-key value pairs to organize, track, or control access for this policy.
9. Choose **Create policy**. The system returns you to the **Policies** page.
10. Repeat this procedure in each AWS account that will vend configuration data for the retrieval account.

Create the IAM role

Use the following procedure to create an IAM role using the IAM console. Complete the procedure in each AWS account that will vend configuration data for the retrieval account.

To create an IAM role

1. Sign in to the AWS Management Console in a vendor account.
2. Open the IAM console at <https://console.aws.amazon.com/iam/>.
3. In the navigation pane, choose **Roles**, and then choose **Create policy**.

4. For **Trusted entity type**, choose **AWS account**.
5. In the **AWS account** section, choose **Another AWS account**.
6. In the **Account ID** field, enter the retrieval account ID.
7. (Optional) As a security best practice for this assume role, choose **Require external ID** and enter a string.
8. Choose **Next**.
9. On the **Add permissions** page, use the **Search** field to locate the policy you created in the previous procedure. Select the check box next to its name.
10. Choose **Next**.
11. For **Role name**, enter a name.
12. (Optional) For **Description**, enter a description.
13. For **Step 1: Select trusted entities**, choose **Edit**. Replace the default JSON trust policy with the following policy. Update each *example resource placeholder* with information from your retrieval account.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS":
          "arn:aws:iam::111122223333:role/appconfig_role_in_retrieval_account"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

14. (Optional) For **Tags**, add one or more tag-key value pairs to organize, track, or control access for this role.
15. Choose **Create role**. The system returns you to the **Roles** page.
16. Search for the role you just created. Choose it. In the **ARN** section, copy the ARN. You'll specify this information in the next procedure.

Add credential overrides to the manifest

After you create the IAM role in your vendor account, update the manifest in the retrieval account. Specifically, add the credentials block and the IAM role ARN for retrieving configuration data from the vendor account. Here is the JSON format:

```
{
  "vendor_application_name:vendor_environment_name:vendor_configuration_name": {
    "credentials": {
      "roleArn":
"arn:partition:iam::vendor_account_ID:role/name_of_role_created_in_vendor_account",
      "roleExternalId": "string",
      "roleSessionName": "string",
      "credentialsDuration": "time_in_hours"
    }
  }
}
```

Here is an example:

```
{
  "My2ndApp:Beta:MyEnableMobilePaymentsFeatureFlagConfiguration": {
    "credentials": {
      "roleArn": "arn:aws:us-west-1:iam::123456789012:role/MyTestRole",
      "roleExternalId": "00b148e2-4ea4-46a1-ab0f-c422b54d0aac",
      "roleSessionName": "AwsAppConfigAgent",
      "credentialsDuration": "2h"
    }
  }
}
```

Validate that multi-account retrieval is working

You can validate that that agent is able to retrieve configuration data from multiple accounts by reviewing the AWS AppConfig agent logs. The INFO level log for retrieved initial data for 'YourApplicationName:YourEnvironmentName:YourConfigurationName' is the best indicator for successful retrievals. If retrievals are failing, you should see an ERROR level log indicating the failure reason. Here is an example for a successful retrieval from a vendor account:

```
[appconfig agent] 2023/11/13 11:33:27 INFO AppConfig Agent 2.0.x
[appconfig agent] 2023/11/13 11:33:28 INFO serving on localhost:2772
```

```
[appconfig agent] 2023/11/13 11:33:28 INFO retrieved initial data for  
'MyTestApplication:MyTestEnvironment:MyDenyListConfiguration' in XX.Xms
```

Configuring AWS AppConfig Agent to write configuration copies to disk

You can configure AWS AppConfig Agent to automatically store a copy of a configuration to disk in plain text. This feature enables customers with applications that read configuration data from disk to integrate with AWS AppConfig.

This feature is not designed to be used as a configuration backup feature. AWS AppConfig Agent doesn't read from the configuration files copied to disk. If you want to back up configurations to disk, see the `BACKUP_DIRECTORY` and `PRELOAD_BACKUP` environment variables for [Using AWS AppConfig Agent with Amazon EC2](#) or [Using AWS AppConfig Agent with Amazon ECS and Amazon EKS](#).

Warning

Note the following important information about this feature:

- Configurations saved to disk are stored in *plain text* and are human readable. Don't enable this feature for configurations that include sensitive data.
- This feature writes to the local disk. Use the principle of least privilege for filesystem permissions. For more information, see [Implement least privilege access](#).

To enable write configuration copy to disk

1. Edit the manifest.
2. Choose the configuration that you want AWS AppConfig to write to disk and add a `writeTo` element. Here is an example:

```
{  
  "application_name:environment_name:configuration_name": {  
    "writeTo": {  
      "path": "path_to_configuration_file"  
    }  
  }  
}
```

Here is an example:

```
{
  "MyTestApp:MyTestEnvironment:MyNewConfiguration": {
    "writeTo": {
      "path": "/tmp/aws-appconfig/mobile-app/beta/enable-mobile-payments"
    }
  }
}
```

3. Save your changes. The configuration.json file will be updated each time new configuration data is deployed.

Validate that write configuration copy to disk is working

You can validate that copies of a configuration are being written to disk by looking by reviewing the AWS AppConfig agent logs. The INFO log entry with the phrasing "INFO wrote configuration '*application:environment:configuration*' to *file_path*" indicates that AWS AppConfig Agent writes configuration copies to disk.

Here is an example:

```
[appconfig agent] 2023/11/13 11:33:27 INFO AppConfig Agent 2.0.x
[appconfig agent] 2023/11/13 11:33:28 INFO serving on localhost:2772
[appconfig agent] 2023/11/13 11:33:28 INFO retrieved initial data for
'MobileApp:Beta:EnableMobilePayments' in XX.Xms
[appconfig agent] 2023/11/13 17:05:49 INFO wrote configuration
'MobileApp:Beta:EnableMobilePayments' to /tmp/configs/your-app/your-env/your-
config.json
```

Generating a client using the OpenAPI specification

You can use the following YAML specification for OpenAPI to create an SDK using a tool like [OpenAPI Generator](#). You can update this specification to include hardcoded values for Application, Environment, or Configuration. You can also add additional paths (if you have multiple configuration types) and include configuration schemas to generate configuration-specific typed models for your SDK clients. For more information about OpenAPI (which is also known as Swagger), see the [OpenAPI specification](#).

```
openapi: 3.0.0
```

```
info:
  version: 1.0.0
  title: AWS AppConfig Agent API
  description: An API model for AWS AppConfig Agent.
servers:
  - url: http://localhost:{port}/
    variables:
      port:
        default:
          '2772'
paths:
  /applications/{Application}/environments/{Environment}/configurations/
  {Configuration}:
    get:
      operationId: getConfiguration
      tags:
        - configuration
      parameters:
        - in: path
          name: Application
          description: The application for the configuration to get. Specify either the
application name or the application ID.
          required: true
          schema:
            type: string
        - in: path
          name: Environment
          description: The environment for the configuration to get. Specify either the
environment name or the environment ID.
          required: true
          schema:
            type: string
        - in: path
          name: Configuration
          description: The configuration to get. Specify either the configuration name
or the configuration ID.
          required: true
          schema:
            type: string
        - in: query
          name: flag
          description: The key(s) of the feature flag(s) to retrieve. If not provided,
all flags are returned.
          required: false
```

```

    schema:
      type: array
      items:
        type: string
  - in: header
    name: context
    description: Request context used to evaluate multi-variant feature flags.
    required: false
    schema:
      type: array
      items:
        type: string
        pattern: '^\\w+=\\w+$'
responses:
  200:
    headers:
      ConfigurationVersion:
        schema:
          type: string
    content:
      application/octet-stream:
        schema:
          type: string
          format: binary
    description: successful config retrieval
  400:
    description: BadRequestException
    content:
      application/text:
        schema:
          $ref: '#/components/schemas/Error'
  404:
    description: ResourceNotFoundException
    content:
      application/text:
        schema:
          $ref: '#/components/schemas/Error'
  500:
    description: InternalServerErrorException
    content:
      application/text:
        schema:
          $ref: '#/components/schemas/Error'
  502:

```

```
description: BadGatewayException
content:
  application/text:
    schema:
      $ref: '#/components/schemas/Error'
504:
description: GatewayTimeoutException
content:
  application/text:
    schema:
      $ref: '#/components/schemas/Error'

components:
  schemas:
    Error:
      type: string
      description: The response error
```

Working with AWS AppConfig Agent local development mode

AWS AppConfig Agent supports a *local development mode*. If you enable local development mode, the agent reads configuration data from a specified directory on disk. It doesn't retrieve configuration data from AWS AppConfig. You can simulate configuration deployments by updating files in the specified directory. We recommend local development mode for the following use cases:

- Test different configuration versions before deploying them using AWS AppConfig.
- Test different configuration options for a new feature before committing changes to your code repository.
- Test different configuration scenarios to verify they work as expected.

Warning

Don't use local development mode in production environments. This mode doesn't support important AWS AppConfig safety features like deployment validation and automated rollbacks.

Use the following procedure to configure AWS AppConfig Agent for local development mode.

To configure AWS AppConfig Agent for local development mode

1. Install the agent using the method described for your compute environment. AWS AppConfig Agent works with the following AWS services:
 - [AWS Lambda](#)
 - [Amazon EC2](#)
 - [Amazon ECS and Amazon EKS](#)
2. If the agent is running, stop it.
3. Add LOCAL_DEVELOPMENT_DIRECTORY to the list of environment variables. Specify a directory on the filesystem that provides the agent with read permissions. For example, /tmp/local_configs.
4. Create a file in the directory. The file name must use the following format:

```
application_name:environment_name:configuration_profile_name
```

Here is an example:

```
Mobile:Development:EnableMobilePaymentsFeatureFlagConfiguration
```

Note

- To view feature flag samples you can add to a file in your LOCAL_DEVELOPMENT_DIRECTORY directory, see [Feature flag samples for AWS AppConfig Agent local development mode](#).
- (Optional) You can control the content type the agent returns for your configuration data based on the extension you give the file. For example, if you name the file with a .json extension, the agent returns a content type of application/json when your application requests it. If you omit the extension, the agent uses application/octet-stream for the content type. If you need precise control, you can provide an extension in the format *.type%subtype*. The agent will return a content type of .type/subtype.

5. Run the following command to restart the agent and request the configuration data.

```
curl http://localhost:2772/applications/application_name/
environments/environment_name/configurations/configuration_name
```

The agent checks for changes to the local file at the poll interval specified for the agent. If the poll interval isn't specified, the agent uses the default interval of 45 seconds. This check at the poll interval ensures that the agent behaves the same in a local development environment as it does when configured to interact with the AWS AppConfig service.

Note

To deploy a new version of a local development configuration file, update the file with new data.

Feature flag samples for AWS AppConfig Agent local development mode

This section includes feature flag samples you can use with AWS AppConfig Agent in local development mode. Local development mode expects feature flag data in the data's retrieval-time format. *Retrieval-time format* is the format returned when the flag is retrieved from the [GetLatestConfiguration](#) API, which only contains the flag's value. Retrieval-time format doesn't include a flag's complete definition (as passed to the [CreateHostedConfigurationVersion](#) API). The complete definition for a flag also contains information such as attribute names and values, constraints, and the flag's enabled state.

Topics

- [Basic feature flag samples](#)
- [Multi-variant feature flag samples](#)

Basic feature flag samples

Use the following basic feature flag samples with AWS AppConfig Agent in local development mode.

Note

If you want the agent to report the content type of your local feature flag data as `application/json` (as it would when retrieving flag data from AWS AppConfig in an environment that isn't local development mode), your local feature flag files must use the `.json` extension. For example, `Local:MyFeatureFlags:SampleB1.json`.

Sample 1: A single flag representing a UI refresh.

```
{
  "ui_refresh": {
    "enabled": true,
    "new_styleguide_colors": true
  }
}
```

Sample 2: Multiple flags representing operational feature flags.

```
{
  "background_worker": {
    "enabled": true,
    "num_threads": 4,
    "queue_name": "MyWorkQueue"
  },
  "emergency_shutoff_switch": {
    "enabled": false
  },
  "logger_settings": {
    "enabled": true,
    "level": "INFO"
  }
}
```

Multi-variant feature flag samples

The retrieval-time format of a feature flag configuration that contains at least one multi-variant feature flag is represented as [Amazon Ion](#) data instead of JSON data. In this format, multi-variant flags are represented as an annotated list, and basic flags are represented as an annotated string. The list elements of a multi-variant flag are either a tuple (a list with a length of two), which

represents a single variant, or a string, which represents the default variant. Within a variant tuple, the first element is an s-expression that represents the variant's rule, and the second element is a string that represents the variant's content.

In order for the agent to properly interpret these files, your local feature flag files must use the following extension: `.application%ion%type=AWS.AppConfig.FeatureFlags`.

For example, `Local:MyFeatureFlags:SampleMV1.application%ion%type=AWS.AppConfig.FeatureFlags`.

Sample 1: A multi-variant flag representing a tiered release of a new feature.

```
'tiered_release'::[
  [
    (or (and (eq $group "Tier1") (split by::$userId pct::1 seed::"2025.01.01")) (and
(eq $group "Tier2") (split by::$userId pct::7 seed::"2025.01.01"))),
    '{"_variant": "ShowFeature", "enabled": true}'
  ],
  '{"_variant": "HideFeature", "enabled": false}'
]
```

Sample 2: Multiple flags representing different UX displays based on the user's ID. The first two flags are multi-variant and the final flag is basic.

```
'colorway'::[
  [
    (contains $userId "beta"),
    '{"_variant": "BetaTesters", "enabled": true, "background": "blue", "foreground":
"red"}',
  ],
  [
    (split by::$userId pct::10),
    '{"_variant": "SplitRollOutRedAndBlue", "enabled": true, "background": "blue",
"foreground": "red"}',
  ],
  '{"_variant": "default", "enabled": true, "background": "green", "foreground":
"green"}'
]

'simple_feature'::[
  [
    (contains $userId "beta"),
    '{"_variant": "BetaTesters", "enabled": true}'
  ]
]
```

```
],  
  '''{"_variant": "default", "enabled": false}'''  
]  
  
'button_color'::'''{"enabled": true, "color": "orange"}'''
```

AWS AppConfig browser and mobile use considerations

Feature flags enable you to update the experience of your web pages and mobile application on the fly, without the overhead, risk, or rigidity of an app store release. Using feature flags, you can gradually release a change to your user base at a time of your choosing. If you encounter an error, you can instantly roll back the change without requiring users to upgrade to a new software version. In short, feature flags provide greater control and flexibility when deploying changes to your application.

The following sections describe important considerations for using AWS AppConfig feature flags with web pages and mobiles devices.

Topics

- [Configuration data and flag retrieval](#)
- [Authentication and Amazon Cognito](#)
- [Caching](#)
- [Segmentation](#)
- [Bandwidth \(mobile use cases\)](#)
- [Additional flag use cases](#)

Configuration data and flag retrieval

For browser and mobile use cases, many customers choose to employ a proxy layer between the web or the mobile application and AWS AppConfig. Doing so decouples your AWS AppConfig call volume from the size of your user base, which reduces costs. It also enables you to leverage the [AWS AppConfig Agent](#), which optimizes flag-retrieval performance and supports features like [multi-variant flags](#). AWS AppConfig recommends using AWS Lambda to create the proxy. Instead of retrieving flags directly from AWS AppConfig, configure the [AWS AppConfig Lambda extension](#) to retrieve your feature flags within a Lambda function. Write the function to accept AWS AppConfig retrieval parameters from the event request and to return the corresponding configuration data in the Lambda response. Expose your proxy to the internet using [Lambda function URLs](#).

After you configure your proxy, consider the frequency in which you retrieve data. Mobile uses cases typically do not require high-frequency polling intervals. Configure AWS AppConfig Agent to refresh data from AWS AppConfig more frequently than your application refreshes from the proxy.

Authentication and Amazon Cognito

Lambda function URLs support [two forms of access control](#), `AWS_IAM` and `NONE`. Use `NONE` if you prefer to implement your own authentication and authorization in your Lambda function. `NONE` is also the recommended option if your use case allows exposing your endpoint to the public and your configuration data doesn't contain sensitive data. For all other use cases, use `AWS_IAM`.

Important

If you expose your endpoint to the internet without authentication, ensure that your configuration data does not leak sensitive data, including personally identifiable information (PII), user IDs, or unreleased feature names.

If you choose to use `AWS_IAM`, you'll need to manage credentials with [Amazon Cognito](#). To get started with Amazon Cognito, you create an identity pool. An identity pool allows you to vend short-term credentials to your application for authenticated or guest users. You will need to add roles in the identity pool that permit users to use the `InvokeFunctionUrl` for your Lambda function. Doing so enables instances of your application to access the credentials necessary to retrieve your configuration data.

When working with Amazon Cognito in your application, consider using [AWS Amplify](#). Amplify simplifies mobile/web application interactions with AWS and provides built-in support for Amazon Cognito.

Caching

When using AWS AppConfig, you should always cache your configuration data locally on the device or in the browser. Caching offers the following advantages:

- Improves performance by reducing latency and battery drain
- Offers stability by eliminating dependencies to network access
- Lowers cost by reducing data retrieval frequency

For mobile use cases, we recommend you implement in-memory and persistent on-device caches. Configure your application to attempt to retrieve the desired configuration from the in-memory cache and fall back to fetching from your proxy, if necessary. Upon successful retrieval from your proxy, update the in-memory cache and then persist the configuration to the device. Use a background process to iterate through the cache and refresh each configuration. When fetching configuration for the first time after application startup, if a retrieval is unsuccessful, defer to the persistent configuration (and use it to seed the in-memory cache).

Segmentation

When using feature flags, you may want to segment the feature flagging experience across your customer base. To do so, supply context to your flag retrieval calls, and configure rules to return different [variants of your feature flags](#) based on the provided context. For example, you may have a feature flag variant for iOS 18.X users, a variant for iOS 17.X users, and a default flag for all other versions of iOS. With variants, you can configure every iOS version of your application to target the same configuration in the same environment, but based on the context supplied in the retrieval call (for example, "version": "iOS18.1"), the devices will receive the appropriate variant of the configuration.

Note

If you are using AWS AppConfig feature flag variants for a mobile use case, you must use the AWS AppConfig Agent and a proxy for retrieving feature flags.

If you choose not to use AWS AppConfig Agent to retrieve feature flags, you can leverage AWS AppConfig [environments](#) for simple, low-cardinality segmentation. An environment is a logical deployment group for your targets. In addition to partitioning your configurations to development, testing, and production environments, you can subdivide your customer base by creating mobile-specific environments, such as device type (tablet vs phone) or OS major versions. With separate environments, you can deploy the same or different sets of configuration data to meet the particular requirements of your customer base.

Bandwidth (mobile use cases)

In general, aim to keep the size of each flag set small. Mobile use cases tend to involve low-bandwidth constraints. Minimizing the size of your data will help you maintain a consistent experience across your user base. Also, consider that because mobile devices often operate

between low- and no-bandwidth environments, on-device caching is critical. Application code that fails gracefully if no configuration data can be retrieved is also critical.

Additional flag use cases

The power of feature flags extends beyond feature-release convenience. Long standing operational flags can be used to improve the operational posture of your application. For example, you can create a performance monitoring toggle that emits additional metrics and debug data during an event. Alternatively, you may want to maintain and adjust your application refresh rates for a segment of your customer base.

Retrieving configuration data without AWS AppConfig Agent

The recommended way to retrieve configuration data from AWS AppConfig is by using the Amazon-developed and managed AWS AppConfig Agent. With the agent, you can cache configuration data locally and asynchronously poll the AWS AppConfig data plane service for updates. This caching/polling process ensures that your configuration data is always available for your application while minimizing latency and cost. If you prefer not to use the agent, you can call public APIs directly from the AWS AppConfig data plane service.

The data plane service uses two API actions, [StartConfigurationSession](#) and [GetLatestConfiguration](#). The data plane service also uses [separate endpoints](#) from the AWS AppConfig control plane.

Note

The data plane service replaces the previous process of retrieving configuration data by using the `GetConfiguration` API action. The `GetConfiguration` API is deprecated.

How it works

Here's how the process of directly calling AWS AppConfig APIs using the data plane service works.

Your application retrieves configuration data by first establishing a configuration session using the [StartConfigurationSession](#) API operation. Your session's client then makes periodic calls to [GetLatestConfiguration](#) to check for and retrieve the latest data available.

When calling `StartConfigurationSession`, your code sends the following information:

- Identifiers (ID or name) of an AWS AppConfig application, environment, and configuration profile that the session tracks.
- (Optional) The minimum amount of time the session's client must wait between calls to `GetLatestConfiguration`.

In response, AWS AppConfig provides an `InitialConfigurationToken` to be given to the session's client and used the first time it calls `GetLatestConfiguration` for that session.

Important

This token should only be used once in your first call to `GetLatestConfiguration`. You *must* use the new token in the `GetLatestConfiguration` response (`NextPollConfigurationToken`) in each subsequent call to `GetLatestConfiguration`. To support long poll use cases, the tokens are valid for up to 24 hours. If a `GetLatestConfiguration` call uses an expired token, the system returns `BadRequestException`.

When calling `GetLatestConfiguration`, your client code sends the most recent `ConfigurationToken` value it has and receives in response:

- `NextPollConfigurationToken`: the `ConfigurationToken` value to use on the next call to `GetLatestConfiguration`.
- `NextPollIntervalInSeconds`: the duration the client should wait before making its next call to `GetLatestConfiguration`.
- The configuration: the latest data intended for the session. This may be empty if the client already has the latest version of the configuration.

Important

Note the following important information.

- The [StartConfigurationSession](#) API should only be called once per application, environment, configuration profile, and client to establish a session with the service. This is typically done in the startup of your application or immediately prior to the first retrieval of a configuration.

- If your configuration is deployed using a `KmsKeyIdentifier`, your request to receive the configuration must include permission to call `kms:Decrypt`. For more information, see [Decrypt](#) in the *AWS Key Management Service API Reference*.
- The API operation previously used to retrieve configuration data, `GetConfiguration`, is deprecated. The `GetConfiguration` API operation does not support encrypted configurations.

(Example) Retrieving a configuration by calling AWS AppConfig APIs

The following AWS CLI example demonstrates how to retrieve configuration data by using the AWS AppConfig Data `StartConfigurationSession` and `GetLatestConfiguration` API operations. The first command starts a configuration session. This call includes the IDs (or names) of the AWS AppConfig application, the environment, and the configuration profile. The API returns an `InitialConfigurationToken` used to fetch your configuration data.

```
aws appconfigdata start-configuration-session \  
  --application-identifier application_name_or_ID \  
  --environment-identifier environment_name_or_ID \  
  --configuration-profile-identifier configuration_profile_name_or_ID
```

The system responds with information in the following format.

```
{  
  "InitialConfigurationToken": initial configuration token  
}
```

After starting a session, use [InitialConfigurationToken](#) to call [GetLatestConfiguration](#) to fetch your configuration data. The configuration data is saved to the `mydata.json` file.

```
aws appconfigdata get-latest-configuration \  
  --configuration-token initial configuration token mydata.json
```

The first call to `GetLatestConfiguration` uses the `ConfigurationToken` obtained from `StartConfigurationSession`. The following information is returned.

```
{  
  "NextPollConfigurationToken" : next configuration token,
```

```
"ContentType" : content type of configuration,  
"NextPollIntervalInSeconds" : 60  
}
```

Subsequent calls to `GetLatestConfiguration` *must* provide `NextPollConfigurationToken` from the previous response.

```
aws appconfigdata get-latest-configuration \  
--configuration-token next configuration token mydata.json
```

Important

Note the following important details about the `GetLatestConfiguration` API operation:

- The `GetLatestConfiguration` response includes a `Configuration` section that shows the configuration data. The `Configuration` section only appears if the system finds new or updated configuration data. If the system doesn't find new or updated configuration data, then the `Configuration` data is empty.
- You receive a new `ConfigurationToken` in every response from `GetLatestConfiguration`.
- We recommend tuning the polling frequency of your `GetLatestConfiguration` API calls based on your budget, the expected frequency of your configuration deployments, and the number of targets for a configuration.

Extending AWS AppConfig workflows using extensions

An extension augments your ability to inject logic or behavior at different points during the AWS AppConfig workflow of creating or deploying a configuration. For example, you can use extensions to perform the following types of tasks (to name a few):

- Send a notification to an Amazon Simple Notification Service (Amazon SNS) topic when a configuration profile is deployed.
- Scrub the contents of a configuration profile for sensitive data before a deployment starts.
- Create or update an Atlassian Jira issue whenever a change is made to a feature flag.
- Merge content from a service or data source into your configuration data when you start a deployment.
- Back up a configuration to an Amazon Simple Storage Service (Amazon S3) bucket whenever a configuration is deployed.

You can associate these types of tasks with AWS AppConfig applications, environments, and configuration profiles.

Contents

- [Understanding AWS AppConfig extensions](#)
- [Working with AWS authored extensions](#)
- [Walkthrough: Creating custom AWS AppConfig extensions](#)

Understanding AWS AppConfig extensions

This topic introduces AWS AppConfig extension concepts and terminology. The information is discussed in the context of each step required to set up and use AWS AppConfig extensions.

Topics

- [Step 1: Determine what you want to do with extensions](#)
- [Step 2: Determine when you want the extension to run](#)
- [Step 3: Create an extension association](#)
- [Step 4: Deploy a configuration and verify the extension actions are performed](#)

Step 1: Determine what you want to do with extensions

Do you want to receive a notification to a webhook that sends messages to Slack anytime an AWS AppConfig deployment completes? Do you want to back up a configuration profile to an Amazon Simple Storage Service (Amazon S3) bucket before a configuration is deployed? Do you want to scrub configuration data for sensitive information before the configuration is deployed? You can use extensions to perform these types of tasks and more. You can create custom extensions or use the AWS authored extensions included with AWS AppConfig.

Note

For most use cases, to create a custom extension, you must create an AWS Lambda function to perform any computation and processing defined in the extension. For more information, see [Walkthrough: Creating custom AWS AppConfig extensions](#).

The following AWS authored extensions can help you quickly integrate configuration deployments with other services. You can use these extensions in the AWS AppConfig console or by calling extension [API actions](#) directly from the AWS CLI, AWS Tools for PowerShell, or the SDK.

Extension	Description
AWS AppConfig deployment events to EventBridge	This extension sends events to the EventBridge default event bus when a configuration is deployed.
AWS AppConfig deployment events to Amazon Simple Notification Service (Amazon SNS)	This extension sends messages to an Amazon SNS topic that you specify when a configuration is deployed.
AWS AppConfig deployment events to Amazon Simple Queue Service (Amazon SQS)	This extension enqueues messages into your Amazon SQS queue when a configuration is deployed.
Integration extension—Atlassian Jira	This extensions allows AWS AppConfig to create and update issues whenever you make changes to a feature flag .

Step 2: Determine when you want the extension to run

An extension defines one or more actions that it performs during an AWS AppConfig workflow. For example, the AWS authored AWS AppConfig deployment events to Amazon SNS extension includes an action to send a notification to an Amazon SNS topic. Each action is invoked either when you interact with AWS AppConfig or when AWS AppConfig is performing a process on your behalf. These are called *action points*. AWS AppConfig extensions support the following action points:

PRE_* action points: Extension actions configured on PRE_* action points are applied after request validation, but before AWS AppConfig performs the activity that corresponds to the action point name. These action invocations are processed at the same time as a request. If more than one request is made, action invocations run sequentially. Also note that PRE_* action points receive and can change the contents of a configuration. PRE_* action points can also respond to an error and prevent an action from happening.

- PRE_CREATE_HOSTED_CONFIGURATION_VERSION
- PRE_START_DEPLOYMENT

ON_* action points: An extension can also run in parallel with an AWS AppConfig workflow by using an ON_* action point. ON_* action points are invoked asynchronously. ON_* action points don't receive the contents of a configuration. If an extension experiences an error during an ON_* action point, the service ignores the error and continues the workflow.

- ON_DEPLOYMENT_START
- ON_DEPLOYMENT_STEP
- ON_DEPLOYMENT_BAKING
- ON_DEPLOYMENT_COMPLETE
- ON_DEPLOYMENT_ROLLED_BACK

AT_* action points: Extension actions configured on AT_* action points are invoked synchronously and in parallel to an AWS AppConfig workflow. If an extension experiences an error during an AT_* action point, the service stops the workflow and rolls back the deployment.

- AT_DEPLOYMENT_TICK

Step 3: Create an extension association

To create an extension, or configure an AWS authored extension, you define the action points that invoke an extension when a specific AWS AppConfig resource is used. For example, you can choose to run the AWS AppConfig deployment events to Amazon SNS extension and receive notifications on an Amazon SNS topic anytime a configuration deployment is started for a specific application. Defining which action points invoke an extension for a specific AWS AppConfig resource is called an *extension association*. An extension association is a specified relationship between an extension and an AWS AppConfig resource, such as an application or a configuration profile.

A single AWS AppConfig application can include multiple environments and configuration profiles. If you associate an extension to an application or an environment, AWS AppConfig invokes the extension for any workflows that relate to the application or environment resources, if applicable.

For example, say you have an AWS AppConfig application called MobileApps that includes a configuration profile called AccessList. And say the MobileApps application includes Beta, Integration, and Production environments. You create an extension association for the AWS authored Amazon SNS notification extension and associate the extension to the MobileApps application. The Amazon SNS notification extension is invoked anytime the configuration is deployed for the application to any of the three environments.

Note

You don't have to create an extension to use AWS authored extensions, but you do have to create an extension association.

Step 4: Deploy a configuration and verify the extension actions are performed

After you create an association, when a hosted configuration is created or a configuration is deployed, AWS AppConfig invokes the extension and performs the specified actions. When an extension is invoked, if the system experiences an error during a PRE - * action point, AWS AppConfig returns information about that error.

Working with AWS authored extensions

AWS AppConfig includes the following AWS authored extensions. These extensions can help you integrate the AWS AppConfig workflow with other services. You can use these extensions in the AWS Management Console or by calling extension [API actions](#) directly from the AWS CLI, AWS Tools for PowerShell, or the SDK.

Extension	Description
AWS AppConfig deployment events to EventBridge	This extension sends events to the EventBridge default event bus when a configuration is deployed.
AWS AppConfig deployment events to Amazon Simple Notification Service (Amazon SNS)	This extension sends messages to an Amazon SNS topic that you specify when a configuration is deployed.
AWS AppConfig deployment events to Amazon Simple Queue Service (Amazon SQS)	This extension enqueues messages into your Amazon SQS queue when a configuration is deployed.
Integration extension—Atlassian Jira	This extension allows AWS AppConfig to create and update issues whenever you make changes to a feature flag .

Using the AWS AppConfig deployment events to Amazon EventBridge extension

The AWS AppConfig deployment events to Amazon EventBridge extension is an AWS authored extension that helps you monitor and act on the AWS AppConfig configuration deployment workflow. The extension sends event notifications to the EventBridge default events bus whenever a configuration is deployed. After you've associated the extension to one of your AWS AppConfig applications, environments, or configuration profiles, AWS AppConfig sends event notifications to the event bus after every configuration deployment start, end, and rollback.

If you want more control over which action points send EventBridge notifications, you can create a custom extension and enter the EventBridge default events bus Amazon Resource Name (ARN) for

the URI field. For information about creating an extension, see [Walkthrough: Creating custom AWS AppConfig extensions](#).

Important

This extension supports only the EventBridge default events bus.

Using the extension

To use the AWS AppConfig deployment events to Amazon EventBridge extension, you first attach the extension to one of your AWS AppConfig resources by creating an extension association. You create the association by using the AWS AppConfig console or the [CreateExtensionAssociation](#) API action. When you create the association, you specify the ARN of an AWS AppConfig application, environment, or configuration profile. If you associate the extension to an application or an environment, an event notification is sent for any configuration profile contained within the specified application or environment.

After you create the association, when a configuration for the specified AWS AppConfig resource is deployed, AWS AppConfig invokes the extension and sends notifications according to the action points specified in the extension.

Note

This extension is invoked by the following action points:

- ON_DEPLOYMENT_START
- ON_DEPLOYMENT_COMPLETE
- ON_DEPLOYMENT_ROLLED_BACK

You can't customize the actions points for this extension. To invoke different action points, you can create your own extension. For more information, see [Walkthrough: Creating custom AWS AppConfig extensions](#).

Use the following procedures to create an AWS AppConfig extension association by using either the AWS Systems Manager console or the AWS CLI.

To create an extension association (console)

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **AWS AppConfig**.
3. On the **Extensions** tab, choose **Add to resource**.
4. In the **Extension resource details** section, for **Resource type**, choose an AWS AppConfig resource type. Depending on the resource you choose, AWS AppConfig prompts you to choose other resources.
5. Choose **Create association to resource**.

Here's a sample event sent to EventBridge when the extension is invoked.

```
{
  "version": "0",
  "id": "c53dbd72-c1a0-2302-9ed6-c076e9128277",
  "detail-type": "On Deployment Complete",
  "source": "aws.appconfig",
  "account": "111122223333",
  "time": "2022-07-09T01:44:15Z",
  "region": "us-east-1",
  "resources": [
    "arn:aws:appconfig:us-east-1:111122223333:extensionassociation/z763ff5"
  ],
  "detail": {
    "InvocationId": "5tfjcig",
    "Parameters": {

    },
    "Type": "OnDeploymentComplete",
    "Application": {
      "Id": "ba8to7",
      "Name": "MyApp"
    },
    "Environment": {
      "Id": "pgil2o7",
      "Name": "MyEnv"
    },
    "ConfigurationProfile": {
      "Id": "ga3tqep",
```

```
    "Name": "MyConfigProfile"
  },
  "DeploymentNumber": 1,
  "ConfigurationVersion": "1"
}
```

Using the AWS AppConfig deployment events to Amazon SNS extension

The AWS AppConfig deployment events to Amazon SNS extension is an AWS authored extension that helps you monitor and act on the AWS AppConfig configuration deployment workflow. The extension publishes messages to an Amazon SNS topic whenever a configuration is deployed. After you associate the extension to one of your AWS AppConfig applications, environments, or configuration profiles, AWS AppConfig publishes a message to the topic after every configuration deployment start, end, and rollback.

If you want more control over which action points send Amazon SNS notifications, you can create a custom extension and enter an Amazon SNS topic Amazon Resource Name (ARN) for the URI field. For information about creating an extension, see [Walkthrough: Creating custom AWS AppConfig extensions](#).

Using the extension

This section describes how to use the AWS AppConfig deployment events to Amazon SNS extension.

Step 1: Configure AWS AppConfig to publish messages to a topic

Add an access control policy to your Amazon SNS topic that gives AWS AppConfig (appconfig.amazonaws.com) permission to publish messages (sns:Publish). To make sure only your own account can send these messages, add an `aws:SourceAccount` condition to the policy and set it to your AWS account ID. This prevents someone in a different AWS account from using the extension to send messages to your topic. To scope access even more tightly, you can use an `aws:SourceArn` condition set to the ARN of the extension association instead. For more information, see [Example cases for Amazon SNS access control](#).

The following example policy lets AWS AppConfig publish to your topic only when the request comes from your own account. Replace `MySNSTopic` and the example account ID with your own values.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowAppConfigPublish",
      "Effect": "Allow",
      "Principal": {
        "Service": "appconfig.amazonaws.com"
      },
      "Action": "SNS:Publish",
      "Resource": "arn:aws:sns:us-east-1:111122223333:MySNSTopic",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "111122223333"
        }
      }
    }
  ]
}
```

Prevent cross-account access

You must include either the `aws:SourceAccount` or `aws:SourceArn` condition in your Amazon SNS topic policy to prevent cross-account access. Without these conditions, another AWS account could configure its AWS AppConfig deployment events to Amazon SNS extension to send deployment events to your topic. Always limit the policy to your own account.

Step 2: Create an extension association

Attach the extension to one of your AWS AppConfig resources by creating an extension association. You create the association by using the AWS AppConfig console or the [CreateExtensionAssociation](#) API action. When you create the association, you specify the ARN of an AWS AppConfig application, environment, or configuration profile. If you associate the extension to an application or an environment, a notification is sent for any configuration profile contained within the specified application or environment. When you create the association, you must enter a value for the `topicArn` parameter that contains the ARN of the Amazon SNS topic you want to use.

After you create the association, when a configuration for the specified AWS AppConfig resource is deployed, AWS AppConfig invokes the extension and sends notifications according to the action points specified in the extension.

Note

This extension is invoked by the following action points:

- ON_DEPLOYMENT_START
- ON_DEPLOYMENT_COMPLETE
- ON_DEPLOYMENT_ROLLED_BACK

You can't customize the actions points for this extension. To invoke different action points, you can create your own extension. For more information, see [Walkthrough: Creating custom AWS AppConfig extensions](#).

Use the following procedures to create an AWS AppConfig extension association by using either the AWS Systems Manager console or the AWS CLI.

To create an extension association (console)

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **AWS AppConfig**.
3. On the **Extensions** tab, choose **Add to resource**.
4. In the **Extension resource details** section, for **Resource type**, choose an AWS AppConfig resource type. Depending on the resource you choose, AWS AppConfig prompts you to choose other resources.
5. Choose **Create association to resource**.

Here's a sample of the message sent to the Amazon SNS topic when the extension is invoked.

```
{
  "Type": "Notification",
  "MessageId": "ae9d702f-9a66-51b3-8586-2b17932a9f28",
  "TopicArn": "arn:aws:sns:us-east-1:111122223333:MySNSTopic",
```

```

    "Message": {
      "InvocationId": "7itcaxp",
      "Parameters": {
        "topicArn": "arn:aws:sns:us-east-1:111122223333:MySNSTopic"
      },
      "Application": {
        "Id": "1a2b3c4d",
        "Name": MyApp
      },
      "Environment": {
        "Id": "1a2b3c4d",
        "Name": MyEnv
      },
      "ConfigurationProfile": {
        "Id": "1a2b3c4d",
        "Name": "MyConfigProfile"
      },
      "Description": null,
      "DeploymentNumber": "3",
      "ConfigurationVersion": "1",
      "Type": "OnDeploymentComplete"
    },
    "Timestamp": "2022-06-30T20:26:52.067Z",
    "SignatureVersion": "1",
    "Signature": "<...>",
    "SigningCertURL": "<...>",
    "UnsubscribeURL": "<...>",
    "MessageAttributes": {
      "MessageType": {
        "Type": "String",
        "Value": "OnDeploymentStart"
      }
    }
  }
}

```

Using the AWS AppConfig deployment events to Amazon SQS extension

The AWS AppConfig deployment events to Amazon SQS extension is an AWS authored extension that helps you monitor and act on the AWS AppConfig configuration deployment workflow. The extension enqueues messages into your Amazon Simple Queue Service (Amazon SQS) queue whenever a configuration is deployed. After you associate the extension to one of your

AWS AppConfig applications, environments, or configuration profiles, AWS AppConfig enqueues a message into the queue after every configuration deployment start, end, and rollback.

If you want more control over which action points send Amazon SQS notifications, you can create a custom extension and enter an Amazon SQS queue Amazon Resource Name (ARN) for the URI field. For information about creating an extension, see [Walkthrough: Creating custom AWS AppConfig extensions](#).

Using the extension

This section describes how to use the AWS AppConfig deployment events to Amazon SQS extension.

Step 1: Configure AWS AppConfig to enqueue messages

Add an Amazon SQS policy to your Amazon SQS queue that gives AWS AppConfig (appconfig.amazonaws.com) permission to send messages (sqs:SendMessage). To make sure only your own account can send these messages, add an `aws:SourceAccount` condition to the policy and set it to your AWS account ID. This prevents someone in a different AWS account from using the extension to send messages to your queue. To scope access even more tightly, you can use an `aws:SourceArn` condition set to the ARN of the extension association instead. For more information, see [Basic examples of Amazon SQS policies](#).

The following example policy lets AWS AppConfig send messages to your queue only when the request comes from your own account. Replace `MySQSQueue` and the example account ID with your own values.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowAppConfigSendMessage",
      "Effect": "Allow",
      "Principal": {
        "Service": "appconfig.amazonaws.com"
      },
      "Action": "SQS:SendMessage",
      "Resource": "arn:aws:sqs:us-east-1:111122223333:MySQSQueue",
      "Condition": {
        "StringEquals": {
```

```
    "aws:SourceAccount": "111122223333"
  }
}
]
```

Prevent cross-account access

You must include either the `aws:SourceAccount` or `aws:SourceArn` condition in your Amazon SQS queue policy to prevent cross-account access. Without these conditions, another AWS account could configure its AWS AppConfig deployment events to Amazon SQS extension to send deployment events to your queue. Always limit the policy to your own account.

Step 2: Create an extension association

Attach the extension to one of your AWS AppConfig resources by creating an extension association. You create the association by using the AWS AppConfig console or the [CreateExtensionAssociation](#) API action. When you create the association, you specify the ARN of an AWS AppConfig application, environment, or configuration profile. If you associate the extension to an application or an environment, a notification is sent for any configuration profile contained within the specified application or environment. When you create the association, you must enter a `Here` parameter that contains the ARN of the Amazon SQS queue you want to use.

After you create the association, when a configuration for the specified AWS AppConfig resource is created or deployed, AWS AppConfig invokes the extension and sends notifications according to the action points specified in the extension.

Note

This extension is invoked by the following action points:

- `ON_DEPLOYMENT_START`
- `ON_DEPLOYMENT_COMPLETE`
- `ON_DEPLOYMENT_ROLLED_BACK`

You can't customize the actions points for this extension. To invoke different action points, you can create your own extension. For more information, see [Walkthrough: Creating custom AWS AppConfig extensions](#).

Use the following procedures to create an AWS AppConfig extension association by using either the AWS Systems Manager console or the AWS CLI.

To create an extension association (console)

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **AWS AppConfig**.
3. On the **Extensions** tab, choose **Add to resource**.
4. In the **Extension resource details** section, for **Resource type**, choose an AWS AppConfig resource type. Depending on the resource you choose, AWS AppConfig prompts you to choose other resources.
5. Choose **Create association to resource**.

Here's an example of the message sent to the Amazon SQS queue when the extension is invoked.

```
{
  "InvocationId": "7itcaxp",
  "Parameters": {
    "queueArn": "arn:aws:sqs:us-east-1:111122223333:MySQSQueue"
  },
  "Application": {
    "Id": "1a2b3c4d",
    "Name": "MyApp"
  },
  "Environment": {
    "Id": "1a2b3c4d",
    "Name": "MyEnv"
  },
  "ConfigurationProfile": {
    "Id": "1a2b3c4d",
    "Name": "MyConfigProfile"
  },
}
```

```
"Description":null,
"DeploymentNumber":"3",
"ConfigurationVersion":"1",
"Type":"OnDeploymentComplete"
}
```

Using the Atlassian Jira extension for AWS AppConfig

By integrating with Atlassian Jira, AWS AppConfig can create and update issues in the Atlassian console whenever you make changes to a [feature flag](#) in your AWS account for the specified AWS Region. Each Jira issue includes the flag name, application ID, configuration profile ID, and flag values. After you update, save, and deploy your flag changes, Jira updates the existing issues with the details of the change.

Note

Jira updates issues whenever you create or update a feature flag. Jira also updates issues when you delete a child-level flag attribute from a parent-level flag. Jira does not record information when you delete a parent-level flag.

To configure integration, you must do the following:

- [Configuring permissions for AWS AppConfig Jira integration](#)
- [Configuring the AWS AppConfig Jira integration application](#)

Configuring permissions for AWS AppConfig Jira integration

When you configure AWS AppConfig integration with Jira, you specify credentials for a user. Specifically, you enter the user's access key ID and secret key in the **AWS AppConfig for Jira** application. This user gives Jira permission to communicate with AWS AppConfig. AWS AppConfig uses these credentials one time to establish an association between AWS AppConfig and Jira. The credentials are not stored. You can remove the association by uninstalling the AWS AppConfig for Jira application.

The user account requires a permission policy that includes the following actions:

- `appconfig:CreateExtensionAssociation`

- `appconfig:GetConfigurationProfile`
- `appconfig:ListApplications`
- `appconfig:ListConfigurationProfiles`
- `appconfig:ListExtensionAssociations`
- `sts:GetCallerIdentity`

Complete the following tasks to create an IAM permission policy and a user for AWS AppConfig and Jira integration:

Tasks

- [Task 1: Create an IAM permission policy for AWS AppConfig and Jira integration](#)
- [Task 2: Create a user for AWS AppConfig and Jira integration](#)

Task 1: Create an IAM permission policy for AWS AppConfig and Jira integration

Use the following procedure to create an IAM permission policy that allows Atlassian Jira to communicate with AWS AppConfig. We recommend that you create a new policy and attach this policy to a new IAM role. Adding the required permission to an existing IAM policy and role goes against the principle of least privilege and is not recommended.

To create an IAM policy for AWS AppConfig and Jira integration

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Policies**, and then choose **Create policy**.
3. On the **Create policy** page, choose the **JSON** tab and replace the default content with the following policy. In the following policy, replace *Region*, *account_ID*, *application_ID*, and *configuration_profile_ID* with information from your AWS AppConfig feature flag environment.

JSON

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

```

        "Action": [
            "appconfig:CreateExtensionAssociation",
            "appconfig:ListExtensionAssociations",
            "appconfig:GetConfigurationProfile"
        ],
        "Resource": [
            "arn:aws:appconfig:us-  
east-1:111122223333:application/application_ID",
            "arn:aws:appconfig:us-  
east-1:111122223333:application/application_ID/  
configurationprofile/configuration_profile_ID"
        ]
    },
    {
        "Effect": "Allow",
        "Action": [
            "appconfig:ListApplications"
        ],
        "Resource": [
            "arn:aws:appconfig:us-east-1:111122223333:*"
        ]
    },
    {
        "Effect": "Allow",
        "Action": [
            "appconfig:ListConfigurationProfiles"
        ],
        "Resource": [
            "arn:aws:appconfig:us-  
east-1:111122223333:application/application_ID"
        ]
    },
    {
        "Effect": "Allow",
        "Action": "sts:GetCallerIdentity",
        "Resource": "*"
    }
]
}

```

4. Choose **Next: Tags**.
5. (Optional) Add one or more tag-key value pairs to organize, track, or control access for this policy, and then choose **Next: Review**.

6. On the **Review policy** page, enter a name in the **Name** box, such as **AppConfigJiraPolicy**, and then enter an optional description.
7. Choose **Create policy**.

Task 2: Create a user for AWS AppConfig and Jira integration

Use the following procedure to create a user for AWS AppConfig and Atlassian Jira integration. After you create the user, you can copy the access key ID and secret key, which you will specify when you complete the integration.

To create a user for AWS AppConfig and Jira integration

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Users**, and then choose **Add users**.
3. In the **User name** field, enter a name, such as **AppConfigJiraUser**.
4. For **Select AWS credential type**, choose **Access key - Programmatic access**.
5. Choose **Next: Permissions**.
6. Under **Set permissions** page, choose **Attach existing policies directly**. Search for and select the check box for the policy that you created in [Task 1: Create an IAM permission policy for AWS AppConfig and Jira integration](#), and then choose **Next: Tags**.
7. On the **Add tags (optional)** page, add one or more tag-key value pairs to organize, track, or control access for this user. Choose **Next: Review**.
8. On the **Review** page, verify the user details.
9. Choose **Create user**. The system displays the user's access key ID and secret key. Either download the .csv file or copy these credentials to a separate location. You will specify these credentials when you configure integration.

Configuring the AWS AppConfig Jira integration application

Use the following procedure to configure required options in the AWS AppConfig for Jira application. After you complete this procedure, Jira creates a new issue for each feature flag in your AWS account for the specified AWS Region. If you make changes to a feature flag in AWS AppConfig, Jira records the details in the existing issues.

 Note

An AWS AppConfig feature flag can include multiple child-level flag attributes. Jira creates one issue for each parent-level feature flag. If you change a child-level flag attribute, you can view the details of that change in the Jira issue for the parent-level flag.

To configure integration

1. Log in to the [Atlassian Marketplace](#).
2. Type **AWS AppConfig** in the search field and press **Enter**.
3. Install the application on your Jira instance.
4. In the Atlassian console, choose **Manage apps**, and then choose **AWS AppConfig for Jira**.
5. Choose **Configure**.
6. Under **Configuration details**, choose **Jira project** and then choose the project that you want to associate with your AWS AppConfig feature flag.
7. Choose **AWS Region**, and then choose the Region where your AWS AppConfig feature flag is located.
8. In the **Application ID** field, enter the name of the AWS AppConfig application that contains your feature flag.
9. In the **Configuration profile ID** field, enter the name of the AWS AppConfig configuration profile for your feature flag.
10. In the **Access key ID** and **Secret key** fields, enter the credentials you copied in [Task 2: Create a user for AWS AppConfig and Jira integration](#). Optionally, you can also specify a session token.
11. Choose **Submit**.
12. In the Atlassian console, choose **Projects**, and then choose the project you selected for AWS AppConfig integration. The **Issues** page displays an issue for each feature flag in the specified AWS account and AWS Region.

Deleting the AWS AppConfig for Jira application and data

If you no longer want to use Jira integration with AWS AppConfig feature flags, you can delete the AWS AppConfig for Jira application in the Atlassian console. Deleting the integration application does the following:

- Deletes the association between your Jira instance and AWS AppConfig
- Deletes your Jira instance details from AWS AppConfig

To delete the AWS AppConfig for Jira application

1. In the Atlassian console, choose **Manage apps**.
2. Choose **AWS AppConfig for Jira**.
3. Choose **Uninstall**.

Walkthrough: Creating custom AWS AppConfig extensions

To create a custom AWS AppConfig extension, complete the following tasks. Each task is described in more detail in later topics.

Note

You can view samples of custom AWS AppConfig extensions on GitHub:

- [Sample extension that prevents deployments with a blocked day moratorium calendar using Systems Manager Change Calendar](#)
- [Sample extension that prevents secrets from leaking into configuration data using git-secrets](#)
- [Sample extension that prevents personally identifiable information \(PII\) from leaking into configuration data using Amazon Comprehend](#)

1. [Create an AWS Lambda function](#)

For most use cases, to create a custom extension, you must create an AWS Lambda function to perform any computation and processing defined in the extension. An exception to this rule is if you create *custom* versions of the [AWS authored notification extensions](#) to add or remove action points. For more details about this exception, see [Step 3: Create a custom AWS AppConfig extension](#).

2. [Configure permissions for your custom extension](#)

To configure permissions for your custom extension, you can do one of the following:

- Create an AWS Identity and Access Management (IAM) service role that includes `InvokeFunction` permissions.
- Create a resource policy by using the Lambda [AddPermission](#) API action.

This walkthrough describes how to create the IAM service role.

3. [Create an extension](#)

You can create an extension by using the AWS AppConfig console or by calling the [CreateExtension](#) API action from the AWS CLI, AWS Tools for PowerShell, or the SDK. The walkthrough uses the console.

4. [Create an extension association](#)

You can create an extension association by using the AWS AppConfig console or by calling the [CreateExtensionAssociation](#) API action from the AWS CLI, AWS Tools for PowerShell, or the SDK. The walkthrough uses the console.

5. Perform an action that invokes the extension

After you create the association, AWS AppConfig invokes the extension when the action points defined by the extension occur for that resource. For example, if you associate an extension that contains a `PRE_CREATE_HOSTED_CONFIGURATION_VERSION` action, the extension is invoked every time you create a new hosted configuration version.

The topics in this section describe each task involved in creating a custom AWS AppConfig extension. Each task is described in the context of a use case where a customer wants to create an extension that automatically backs up a configuration to an Amazon Simple Storage Service (Amazon S3) bucket. The extension runs whenever a hosted configuration is created (`PRE_CREATE_HOSTED_CONFIGURATION_VERSION`) or deployed (`PRE_START_DEPLOYMENT`).

Topics

- [Step 1: Create a Lambda function for a custom AWS AppConfig extension](#)
- [Step 2: Configure permissions for a custom AWS AppConfig extension](#)
- [Step 3: Create a custom AWS AppConfig extension](#)
- [Step 4: Create an extension association for a custom AWS AppConfig extension](#)

Step 1: Create a Lambda function for a custom AWS AppConfig extension

For most use-cases, to create a custom extension, you must create an AWS Lambda function to perform any computation and processing defined in the extension. This section includes Lambda function sample code for a custom AWS AppConfig extension. This section also includes payload request and response reference details. For information about creating a Lambda function, see [Getting started with Lambda](#) in the *AWS Lambda Developer Guide*.

Sample code

The following sample code for a Lambda function, when invoked, automatically backs up an AWS AppConfig configuration to an Amazon S3 bucket. The configuration is backed up whenever a new configuration is created or deployed. The sample uses extension parameters so the bucket name doesn't have to be hardcoded in the Lambda function. By using extension parameters, the user can attach the extension to multiple applications and back up configurations to different buckets. The code sample includes comments to further explain the function.

Sample Lambda function for an AWS AppConfig extension

```
from datetime import datetime
import base64
import json

import boto3

def lambda_handler(event, context):
    print(event)

    # Extensions that use the PRE_CREATE_HOSTED_CONFIGURATION_VERSION and
    PRE_START_DEPLOYMENT
    # action points receive the contents of AWS AppConfig configurations in Lambda
    event parameters.
    # Configuration contents are received as a base64-encoded string, which the lambda
    needs to decode
    # in order to get the configuration data as bytes. For other action points, the
    content
    # of the configuration isn't present, so the code below will fail.
    config_data_bytes = base64.b64decode(event["Content"])
```

```

    # You can specify parameters for extensions. The CreateExtension API action lets
    you define
    # which parameters an extension supports. You supply the values for those
    parameters when you
    # create an extension association by calling the CreateExtensionAssociation API
    action.
    # The following code uses a parameter called S3_BUCKET to obtain the value
    specified in the
    # extension association. You can specify this parameter when you create the
    extension
    # later in this walkthrough.
    extension_association_params = event.get('Parameters', {})
    bucket_name = extension_association_params['S3_BUCKET']
    write_backup_to_s3(bucket_name, config_data_bytes)

    # The PRE_CREATE_HOSTED_CONFIGURATION_VERSION and PRE_START_DEPLOYMENT action
    points can
    # modify the contents of a configuration. The following code makes a minor change
    # for the purposes of a demonstration.
    old_config_data_string = config_data_bytes.decode('utf-8')
    new_config_data_string = old_config_data_string.replace('hello', 'hello!')
    new_config_data_bytes = new_config_data_string.encode('utf-8')

    # The lambda initially received the configuration data as a base64-encoded string
    # and must return it in the same format.
    new_config_data_base64string =
    base64.b64encode(new_config_data_bytes).decode('ascii')

    return {
        'statusCode': 200,
        # If you want to modify the contents of the configuration, you must include the
    new contents in the
        # Lambda response. If you don't want to modify the contents, you can omit the
    'Content' field shown here.
        'Content': new_config_data_base64string
    }

def write_backup_to_s3(bucket_name, config_data_bytes):
    s3 = boto3.resource('s3')
    new_object = s3.Object(bucket_name,
    f"config_backup_{datetime.now().isoformat()}.txt")
    new_object.put(Body=config_data_bytes)

```

If you want to use this sample during this walkthrough, save it with the name **MyS3ConfigurationBackUpExtension** and copy the Amazon Resource Name (ARN) for the function. You specify the ARN when you create the AWS Identity and Access Management (IAM) assume role in the next section. You specify the ARN and the name when you create the extension.

Payload reference

This section includes payload request and response reference details for working with custom AWS AppConfig extensions.

Request structure

AtDeploymentTick

```
{
  'InvocationId': 'o2xbtm7',
  'Parameters': {
    'ParameterOne': 'ValueOne',
    'ParameterTwo': 'ValueTwo'
  },
  'Type': 'OnDeploymentStart',
  'Application': {
    'Id': 'abcd123'
  },
  'Environment': {
    'Id': 'efgh456'
  },
  'ConfigurationProfile': {
    'Id': 'ijkl789',
    'Name': 'ConfigurationName'
  },
  'DeploymentNumber': 2,
  'Description': 'Deployment description',
  'ConfigurationVersion': '2',
  'DeploymentState': 'DEPLOYING',
  'PercentageComplete': '0.0'
}
```

Request structure

PreCreateHostedConfigurationVersion

```
{
```

```

    'InvocationId': 'vlns753', // id for specific invocation
    'Parameters': {
      'ParameterOne': 'ValueOne',
      'ParameterTwo': 'ValueTwo'
    },
    'ContentType': 'text/plain',
    'ContentVersion': '2',
    'Content': 'SGVsbG8gZWYdGgh', // Base64 encoded content
    'Application': {
      'Id': 'abcd123',
      'Name': 'ApplicationName'
    },
    'ConfigurationProfile': {
      'Id': 'ijkl789',
      'Name': 'ConfigurationName'
    },
    'Description': '',
    'Type': 'PreCreateHostedConfigurationVersion',
    'PreviousContent': {
      'ContentType': 'text/plain',
      'ContentVersion': '1',
      'Content': 'SGVsbG8gd29ybGQh'
    }
  }
}

```

PreStartDeployment

```

{
  'InvocationId': '765ahdm',
  'Parameters': {
    'ParameterOne': 'ValueOne',
    'ParameterTwo': 'ValueTwo'
  },
  'ContentType': 'text/plain',
  'ContentVersion': '2',
  'Content': 'SGVsbG8gZWYdGgh',
  'Application': {
    'Id': 'abcd123',
    'Name': 'ApplicationName'
  },
  'Environment': {
    'Id': 'ibpnqlq',
    'Name': 'EnvironmentName'
  }
}

```

```
    },
    'ConfigurationProfile': {
      'Id': 'ijkl789',
      'Name': 'ConfigurationName'
    },
    'DeploymentNumber': 2,
    'Description': 'Deployment description',
    'Type': 'PreStartDeployment'
  }
```

Asynchronous events

OnStartDeployment, OnDeploymentStep, OnDeployment

```
{
  'InvocationId': 'o2xbtm7',
  'Parameters': {
    'ParameterOne': 'ValueOne',
    'ParameterTwo': 'ValueTwo'
  },
  'Type': 'OnDeploymentStart',
  'Application': {
    'Id': 'abcd123'
  },
  'Environment': {
    'Id': 'efgh456'
  },
  'ConfigurationProfile': {
    'Id': 'ijkl789',
    'Name': 'ConfigurationName'
  },
  'DeploymentNumber': 2,
  'Description': 'Deployment description',
  'ConfigurationVersion': '2'
}
```

Response structure

The following examples show what your Lambda function returns in response to the request from a custom AWS AppConfig extension.

PRE_ Synchronous events - successful response*

If you want to transform the content, use the following:

```
"Content": "SomeBase64EncodedByteArray"
```

AT_ Synchronous events - successful response*

If you want to control the next steps of a deployment (continue a deployment or roll it back) set **Directive** and **Description** attributes in the response.

```
"Directive": "ROLL_BACK"  
"Description": "Deployment event log description"
```

Directive supports two values: **CONTINUE** or **ROLL_BACK**. Use these enums in your payload response to control the next steps of a deployment.

Synchronous events - successful response

If you want to transform the content, use the following:

```
"Content": "SomeBase64EncodedByteArray"
```

If you don't want to transform the content, return nothing.

Asynchronous events - successful response

Return nothing.

All error events

```
{  
  "Error": "BadRequestError",  
  "Message": "There was malformed stuff in here",  
  "Details": [{  
    "Type": "Malformed",  
    "Name": "S3 pointer",  
    "Reason": "S3 bucket did not exist"  
  }]  
}
```

Step 2: Configure permissions for a custom AWS AppConfig extension

Use the following procedure to create and configure an AWS Identity and Access Management (IAM) service role (or *assume role*). AWS AppConfig uses this role to invoke the Lambda function.

To create an IAM service role and allow AWS AppConfig to assume it

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Roles**, and then choose **Create role**.
3. Under **Select type of trusted entity**, choose **Custom trust policy**.
4. Paste the following JSON policy into the **Custom trust policy** field.

JSON

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

Choose **Next**.

5. On the **Add permissions** page, choose **Create policy**. The **Create policy** page opens in a new tab.
6. Choose the **JSON** tab, and then paste the following permission policy into the editor. The `lambda:InvokeFunction` action is used for `PRE_*` action points. The `lambda:InvokeAsync` action is used for `ON_*` action points. Replace *Your Lambda ARN* with the Amazon Resource Name (ARN) of your Lambda.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "lambda:InvokeFunction",
        "lambda:InvokeAsync"
      ],
      "Resource": "arn:aws:lambda:us-east-1:111122223333:function:func-  
name"
    }
  ]
}
```

7. Choose **Next: Tags**.
8. On the **Add tags (Optional)** page, add one or more key-value pairs and then choose **Next: Review**.
9. On the **Review policy** page enter a name and a description, and then choose **Create policy**.
10. On the browser tab for your custom trust policy, choose the Refresh icon and then search for the permission policy you just created.
11. Select the check box for your permission policy and then choose **Next**.
12. On the **Name, review, and create** page, enter a name in the **Role name** box, and then enter a description.
13. Choose **Create role**. The system returns you to the **Roles** page. Choose **View role** in the banner.
14. Copy the ARN. You specify this ARN when you create the extension.

Step 3: Create a custom AWS AppConfig extension

An extension defines one or more actions that it performs during an AWS AppConfig workflow. For example, the AWS authored AWS AppConfig deployment events to Amazon SNS extension includes an action to send a notification to an Amazon SNS topic. Each action is invoked either

when you interact with AWS AppConfig or when AWS AppConfig is performing a process on your behalf. These are called *action points*. AWS AppConfig extensions support the following action points:

PRE_* action points: Extension actions configured on PRE_* action points are applied after request validation, but before AWS AppConfig performs the activity that corresponds to the action point name. These action invocations are processed at the same time as a request. If more than one request is made, action invocations run sequentially. Also note that PRE_* action points receive and can change the contents of a configuration. PRE_* action points can also respond to an error and prevent an action from happening.

- PRE_CREATE_HOSTED_CONFIGURATION_VERSION
- PRE_START_DEPLOYMENT

ON_* action points: An extension can also run in parallel with an AWS AppConfig workflow by using an ON_* action point. ON_* action points are invoked asynchronously. ON_* action points don't receive the contents of a configuration. If an extension experiences an error during an ON_* action point, the service ignores the error and continues the workflow.

- ON_DEPLOYMENT_START
- ON_DEPLOYMENT_STEP
- ON_DEPLOYMENT_BAKING
- ON_DEPLOYMENT_COMPLETE
- ON_DEPLOYMENT_ROLLED_BACK

AT_* action points: Extension actions configured on AT_* action points are invoked synchronously and in parallel to an AWS AppConfig workflow. If an extension experiences an error during an AT_* action point, the service stops the workflow and rolls back the deployment.

- AT_DEPLOYMENT_TICK

The AT_DEPLOYMENT_TICK action point supports third-party monitoring integration.

AT_DEPLOYMENT_TICK is invoked during configuration deployment processing orchestration. If you use a third-party monitoring solution (for example, Datadog or New Relic), you can create an AWS AppConfig extension that checks for alarms at the AT_DEPLOYMENT_TICK action point and, as a safety guardrail, rolls back the deployment if it triggered an alarm.

If you use a third-party monitoring solution like Datadog or New Relic, you can create an AWS AppConfig extension that checks for alarms at the `AT_DEPLOYMENT_TICK` action point and, as a safety guardrail, rolls back the deployment if it triggered an alarm. For more information, see the following Datadog and New Relic integration examples on GitHub:

- [Datadog](#)
- [New Relic](#)

For more information about AWS AppConfig extensions, see the following topics:

- [Extending AWS AppConfig workflows using extensions](#)
- [Walkthrough: Creating custom AWS AppConfig extensions](#)

Sample extension

The following sample extension defines one action that calls the `PRE_CREATE_HOSTED_CONFIGURATION_VERSION` action point. In the `Uri` field, the action specifies the Amazon Resource Name (ARN) of the `MyS3ConfigurationBackUpExtension` Lambda function created earlier in this walkthrough. The action also specifies the AWS Identity and Access Management (IAM) assume role ARN created earlier in this walkthrough.

Sample AWS AppConfig extension

```
{
  "Name": "MySampleExtension",
  "Description": "A sample extension that backs up configurations to an S3 bucket.",
  "Actions": {
    "PRE_CREATE_HOSTED_CONFIGURATION_VERSION": [
      {
        "Name": "PreCreateHostedConfigVersionActionForS3Backup",
        "Uri": "arn:aws:lambda:aws-  
region:111122223333:function:MyS3ConfigurationBackUpExtension",
        "RoleArn": "arn:aws:iam::111122223333:role/ExtensionsTestRole"
      }
    ]
  },
  "Parameters" : {
    "S3_BUCKET": {
      "Required": false
    }
  }
}
```

```
}  
}
```

Note

To view request syntax and field descriptions when creating an extension, see the [CreateExtension](#) topic in the *AWS AppConfig API Reference*.

To create an extension (console)

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **AWS AppConfig**.
3. On the **Extensions** tab, choose **Create extension**.
4. For **Extension name**, enter a unique name. For the purposes of this walkthrough, enter **MyS3ConfigurationBackUpExtension**. Optionally, enter a description.
5. In the **Actions** section, choose **Add new action**.
6. For **Action name**, enter a unique name. For the purposes of this walkthrough, enter **PreCreateHostedConfigVersionActionForS3Backup**. This name describes the action point used by the action and the extension purpose.
7. In the **Action point** list, choose **PRE_CREATE_HOSTED_CONFIGURATION_VERSION**.
8. For **Uri**, choose **Lambda function** and then choose the function in the **Lambda function** list. If you don't see your function, verify that you are in the same AWS Region where you created the function.
9. For **IAM Role**, choose the role you created earlier in this walkthrough.
10. In the **Extension parameters (optional)** section, choose **Add new parameter**.
11. For **Parameter name**, enter a name. For the purposes of this walkthrough, enter **S3_BUCKET**.
12. Repeat steps 5–11 to create a second action for the **PRE_START_DEPLOYMENT** action point.
13. Choose **Create extension**.

Customizing AWS authored notification extensions

You don't have to create a Lambda or an extension to use [AWS authored notification extensions](#). You can simply create an extension association and then perform an operation that calls one of

the supported action points. By default, the AWS authored notification extensions support the following actions points:

- ON_DEPLOYMENT_START
- ON_DEPLOYMENT_COMPLETE
- ON_DEPLOYMENT_ROLLED_BACK

If you create custom versions of the AWS AppConfig deployment events to Amazon SNS extension and AWS AppConfig deployment events to Amazon SQS extensions, you can specify the action points for which you want to receive notifications.

 Note

The AWS AppConfig deployment events to EventBridge extension doesn't support the PRE_* action points. You can create a custom version if you want to remove some of the default actions points assigned to the AWS authored version.

You don't need to create a Lambda function if you create custom versions of the AWS authored notification extensions. You only need to specify an Amazon Resource Name (ARN) in the `Uri` field for the new extension version.

- For a custom EventBridge notification extension, enter the ARN of the EventBridge default events in the `Uri` field.
- For a custom Amazon SNS notification extension, enter the ARN of an Amazon SNS topic in the `Uri` field.
- For a custom Amazon SQS notification extension, enter the ARN of an Amazon SQS message queue in the `Uri` field.

Step 4: Create an extension association for a custom AWS AppConfig extension

To create an extension, or configure an AWS authored extension, you define the action points that invoke an extension when a specific AWS AppConfig resource is used. For example, you can choose to run the AWS AppConfig deployment events to Amazon SNS extension and receive notifications on an Amazon SNS topic anytime a configuration deployment is started for a

specific application. Defining which action points invoke an extension for a specific AWS AppConfig resource is called an *extension association*. An extension association is a specified relationship between an extension and an AWS AppConfig resource, such as an application or a configuration profile.

A single AWS AppConfig application can include multiple environments and configuration profiles. If you associate an extension to an application or an environment, AWS AppConfig invokes the extension for any workflows that relate to the application or environment resources, if applicable.

For example, say you have an AWS AppConfig application called MobileApps that includes a configuration profile called AccessList. And say the MobileApps application includes Beta, Integration, and Production environments. You create an extension association for the AWS authored Amazon SNS notification extension and associate the extension to the MobileApps application. The Amazon SNS notification extension is invoked anytime the configuration is deployed for the application to any of the three environments.

Use the following procedures to create an AWS AppConfig extension association by using the AWS AppConfig console.

To create an extension association (console)

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **AWS AppConfig**.
3. On the **Extensions** tab, choose an option button for an extension and then choose **Add to resource**. For the purposes of this walkthrough, choose **MyS3ConfigurationBackUpExtension**.
4. In the **Extension resource details** section, for **Resource type**, choose an AWS AppConfig resource type. Depending on the resource you choose, AWS AppConfig prompts you to choose other resources. For the purposes of this walkthrough, choose **Application**.
5. Choose an application in the list.
6. In the **Parameters** section, verify that **S3_BUCKET** is listed in the **Key** field. In the **Value** field, paste the ARN of the Lambda extensions. For example: `arn:aws:lambda:aws-region:111122223333:function:MyS3ConfigurationBackUpExtension`.
7. Choose **Create association to resource**.

After you create the association, you can invoke the MyS3ConfigurationBackUpExtension extension by creating a new configuration profile that specifies hosted for its SourceUri.

As a part of the workflow to create the new configuration, AWS AppConfig encounters the `PRE_CREATE_HOSTED_CONFIGURATION_VERSION` action point. Encountering this action point invokes the `MyS3ConfigurationBackUpExtension` extension, which automatically backs up the newly created configuration to the S3 bucket specified in the `Parameter` section of the extension association.

AWS AppConfig experimentation

Application experimentation is the systematic process of testing hypotheses on software applications. It uses A/B testing, feature flags, and analytics to measure user behavior, optimize performance, and drive data-driven product decisions. Experimentation enables teams to validate changes, such as new features or UI changes, before a full rollout, thereby reducing risk and improving engagement.

Built with the expertise of 25+ years of software experimentation in Amazon (including on Amazon.com), AWS AppConfig experimentation integrates seamlessly with your existing AWS AppConfig workflows, providing a comprehensive solution for feature management and experimentation in cloud-native applications. AWS AppConfig experimentation offers the following features:

- Easy experiment setup and management through the AWS Management Console.
- Powerful and flexible segmentation options for experiment participants.
- Best-practices for data science analytics between control and treatment groups.
- Data independence – use your existing data warehouses and analytics tools for experiments or use an AWS-native service like CloudWatch.

Note

Use AWS AppConfig experimentation on a pay-as-you-go basis. You are billed for *experiment hours run*. For more information, see [AWS AppConfig pricing](#). (AWS AppConfig is listed as a feature in Systems Manager.)

Experimentation terminology

Familiarize yourself with the following terminology and concepts before you begin working with AWS AppConfig experimentation.

Experiment

A configured test in which a control and one or more treatments are deployed to defined targets, and their performance is measured to evaluate the impact of changes.

Exposure

The event in which an assigned user actually encounters an experiment treatment or control in their application experience. The *audience exposure* setting controls the percentage of eligible users who receive this experience.

Experiment run

A single execution of an experiment in which treatments are delivered to the target audience and metrics are collected. You can start multiple experiment runs from the same experiment definition. AWS AppConfig experimentation have priced calls, and you will be charged per hour per experiment run from when the experiment starts until it is ended. [Learn more about pricing.](#)

Treatments

The alternative configurations, features, or experiences evaluated in an experiment, each representing a distinct variation compared to a baseline (control). For example, the control could be the current version of an LLM model, or an existing add-to-cart button experience. A treatment could be a new LLM version or model, or an experience with a different color button.

Control

A stable, well-understood configuration or experience against which treatments are compared. The control is often the current default behavior of the application, but it can be any known-good baseline suitable for measuring the impact of changes.

Overrides

Configuration rules that supersede experiment assignments by forcing specific targets or conditions to receive a designated configuration instead of their assigned control or treatment. Overrides are often used for internal testing.

Testing

The stage of an experiment in which assigned targets are exposed to a control and one or more treatments. During the test, outcome metrics need to be collected and evaluated to determine the experiment result. Note the following about the testing phase:

- Testing begins when traffic is first exposed to the experiment.
- Testing ends when you stop data collection to make a decision.
- Testing may include gradual traffic increases.
- Test results determine whether to:

- Promote a treatment
- Start another experiment run
- Abort the experiment

Example experiment workflow


At a high level, the process of creating, running, and ending an experiment using AWS AppConfig experimentation works as follows:

1. Define the experiment objective, including a name, hypothesis, and launch criteria. For more information, see [Documenting your hypothesis](#). As an Amazon best practice, spend sufficient time on this: the most successful experiments have robust hypotheses and launch criteria.
2. Complete prerequisites, including configuring IAM permissions and installing AWS AppConfig Agent. For more information, see [Step 1: Configuring prerequisites](#).
3. Configure an analytics platform to capture experiment data. This may be CloudWatch, Snowflake, Datadog, or other.
4. Define the target audience using rules that filter users by attributes such as region, account type, or device. For more information, see [Specifying the target audience for the experiment](#).
5. Select an existing feature flag to control treatment delivery. For more information, see [Selecting the experiment feature flag](#).
6. Add treatments that define the control and one or more variations to evaluate. For more information, see [About controls and treatments](#).
7. Start the experiment at 0% exposure, validate with overrides, then gradually increase exposure while monitoring metrics. For more information, see [About running and monitoring an experiment](#).
8. In your data warehouse of choice, analyze results, roll out the winning treatment, and clean up the experiment. For more information, see [Promoting a winning treatment](#).

For more detailed information about AWS AppConfig experimentation processes and tasks, including screenshots with values for a sample experiment called "Implement add-to-cart button for all products displayed", see [About experiments in AWS AppConfig](#). To get started with an experiment, see [Creating and running an experiment](#).

About experiments in AWS AppConfig

This section provides in-depth information about each of the required tasks for setting up and running an experiment using AWS AppConfig experimentation. To help illustrate how to complete these tasks, this section uses an example called *Increase add-to-cart rate by adding a button for each product displayed* experiment. For this experiment, a DevOps team wants to increase add-to-cart rates (and potentially sales rates) by adding a button beneath *each* item displayed on their site, like the following example from Amazon.com:

 Amazon Basics 2-Ply Flex-Sheets Paper Towels, 12 Basics Rolls = 40 Regular Rolls, Everyday Value with 150 Sheets per roll... 150 sheet (Pack of 12) 4.1 ★★★★★ (69.6K) 100K+ bought in past month \$22⁸⁶ (\$1.27/100 Sheets) Extra 25% off when you subscribe ⚡ 3 hour delivery for \$4.99 or ✓prime FREE delivery Today 2 PM - 6 PM 🌱 Forestry practices ▼ Add to cart	 Amazon Basics Mint Dental Flossers, Shred Resistant, Slides Easily Between Teeth, with Fold-Away Pick, 90 Count, 1-Pack... 90 Count (Pack of 1) 4.7 ★★★★★ (8.5K) 10K+ bought in past month Limited time deal \$2⁷⁶ (\$0.03/count) Typical: \$3.09 10% off on any 4 qualifying items ⚡ 3 hour delivery for \$4.99 or ✓prime FREE delivery Today 2 PM - 6 PM on \$10 of qualifying items Add to cart	 Best Seller Amazon Basics Disinfecting Wipes, Lemon & Fresh Scent, Sanitizes/Cleans/Disinfects/Deodorizes... 85 Count (Pack of 3) 4.7 ★★★★★ (91.4K) 80K+ bought in past month \$12¹⁴ (\$0.05/count) Extra 25% off when you subscribe ⚡ 3 hour delivery for \$4.99 or ✓prime FREE delivery Today 2 PM - 6 PM Add to cart
--	---	--

After you understand how to complete these tasks in AWS AppConfig experimentation, see [Creating and running an experiment](#) to get started.

Topics

- [Documenting your hypothesis](#)
- [Specifying the target audience for the experiment](#)
- [Selecting the experiment feature flag](#)
- [About controls and treatments](#)
- [About running and monitoring an experiment](#)
- [About data collection](#)

Documenting your hypothesis

When you create a new experiment, you specify a name for it and define the goals on the **Document your hypothesis** page in AWS AppConfig experimentation. Enter a name that identifies the experiment goal, for example *Increase add-to-cart rate by adding a button for each product displayed*. If you're creating multiple versions of the same experiment, specify version numbers in the name, for example *V2-Increase add-to-cart rate by adding a button for each product displayed*.

In the **Experiment hypothesis** text box, define the outcome you're trying to achieve and other details so that other users can quickly understand the experiment goals. Focus on a measurable customer or business impact. Here's an example:

- Goal: Increase add-to-cart rate by adding a button for each product displayed
- Hypothesis: Implementing 'Add-to-cart' buttons beneath every product displayed on the page will improve clicks
- Primary metric: Add-to-cart rate
- Guardrail metrics: Page load time, checkout completion rate

After naming and defining the experiment hypothesis, select the application where you want to run the experiment. An application in AWS AppConfig is simply an organizational construct like a folder that identifies the namespace of your application.

Specifying the target audience for the experiment

The target audience defines which users are eligible to be assigned to treatments during an experiment run. Selecting the right audience is important for collecting reliable data, reducing risk, and ensuring that experiment results apply to your production environment. When you configure an experiment, choose an audience that is appropriate for the change you are evaluating and the

outcomes you want to measure. Consider the following important details for selecting a target audience:

Representativeness

Choose an audience that reflects the users you plan to impact. Include key user segments, such as new and returning users, if they are relevant to your experiment. If the audience is not representative, experiment results may not generalize to your broader user base. Conversely, choose an audience that does not bias too much toward your launch criteria.

Audience size

Ensure that the eligible audience is large enough to produce meaningful results. Small audiences can lead to inconclusive or highly variable results. Larger audiences provide more reliable comparisons between treatments. Use the exposure setting to control risk — start at low exposure and increase gradually, rather than restricting the audience pool itself.

Risk level

Consider the potential impact of the experiment on users. Begin with lower-risk segments, such as internal users or a limited subset of your audience. Exclude critical workloads or sensitive user groups if needed. Gradually increasing exposure helps reduce the impact of unexpected issues.

Eligibility criteria

Define clear rules for including users in the target audience. Clear eligibility criteria help prevent unintended exposure. You can create audience rules that filter users based on attributes such as:

- Location or region
- Account type or user segment
- Application context, such as device type, OS type, or app version

Interaction with other experiments

Be aware of overlapping experiments. Avoid assigning the same users to multiple experiments that could conflict. Consider isolating audiences when running concurrent experiments. Overlapping experiments can make results difficult to interpret.

Alignment with rollout plans

Align the experiment audience with your intended rollout strategy. Choose an audience that reflects the users you plan to expose to the final treatment. If needed, start additional

experiment runs to validate results across broader audiences. This helps ensure that experiment results remain valid as exposure increases.

Creating audience rules

The **Audience rule** section includes a **Rule builder** tab and an **Editor** tab for creating rules to target an audience for the experiment.

Topics

- [Using the audience rule builder](#)
- [Using the editor](#)

Using the audience rule builder

The rule builder provides a visual interface for defining audience targeting rules without writing expressions. You can create conditions using attributes and combine them with logical operators to control which users are included in the experiment. In the rule builder, you define rules by combining conditions and groups:

- Conditions evaluate a single attribute (for example, `$plan equals "enterprise"`).
- Groups combine multiple conditions using logical operators such as AND or OR.

Each group evaluates to either `true` or `false`. Users who match the overall rule are included in the target audience.

Use the rule builder when you want to:

- Quickly define or modify audience rules without writing expressions
- Visualize how conditions are grouped and evaluated
- Build and validate targeting logic incrementally

Here are some examples:

A rule for targeting users in the Premium account tier.

Audience rule

Use the rule builder or text editor to define which visitors are eligible for this experiment.

[Rule builder](#) | [Editor](#)

[Rule blueprints](#) ▼

\$accountTier

Equals

"Premium"

⋮ ×

A rule for targeting users in a specific geographic region using a specific email domain.**Audience rule**

Use the rule builder or text editor to define which visitors are eligible for this experiment.

[Rule builder](#) | [Editor](#)

[Rule blueprints](#) ▼

And ▼

⋮ ×

\$email

Ends with ▼

"@example.com"

⋮ ×

\$continent

Equals

"EU"

⋮ ×

[Add](#) ▼

A rule for targeting users in a specific on a specific plan or using a secondary plan while having opted into a beta program.**Audience rule**

Use the rule builder or text editor to define which visitors are eligible for this experiment.

[Rule builder](#) | [Editor](#)

[Rule blueprints](#) ▼

Or ▼

⋮ ×

And ▼

⋮ ×

\$plan

Equals

"enterprise"

⋮ ×

\$region

Equals

"us-east-1"

⋮ ×

[Add](#) ▼

And ▼

⋮ ×

\$plan

Equals

"pro"

⋮ ×

\$betaOptIn

Equals

true

⋮ ×

[Add](#) ▼

[Add](#) ▼

Relationship to the editor

The rule builder and the editor define the same targeting logic in different formats:

- The rule builder provides a visual way to construct rules.
- The editor shows the equivalent logical expression.

Changes made in one view are reflected in the other.

Best practices

- Keep rules readable – Use groups to organize conditions logically and avoid deeply nested structures.
- Validate attribute values – Ensure that attribute names and values match those provided by your application.
- Test before exposure – Use 0% exposure and treatment assignment overrides to confirm that your rules behave as expected.

Using the editor

The editor lets you define targeting rules that determine which users are included in an experiment. You specify these rules as logical expressions that evaluate user or device attributes, such as platform, region, or application version. When an experiment runs, AWS AppConfig evaluates the expression for each user. Users who match the criteria are included in the target audience and can be assigned to treatments.

Expression format

Expressions use a prefix (function-style) syntax:

```
(operator argument1 argument2 ...)
```

For example:

```
(and
  (eq $system "iPhone")
  (lt $osVersion 1.8)
)
```

This expression includes users who:

- Use an iPhone, and
- Have an operating system version earlier than 1.8

The editor supports the same targeting scenarios as the rule builder. For visual examples of audience rules, see [Using the audience rule builder](#).

Defining attributes

Attributes represent properties of the user or environment being evaluated. Attributes are referenced with a \$ prefix.

Common examples include:

- \$system – device or platform type
- \$osVersion – operating system version
- \$region – geographic region
- \$appVersion – application version

The available attributes depend on your application and how data is passed to AWS AppConfig.

Writing expressions

You can combine multiple conditions using logical operators:

- Use `and` to require all conditions to be true
- Use `or` to match any condition
- Use `not` to exclude conditions

Example:

```
(and
  (eq $region "us-east-1")
  (or
    (eq $system "iPhone")
    (eq $system "Android")
  )
)
```

```
)
```

This expression targets users in us-east-1 who are using either iPhone or Android devices.

Best practices

- Start simple – Begin with a small set of conditions and expand as needed.
- Validate attributes – Ensure that the attributes used in expressions are consistently populated.
- Avoid overly complex expressions – Complex logic can make targeting harder to understand and maintain.
- Test before exposure – Use 0% exposure and treatment assignment overrides to confirm that your rules behave as expected.

Expression operators

The following operators are supported in audience expressions.

Logical operators

Operator	Description	Example
and	Returns true if all conditions are true	<pre>(and (eq \$region "us") (eq \$system "iPhone"))</pre>
or	Returns true if any condition is true	<pre>(or (eq \$system "iPhone") (eq \$system "Android"))</pre>
not	Returns true if the condition is false	<pre>(not (eq \$region "eu- west-1"))</pre>

Comparison operators

Operator	Description	Example
eq	Equals	<pre>(eq \$system "iPhone")</pre>

Operator	Description	Example
ne	Not equal	(ne \$region "us-west-2")
lt	Less than	(lt \$osVersion 1.8)
lte	Less than or equal	(lte \$osVersion 2.0)
gt	Greater than	(gt \$appVersion 5.0)
gte	Greater than or equal	(gte \$appVersion 5.0)

String and set operators

Operator	Description	Example
in	Matches any value in a list	(in \$region ["us-east-1" "us-west-2"])
contains	Checks if a value contains a substring	(contains \$deviceName "iPhone")

Note

Available operators may vary depending on how attributes are defined and evaluated in your application.

Selecting the experiment feature flag

AWS AppConfig experiments use feature flags to deliver treatments to users. The feature flag defines the possible experiences that users can receive during an experiment run. Each treatment corresponds to a distinct feature flag configuration or variation.

During the experiment run, AWS AppConfig assigns eligible users to treatments and serves the associated feature flag values. AWS AppConfig experimentation provides the configuration interface for defining experiment parameters, audience rules, and treatment settings.

When you select the feature flag for an experiment, ensure that it includes the flag attributes and constraints needed for the experiment. An *attribute* is a field that you associate with your feature flag to control the feature behavior. Here are a few examples: `buttonColor = "green"` and `layout = "expanded"`. *Constraints* are guardrails. For example, they ensure `buttonColor` must be one of: `"blue" | "green" | "red"`.

The feature flag:

- Defines the experiences being evaluated
- Delivers treatments to users at runtime
- Controls how users are exposed to each treatment
- Enables gradual rollout of a selected treatment after the experiment

In practice, the feature flag acts as the operational mechanism behind the experiment.

Each treatment should map to a distinct feature flag configuration. For example, the experiment feature flag could include:

- A flag value that retains the current UI experience (the control)
- A second flag value that defines the *Implement add-to-cart button for all products displayed* experience
- A third flag value that defines an alternative add-to-cart button experience, perhaps using a different button color or font

After you've selected the feature flag and configured different flag values for treatments, configure your application to retrieve the flag data from AWS AppConfig Agent running on your compute environment. For more information, see [Using AWS AppConfig Agent to read a specific feature flag](#).

About controls and treatments

In AWS AppConfig experimentation:

- The control treatment represents a stable, well-understood configuration or experience used as the baseline for comparison. This is often the current default behavior, but it can be any known-good state.
- A treatment represents a variation being evaluated against the control.

Each treatment corresponds to a feature flag configuration or value that your application uses to determine the experience shown to users. Experiment results are evaluated by comparing treatment performance against the control. Treatments can represent:

- UI changes
- Feature configurations
- Recommendation algorithms
- Application behavior
- Operational settings

Each treatment should represent a clear and measurable variation from the control.

When you define treatments in AWS AppConfig, provide clear descriptions for controls and treatments. Good descriptions help teams understand experiment intent and results. Descriptions should explain:

- What changed
- What users will experience
- What behavior or outcome is being evaluated

Add screenshots when possible to document the user experience associated with each treatment. Screenshots can help:

- Validate implementation
- Clarify visual differences
- Improve collaboration between teams
- Simplify experiment review

Screenshots are especially useful for UI experiments.

Designing effective treatments

Keep treatments focused and measurable. Change one major variable at a time, and avoid combining unrelated changes into a single treatment. For example, implementing an add-to-cart button beneath all displayed items is a good treatment option. Implementing the add-to-cart button while also changing the page layout and recommendation logic at the same time would return uncertain data. Smaller changes make results easier to interpret.

Important

Don't modify treatment behavior during a run. Changing treatment behavior mid-run can invalidate experiment data and make results difficult to interpret. If you need to make changes:

- Stop the run
- Modify the treatment
- Start a new experiment run
- Use meaningful names that clearly identify the purpose of the new or revised treatment

Ensure that users receive the same treatment assignment throughout an experiment run. Inconsistent assignment can affect both user experience and experiment validity.

About treatment assignment overrides

Treatment assignment overrides allow you to assign specific entity IDs to a treatment directly. Overrides bypass random treatment assignment and ensure that selected entities consistently receive the specified treatment. Overrides are commonly used to test treatments before exposing them to your audience. By assigning overrides and setting 0% exposure, you can:

- Validate that treatments render correctly
- Verify that metrics and logging work as expected
- Confirm that feature flag evaluation behaves correctly

Overrides also allow you to preview treatments for stakeholders without affecting the broader audience. You can assign entity IDs to any treatment, including the control, which is useful for verifying that control users receive the expected baseline experience.

Overrides are useful for:

- Product reviews
- Design validation
- Executive demonstrations
- QA verification
- Testing specific scenarios

This is one of the safest ways to test an experiment before increasing exposure.

Overrides are configured using entity IDs. An entity ID uniquely identifies the user or system being evaluated during feature flag assignment. An entity ID can be any of the following:

- User IDs
- Session IDs
- Device IDs
- Account identifiers

 Note

Overrides are intended for testing and validation, not long-term audience segmentation. Remove overrides when they are no longer needed, and use audience rules for production targeting logic.

You can add, modify, or remove treatment assignment overrides after an experiment run has started. This allows you to expand testing or grant additional stakeholders access to treatment previews without stopping the run.

About running and monitoring an experiment

Running an experiment in production requires careful monitoring and operational planning. Gradually increasing exposure, validating implementation, and monitoring performance can help reduce risk and improve the reliability of experiment results.

Topics

- [Operational considerations](#)
- [Performance considerations](#)
- [Operational principles](#)

Operational considerations

Start with low or 0% exposure and increase exposure gradually

Before exposing users to treatments:

- Validate treatment rendering
- Confirm feature flag evaluation
- Verify metrics and logging
- Test instrumentation using treatment assignment overrides

Starting with 0% exposure helps detect implementation issues before affecting users. You can then increase exposure gradually and thereby limit the impact of regressions if problems occur. By validating experiment behavior under increasing load, you can monitor experiment results and operational metrics before exposing your entire audience to changes.

Define rollback procedures

Before starting an experiment run, determine:

- When the run should be stopped
- Which metrics indicate unacceptable behavior
- How to disable treatments

You can configure alarms in CloudWatch or third-party monitoring tools such as Datadog, New Relic, or Dynatrace to monitor key metrics like page load time, conversion rate, or error rates during an experiment. If an alarm triggers, evaluate the impact and decide whether to stop the experiment run or continue. For more information about configuring alarms, see [Monitoring deployments for automatic rollback](#).

Coordinate deployments

Avoid introducing major application changes during an active experiment run. Changes to feature flag logic, backend systems, or user flows can invalidate experiment results or complicate analysis. If changes are required, stop the current run and start a new one after updates have been applied.

Also, limit overlapping experiments affecting the same users or functionality. Overlapping experiments can skew results and introduce conflicting behavior. In short, coordinate experiments carefully when audiences overlap.

For details about how your application retrieves an assigned treatment from AWS AppConfig Agent, including the `Entity-Id` and `Context` headers, see [Retrieving experiment treatments](#).

Performance considerations

When you start an experiment, monitor application performance by tracking operational metrics throughout the run. Specifically, monitor:

- Latency
- Error rates
- Throughput
- Resource utilization

A treatment may affect system behavior even if experiment metrics appear positive. Gradual exposure helps identify these issues safely.

Validate client-side performance

For UI experiments, monitor:

- Page load time
- Rendering performance
- Client-side errors
- Device-specific behavior

Different treatments may affect frontend performance differently.

Operational principles

As an experiment is running, keep in mind the following operational principles:

Treat experiments as production changes

Even small experiments can affect application stability, performance, user experience, and data quality. Careful monitoring and gradual exposure help reduce risk while improving confidence in experiment results.

Maintain treatment consistency

Users should receive the same treatment throughout the run. For more information about treatment design, see [About controls and treatments](#).

Validate instrumentation

Before increasing exposure, confirm that metrics are recorded correctly and that logs include treatment information.

Monitor experiment-specific metrics

In addition to operational metrics, monitor metrics directly related to the experiment, such as:

- Conversion rate
- Feature engagement
- Click-through rate
- Error frequency
- User retention

Choose metrics that align with the experiment objective.

About data collection

Although AWS AppConfig experimentation provides aggregate treatment traffic metrics on your experiment, it does not provide analytics for experiment results. The experiment dashboard includes real-time monitoring of experiment traffic. You can view audience exposure levels, treatment allocation over time, and traffic distribution across treatments. For more information, see [Observing treatment traffic](#).

To monitor operational metrics during a run, choose **View in CloudWatch** from the experiment dashboard. In CloudWatch, you can track metrics such as page load time, conversion rate, and error rates, and set alarms to notify you when a metric crosses a threshold.

For experiment results analysis such as conversion rates, user engagement, and feature adoption, use CloudWatch or your existing analytics platforms. You can export experiment data to data

warehouses such as [Amazon Redshift](#), Snowflake, or Databricks. You can also integrate with monitoring tools such as Datadog, New Relic, or Dynatrace, or use AWS services such as Amazon S3 for durable data storage and [Amazon CloudWatch RUM](#) to capture client-side user interactions. This approach gives you full control over how metrics are defined, collected, and interpreted.

Exporting experiment assignment data

To capture which entities are assigned to which treatments during an experiment run, configure AWS AppConfig Agent with the `EXPERIMENT_ASSIGNMENT_LOG_DESTINATION` option (on the Lambda extension, `AWS_APPCONFIG_EXTENSION_EXPERIMENT_ASSIGNMENT_LOG_DESTINATION`). A treatment assignment occurs when a request for the experiment flag is made on behalf of an entity that is currently exposed to the experiment. If you do not set this option, the agent does not emit assignment records.

Important

AWS AppConfig Agent writes the entity ID to assignment records exactly as your application provides it in the `Entity-Id` header. If your entity IDs contain sensitive data or personally identifiable information (PII), such as email addresses, consider hashing or pseudonymizing them before you pass them to the agent so that the values written to your logs are not sensitive.

The entity ID is the key you use to join assignment records to your metric (observations) data. If you obfuscate the entity ID, you must apply the same transformation to both data sets. If you obfuscate the entity ID in your assignment records but not in your metric data, or apply a different transformation to each, you will not be able to join the two.

Emitting records to stderr

When you set the destination to `stderr`, the agent emits each assignment as a single-line JSON record. Each record includes a `type` field so you can distinguish assignment records from the agent application logs, which also go to `stderr` by default.

```
{
  "type": "AWS.AppConfig.TreatmentAssignment",
  "timestamp": "2026-01-01T12:00:00Z",
  "region": "us-east-1",
  "accountId": "111122223333",
```

```

"applicationId": "app1234",
"experimentDefinitionId": "exp1234",
"experimentRunNumber": "1",
"treatmentKey": "t1",
"entityId": "user123"
}

```

The `treatmentKey` field is the key that AWS AppConfig assigns to each treatment (c for the control, and t1, t2, and so on for treatments). This is the bare form of the key. In the agent flag response, the same treatment appears as a flag variant named with a double-underscore pattern (for example, `__t1__`); for more information, see [Retrieving experiment treatments](#).

Writing records to disk

When you set the destination to a value like `file:path/`, the agent writes assignment records to disk under the provided base directory. Records are stored in a file structure that uniquely identifies the experiment run. Each log file name includes a random salt to prevent collisions when files from multiple hosts are combined.

```

path/
us-east-1/           # Region
  111122223333/      # account ID
    app1234/         # application ID
      exp1234/        # experiment definition ID
        1_a4b5.jsonl # log file: <run-number>_<salt>.jsonl

```

Each file contains newline-delimited JSON records:

```

{"timestamp":"2026-01-01T12:00:00Z","treatmentKey":"t1","entityId":"user123"}

```

Observing treatment traffic

AWS AppConfig Agent collects aggregated treatment assignment counts and emits them as CloudWatch metrics in your account through a service-linked role that AWS AppConfig manages. There is no additional charge for these metrics.

- Namespace: AWS/AppConfig
- Metric name: ReturnedCount
- Dimensions: application ID, environment ID, configuration profile ID, flag key, and variant

These metrics show how much traffic each treatment is receiving, which you can use as a real-time check that your experiment is operating as expected. For example, in a 50/50 A/B test, you can confirm that about half of entities are assigned to the control and about half to the treatment.

Note

To emit experiment treatment traffic metrics to your account, you must opt in to the VendedMetrics account setting by using the UpdateAccountSettings operation. Vended metrics incur no additional cost.

Creating and running an experiment

This section describes how to create and run an experiment using AWS AppConfig experimentation. For more detailed information about the concepts discussed in each of the following steps, see [About experiments in AWS AppConfig](#).

Topics

- [Step 1: Configuring prerequisites](#)
- [Step 2: Creating an experiment definition](#)
- [Step 3: Starting an experiment](#)
- [Step 4: Increasing exposure](#)

Step 1: Configuring prerequisites

Before you begin, complete the following tasks:

1. **Deploy the feature flag you want to experiment on:** Before you can run an experiment on a feature flag, the flag must be deployed to the environment where you want to run the experiment.
2. **[Install and configure AWS AppConfig Agent:](#)** AWS AppConfig experimentation requires AWS AppConfig Agent to deliver treatments to users. Experimentation requires AWS AppConfig Agent Lambda extension version 2.0.20159 or later, or Amazon ECS/Amazon EC2 agent version 2.0.175233 or later. The following topics describe how to install and configure AWS AppConfig Agent for each supported compute environment:
 - [Using AWS AppConfig Agent with AWS Lambda](#)

- [Using AWS AppConfig Agent with Amazon EC2 and on-premises machines](#)
- [Using AWS AppConfig Agent with Amazon ECS and Amazon EKS](#)

Note that each section includes information about configuring IAM permissions so the agent can retrieve feature flags and other configuration data.

3. [Configure experiment assignment logging](#): To capture treatment assignment data during an experiment run, set the `EXPERIMENT_ASSIGNMENT_LOG_DESTINATION` agent option (Lambda: `AWS_APPCONFIG_EXTENSION_EXPERIMENT_ASSIGNMENT_LOG_DESTINATION`) to `stderr` or to a disk path such as `file:/var/log/appconfig/experiments/`.

Step 2: Creating an experiment definition

An experiment definition describes the purpose, scope, and operational configuration of an experiment. It captures information about what is being tested, the expected outcome, and the business or technical objective so that stakeholders can quickly understand the intent of the experiment and interpret the results consistently.

The experiment definition also includes the operational details required to run the experiment in AWS AppConfig. This includes the application and feature flag associated with the experiment, the configuration profile used to store experiment settings, and the treatment configurations that define the experiences being evaluated. The experiment definition acts as the central configuration for experiment runs, audience exposure settings, and treatment assignment logic. For more detailed information about the concepts described here, see [About experiments in AWS AppConfig](#).

Topics

- [Creating an experiment definition \(console\)](#)

Creating an experiment definition (console)

Use the following procedure to create an experiment definition using AWS AppConfig experimentation in the AWS Management Console.

To create an experiment definition

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Experiments**, and then choose **Create experiment**.

3. For **Experiment Name**, enter a descriptive name to distinguish this experiment from other experiments.
4. For **Experiment hypothesis**, describe the goal or hypothesis you want to validate.
5. For **Launch criteria**, enter information about metrics or data essential to deploying a winning treatment. This information is helpful for long-running experiments.
6. For **Experiment application**, select the application where this experiment will run.
7. Choose **Next**.
8. On the **Specify target audience** page, for **Audience description**, enter information about the intended audience for the experiment.
9. In the **Audience rule** section, choose either the **Rule builder** tab or the **Editor** tab to create a rule that identifies or filters the audience. Choose **Rule blueprints** to view and select a pre-constructed rule for your experiment. When building a rule:

- Add a condition

Choose an attribute, select an operator, and specify a value.

- Group conditions

Use AND to require all conditions in a group to be true, or OR to match any condition.

- Nest groups

Combine multiple groups to create more advanced targeting logic.

- Refine the rule

Add or remove conditions and groups as needed to match your targeting requirements.

For more information about building rules, including expression formats, supported operators, and rule examples, see [Creating audience rules](#).

10. Choose **Next**.
11. On the **Select experiment feature flag** page, for **Environment**, select the environment where this experiment will run. For **Configuration profile**, select the configuration profile associated with the feature flag.
12. For **Feature flag**, select an existing feature flag to use with the experiment. For more information about experiment feature flags, see [Selecting the experiment feature flag](#).
13. Choose **Next**.

14. On the **Add treatments** page, review the **Default flag value** section. When audience exposure is set between 0% and 99%, users not assigned to a treatment receive the default flag value. This value is always off and cannot be changed.
15. On the **Add treatments** page, in the **Control** section, enter a description and verify the flag value is toggled **ON**.

 Note

The control treatment represents the baseline experience used for comparison during the experiment. It is typically the current default behavior, but can be any stable, well-understood configuration. The *default feature flag value* (which is toggled off by default) defines the fallback behavior for users who are not assigned to a treatment. Although the control treatment and default flag value often deliver the same user experience, only users assigned to the control treatment participate in experiment measurement and analysis.

If one or more keys are listed in the **Attribute values** section, specify their values in the corresponding **Value** field.

16. In the **Treatment 1** section, enter a description and verify the flag value is toggled **ON**. If one or more keys are listed in the **Attribute values** section, specify their values in the corresponding **Value** field. Repeat this process for any additional treatments.
17. AWS AppConfig experimentation automatically allocates traffic to all treatments based on the total number of treatments. If you want to allocate traffic to different treatments, you can expand **Advanced settings**, and choose **Use custom allocation weights** (not recommended).

 Note

Treatment traffic allocation determines how traffic is split among treatments. Audience exposure, which controls what percentage of the eligible audience enters the experiment, is configured separately when you start the experiment run.

18. Choose **Next**.
19. Review the details, and then choose **Save experiment definition**. AWS AppConfig displays the completed definition page for the experiment.

Step 3: Starting an experiment

Before you start an experiment, verify you completed the prerequisites. For more information, see [Step 1: Configuring prerequisites](#). Before running an experiment in production, read [About running and monitoring an experiment](#) to familiarize yourself with recommended practices for safely starting and monitoring the experiment.

To start an experiment

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Experiments**, and then choose an experiment to run. The experiment definition page opens.
3. Choose **Start experiment run**.
4. (Optional) In the **Description** field, enter information about this particular experiment run.
5. For an initial experiment run, verify that **Test first at 0% exposure on start** is selected.
6. If you plan to use treatment assignment overrides, choose **Add entity ID**. For more information, see [About treatment assignment overrides](#).
7. (Optional) Expand **Tags** to assign one or more key-value pairs to the experiment run.
8. (Optional) Expand **Deployment parameters** to assign a AWS Key Management Service (AWS KMS) key and tags to the experiment run.
9. For an initial experiment run, leave audience exposure set to 0%.
10. Choose **Start experiment run**. AWS AppConfig experimentation opens the dashboard for the running experiment.

At 0% exposure, no audience traffic is assigned to treatments unless overrides are configured. Use this stage to validate feature flags, treatment logic, metrics, and logging before increasing exposure. For more information, see [Operational considerations](#).

Step 4: Increasing exposure

After validating the experiment at 0% exposure, gradually increase audience exposure while monitoring both experiment metrics and operational health. As exposure increases, verify that treatments continue to behave as expected under production traffic and monitor metrics such as latency, error rates, resource utilization, conversion rates, and user engagement. Increase exposure

incrementally to reduce operational risk and allow time to detect unexpected behavior, invalid experiment data, or performance regressions before they affect a larger portion of your audience. Monitor alarms configured in CloudWatch or your preferred monitoring tools throughout the run and be prepared to stop the run if operational or business metrics exceed acceptable thresholds.

Important

If you detect operational issues, invalid experiment behavior, or unexpected treatment performance as you increase exposure, stop the experiment run to prevent additional users from being exposed to treatments. Stopping the run ends audience exposure and returns users to the currently deployed feature flag configuration. For more information, see [Stopping an experiment](#).

To increase exposure

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Experiments**, and then choose a running experiment. The experiment dashboard opens.
3. In the **Audience exposure** section, choose **Increase exposure**.
4. On the **Increase exposure** page, set the new exposure percentage by using the exposure dial or entering a value in the **Increase exposure to** field.
5. Review the exposure change summary. The page displays the percentage of traffic that will be served treatments and the percentage that will receive the default experience.

Important

Once audience exposure is increased, it cannot be decreased within the same experiment run. To reduce exposure, you must stop the experiment run and start a new run at a lower percentage.

6. In the **Confirmation** section, type **confirm**.
7. (Optional) Expand **Deployment parameters** to assign a AWS Key Management Service (AWS KMS) key and tags to the experiment run.
8. Choose **Increase exposure**. AWS AppConfig experimentation opens the experiment dashboard with the increased exposure displayed.

Monitor traffic and metrics as described earlier. When you're ready, repeat this procedure to increase audience exposure further.

Stopping an experiment

You can stop an experiment run at any time. Common reasons for stopping include:

- **Operational issues, invalid experiment behavior, or unexpected treatment performance:** Stopping the experiment run prevents additional users from being exposed to treatments and returns users to the currently deployed feature flag configuration.
- **A clear treatment winner:** Results consistently meet the success criteria and operational metrics remain within acceptable thresholds. Before promoting a winner, verify that the treatment performs better than the control for the intended outcome. For information about rolling out the winning treatment, see [Promoting a winning treatment](#).
- **Inconclusive results or changing priorities:** The experiment is no longer relevant, the hypothesis needs revising, or you want to restart with a different configuration.

Use the following procedure to stop an experiment.

To stop an experiment

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Experiments**, and then choose a running experiment. The experiment dashboard opens.
3. Choose **Stop**.
4. (Optional) In the **Experiment results summary** section, document the experiment outcome to help your team make an informed launch decision. This summary is saved with the experiment run for future reference. You can provide:
 - **Executive summary** – Summarize the experiment outcome and key findings.
 - **Reasons to launch** – Evidence in favor of launching the treatment.
 - **Reasons not to launch** – Evidence against launching the treatment.
5. In the **Confirmation** section, type **confirm**, and then choose **Stop experiment run**.

Promoting a winning treatment

After you identify a winning treatment, you can promote it to all users by updating and deploying the feature flag configuration. The order in which you perform these steps determines whether your users experience a seamless transition.

Promoting with a seamless transition

To ensure a seamless transition from your current traffic split to full treatment rollout, update and deploy the feature flag configuration while the experiment is still running:

1. Update the feature flag values in the configuration profile to match the winning treatment.
2. Deploy the updated configuration using a standard AWS AppConfig deployment strategy.
While the experiment is active, this deployment does not change what users see because the experiment is still controlling the flag.
3. Stop the experiment run. AWS AppConfig stops managing the flag, and your application begins serving the configuration you deployed in step 2 – the winning treatment – to all users.

This approach provides a seamless transition from your split traffic state (for example, 50/50) to full treatment release (100%) without any users reverting to the default experience.

Promoting after stopping the experiment

If you stop the experiment before updating the flag, AWS AppConfig stops managing the flag and your application immediately serves whatever version of the flag is currently deployed – typically the default value. Users see the default experience until you update and redeploy the flag with the winning treatment values. This approach does not provide a seamless transition.

For more information about deploying configurations, see [Deploying a configuration](#).

Cleaning up an experiment

After an experiment completes, you can archive or permanently delete the experiment definition.

- **Archive** – The experiment definition is hidden from the active list but can be restored later. The experiment feature flag remains deployed and continues serving its current configuration. To stop serving the flag, disable it separately in your feature flag configuration.
- **Delete permanently** – The experiment definition and all associated run history are permanently deleted. This action cannot be undone.

To archive or delete an experiment

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Experiments**, and then choose an experiment. The experiment dashboard opens.
3. Choose **Actions**, and then choose **Delete experiment definition**.
4. In the **Choose an action** section, choose either **Archive** or **Delete permanently**.
5. In the **Confirmation** section, type **confirm**.
6. Choose the **Archive** or **Delete** button.

Using code samples to perform common AWS AppConfig tasks

This section includes code samples for programmatically performing common AWS AppConfig actions. We recommend you use these samples with the [Java](#), [Python](#), and [JavaScript](#) SDKs to perform the actions in a test environment. This section includes a code sample for cleaning up your test environment after you finish.

Topics

- [Creating or updating a freeform configuration stored in the hosted configuration store](#)
- [Creating a configuration profile for a secret stored in Secrets Manager](#)
- [Deploying a configuration profile](#)
- [Using AWS AppConfig Agent to read a freeform configuration profile](#)
- [Using AWS AppConfig Agent to read a specific feature flag](#)
- [Using AWS AppConfig Agent to retrieve a feature flag with variants](#)
- [Using the GetLatestConfiguration API action to read a freeform configuration profile](#)
- [Cleaning up your environment](#)

Creating or updating a freeform configuration stored in the hosted configuration store

Each of the following samples includes comments about the actions performed by the code. The samples in this section call the following APIs:

- [CreateApplication](#)
- [CreateConfigurationProfile](#)
- [CreateHostedConfigurationVersion](#)

Java

```
public CreateHostedConfigurationVersionResponse createHostedConfigVersion() {  
    AppConfigClient appconfig = AppConfigClient.create();
```

```

        // Create an application
        CreateApplicationResponse app = appconfig.createApplication(req ->
req.name("MyDemoApp"));

        // Create a hosted, freeform configuration profile
        CreateConfigurationProfileResponse configProfile =
appconfig.createConfigurationProfile(req -> req
            .applicationId(app.id())
            .name("MyConfigProfile")
            .locationUri("hosted")
            .type("AWS.Freeform"));

        // Create a hosted configuration version
        CreateHostedConfigurationVersionResponse hcv =
appconfig.createHostedConfigurationVersion(req -> req
            .applicationId(app.id())
            .configurationProfileId(configProfile.id())
            .contentType("text/plain; charset=utf-8")
            .content(SdkBytes.fromUtf8String("my config data")));

        return hcv;
    }

```

Python

```

import boto3

appconfig = boto3.client('appconfig')

# create an application
application = appconfig.create_application(Name='MyDemoApp')

# create a hosted, freeform configuration profile
config_profile = appconfig.create_configuration_profile(
    ApplicationId=application['Id'],
    Name='MyConfigProfile',
    LocationUri='hosted',
    Type='AWS.Freeform')

# create a hosted configuration version
hcv = appconfig.create_hosted_configuration_version(
    ApplicationId=application['Id'],
    ConfigurationProfileId=config_profile['Id'],

```

```
Content=b'my config data',  
ContentType='text/plain')
```

JavaScript

```
import {  
  AppConfigClient,  
  CreateApplicationCommand,  
  CreateConfigurationProfileCommand,  
  CreateHostedConfigurationVersionCommand,  
} from "@aws-sdk/client-appconfig";  
  
const appconfig = new AppConfigClient();  
  
// create an application  
const application = await appconfig.send(  
  new CreateApplicationCommand({ Name: "MyDemoApp" })  
);  
  
// create a hosted, freeform configuration profile  
const profile = await appconfig.send(  
  new CreateConfigurationProfileCommand({  
    ApplicationId: application.Id,  
    Name: "MyConfigProfile",  
    LocationUri: "hosted",  
    Type: "AWS.Freeform",  
  })  
);  
  
// create a hosted configuration version  
await appconfig.send(  
  new CreateHostedConfigurationVersionCommand({  
    ApplicationId: application.Id,  
    ConfigurationProfileId: profile.Id,  
    ContentType: "text/plain",  
    Content: "my config data",  
  })  
);
```

Creating a configuration profile for a secret stored in Secrets Manager

Each of the following samples includes comments about the actions performed by the code. The samples in this section call the following APIs:

- [CreateApplication](#)
- [CreateConfigurationProfile](#)

Java

```
private void createSecretsManagerConfigProfile() {
    AppConfigClient appconfig = AppConfigClient.create();

    // Create an application
    CreateApplicationResponse app = appconfig.createApplication(req ->
    req.name("MyDemoApp"));

    // Create a configuration profile for Secrets Manager Secret
    CreateConfigurationProfileResponse configProfile =
    appconfig.createConfigurationProfile(req -> req
        .applicationId(app.id())
        .name("MyConfigProfile")
        .locationUri("secretsmanager://MySecret")
        .retrievalRoleArn("arn:aws:iam::000000000000:role/
        RoleTrustedByAppConfigThatCanRetrieveSecret")
        .type("AWS.Freeform"));
}
```

Python

```
import boto3

appconfig = boto3.client('appconfig')

# create an application
application = appconfig.create_application(Name='MyDemoApp')

# create a configuration profile for Secrets Manager Secret
config_profile = appconfig.create_configuration_profile(
```

```
ApplicationId=application['Id'],
Name='MyConfigProfile',
LocationUri='secretsmanager://MySecret',
RetrievalRoleArn='arn:aws:iam::000000000000:role/
RoleTrustedByAppConfigThatCanRetrieveSecret',
Type='AWS.Freeform')
```

JavaScript

```
import {
  AppConfigClient,
  CreateConfigurationProfileCommand,
} from "@aws-sdk/client-appconfig";

const appconfig = new AppConfigClient();

// create an application
const application = await appconfig.send(
  new CreateApplicationCommand({ Name: "MyDemoApp" })
);

// create a configuration profile for Secrets Manager Secret
await appconfig.send(
  new CreateConfigurationProfileCommand({
    ApplicationId: application.Id,
    Name: "MyConfigProfile",
    LocationUri: "secretsmanager://MySecret",
    RetrievalRoleArn: "arn:aws:iam::000000000000:role/
RoleTrustedByAppConfigThatCanRetrieveSecret",
    Type: "AWS.Freeform",
  })
);
```

Deploying a configuration profile

Each of the following samples includes comments about the actions performed by the code. The samples in this section call the following APIs:

- [CreateApplication](#)
- [CreateConfigurationProfile](#)
- [CreateHostedConfigurationVersion](#)

- [CreateEnvironment](#)
- [StartDeployment](#)
- [GetDeployment](#)

Java

```
private void createDeployment() throws InterruptedException {
    AppConfigClient appconfig = AppConfigClient.create();

    // Create an application
    CreateApplicationResponse app = appconfig.createApplication(req ->
req.name("MyDemoApp"));

    // Create a hosted, freeform configuration profile
    CreateConfigurationProfileResponse configProfile =
appconfig.createConfigurationProfile(req -> req
        .applicationId(app.id())
        .name("MyConfigProfile")
        .locationUri("hosted")
        .type("AWS.Freeform"));

    // Create a hosted configuration version
    CreateHostedConfigurationVersionResponse hcv =
appconfig.createHostedConfigurationVersion(req -> req
        .applicationId(app.id())
        .configurationProfileId(configProfile.id())
        .contentType("text/plain; charset=utf-8")
        .content(SdkBytes.fromUtf8String("my config data")));

    // Create an environment
    CreateEnvironmentResponse env = appconfig.createEnvironment(req -> req
        .applicationId(app.id())
        .name("Beta")
        // If you have CloudWatch alarms that monitor the health of your
service, you can add them here and they
        // will trigger a rollback if they fire during an appconfig deployment
        // .monitors(Monitor.builder().alarmArn("arn:aws:cloudwatch:us-
east-1:520900602629:alarm:MyAlarm")
        //
        .alarmRoleArn("arn:aws:iam::520900602629:role/MyAppConfigAlarmRole").build())
    );
}
```

```

// Start a deployment
StartDeploymentResponse deploymentResponse = appconfig.startDeployment(req -
> req
    .applicationId(app.id())
    .configurationProfileId(configProfile.id())
    .environmentId(env.id())
    .configurationVersion(hcv.versionNumber().toString())
    .deploymentStrategyId("AppConfig.Linear50PercentEvery30Seconds")
);

// Wait for deployment to complete
List<DeploymentState> nonFinalDeploymentStates = Arrays.asList(
    DeploymentState.DEPLOYING,
    DeploymentState.BAKING,
    DeploymentState.ROLLING_BACK,
    DeploymentState.VALIDATING);
GetDeploymentRequest getDeploymentRequest =
GetDeploymentRequest.builder().applicationId(app.id())

.environmentId(env.id())

.deploymentNumber(deploymentResponse.deploymentNumber()).build();
GetDeploymentResponse deployment =
appconfig.getDeployment(getDeploymentRequest);
while (nonFinalDeploymentStates.contains(deployment.state())) {
    System.out.println("Waiting for deployment to complete: " + deployment);
    Thread.sleep(1000L);
    deployment = appconfig.getDeployment(getDeploymentRequest);
}

System.out.println("Deployment complete: " + deployment);
}

```

Python

```

import boto3

appconfig = boto3.client('appconfig')

# create an application
application = appconfig.create_application(Name='MyDemoApp')

```

```
# create an environment
environment = appconfig.create_environment(
    ApplicationId=application['Id'],
    Name='MyEnvironment')

# create a configuration profile
config_profile = appconfig.create_configuration_profile(
    ApplicationId=application['Id'],
    Name='MyConfigProfile',
    LocationUri='hosted',
    Type='AWS.Freeform')

# create a hosted configuration version
hcv = appconfig.create_hosted_configuration_version(
    ApplicationId=application['Id'],
    ConfigurationProfileId=config_profile['Id'],
    Content=b'my config data',
    ContentType='text/plain')

# start a deployment
deployment = appconfig.start_deployment(
    ApplicationId=application['Id'],
    EnvironmentId=environment['Id'],
    ConfigurationProfileId=config_profile['Id'],
    ConfigurationVersion=str(hcv['VersionNumber']),
    DeploymentStrategyId='AppConfig.Linear20PercentEvery6Minutes')
```

JavaScript

```
import {
    AppConfigClient,
    CreateApplicationCommand,
    CreateEnvironmentCommand,
    CreateConfigurationProfileCommand,
    CreateHostedConfigurationVersionCommand,
    StartDeploymentCommand,
} from "@aws-sdk/client-appconfig";

const appconfig = new AppConfigClient();

// create an application
const application = await appconfig.send(
    new CreateApplicationCommand({ Name: "MyDemoApp" })
```

```
);

// create an environment
const environment = await appconfig.send(
  new CreateEnvironmentCommand({
    ApplicationId: application.Id,
    Name: "MyEnvironment",
  })
);

// create a configuration profile
const config_profile = await appconfig.send(
  new CreateConfigurationProfileCommand({
    ApplicationId: application.Id,
    Name: "MyConfigProfile",
    LocationUri: "hosted",
    Type: "AWS.Freeform",
  })
);

// create a hosted configuration version
const hcv = await appconfig.send(
  new CreateHostedConfigurationVersionCommand({
    ApplicationId: application.Id,
    ConfigurationProfileId: config_profile.Id,
    Content: "my config data",
    ContentType: "text/plain",
  })
);

// start a deployment
await appconfig.send(
  new StartDeploymentCommand({
    ApplicationId: application.Id,
    EnvironmentId: environment.Id,
    ConfigurationProfileId: config_profile.Id,
    ConfigurationVersion: hcv.VersionNumber.toString(),
    DeploymentStrategyId: "AppConfig.Linear20PercentEvery6Minutes",
  })
);
```

Using AWS AppConfig Agent to read a freeform configuration profile

Each of the following samples includes comments about the actions performed by the code.

Java

```
public void retrieveConfigFromAgent() throws Exception {
    /*
        In this sample, we will retrieve configuration data from the AWS AppConfig
        Agent.

        The agent is a sidecar process that handles retrieving configuration data
        from AppConfig
        for you in a way that implements best practices like configuration caching.

        For more information about the agent, see How to use AWS AppConfig Agent
    */

    // The agent runs a local HTTP server that serves configuration data
    // Make a GET request to the agent's local server to retrieve the
    configuration data
    URL url = new URL("http://localhost:2772/applications/MyDemoApp/
environments/Beta/configurations/MyConfigProfile");
    HttpURLConnection con = (HttpURLConnection) url.openConnection();
    con.setRequestMethod("GET");
    // To enable entity-based deployments, add the Entity-Id header to the
    request:
    // con.setRequestProperty("Entity-Id", entityId);
    StringBuilder content;
    try (BufferedReader in = new BufferedReader(new
InputStreamReader(con.getInputStream()))) {
        content = new StringBuilder();
        int ch;
        while ((ch = in.read()) != -1) {
            content.append((char) ch);
        }
    }
    con.disconnect();
    System.out.println("Configuration from agent via HTTP: " + content);
}
```

Python

```
# in this sample, we will retrieve configuration data from the AWS AppConfig Agent.
# the agent is a sidecar process that handles retrieving configuration data from AWS
  AppConfig
# for you in a way that implements best practices like configuration caching.
#
# for more information about the agent, see
# How to use AWS AppConfig Agent
#

import requests

application_name = 'MyDemoApp'
environment_name = 'MyEnvironment'
config_profile_name = 'MyConfigProfile'

# To enable entity-based deployments, add the Entity-Id header to the request:
# response = requests.get(f"http://localhost:2772/applications/{application_name}/
environments/{environment_name}/configurations/{config_profile_name}",
#                           headers={"Entity-Id": entity_id})
# the agent runs a local HTTP server that serves configuration data
# make a GET request to the agent's local server to retrieve the configuration data
response = requests.get(f"http://localhost:2772/applications/{application_name}/
environments/{environment_name}/configurations/{config_profile_name}")
config = response.content
```

JavaScript

```
// in this sample, we will retrieve configuration data from the AWS AppConfig Agent.
// the agent is a sidecar process that handles retrieving configuration data from
  AppConfig
// for you in a way that implements best practices like configuration caching.

// for more information about the agent, see
// How to use AWS AppConfig Agent

const application_name = "MyDemoApp";
const environment_name = "MyEnvironment";
const config_profile_name = "MyConfigProfile";

// the agent runs a local HTTP server that serves configuration data
// make a GET request to the agent's local server to retrieve the configuration data
```

```
const url = `http://localhost:2772/applications/${application_name}/environments/  
${environment_name}/configurations/${config_profile_name}`;  
// To enable entity-based deployments, add the Entity-Id header to the request:  
// const response = await fetch(url, { headers: { "Entity-Id": entityId } });  
const response = await fetch(url);  
const config = await response.text(); // (use `await response.json()` if your config  
is json)
```

Using AWS AppConfig Agent to read a specific feature flag

Each of the following samples includes comments about the actions performed by the code.

Java

```
public void retrieveSingleFlagFromAgent() throws Exception {  
    /*  
        You can retrieve a single flag's data from the agent by providing the  
        "flag" query string parameter.  
        Note: the configuration's type must be AWS.AppConfig.FeatureFlags  
    */  
  
    URL url = new URL("http://localhost:2772/applications/MyDemoApp/  
environments/Beta/configurations/MyFlagsProfile?flag=myFlagKey");  
    HttpURLConnection con = (HttpURLConnection) url.openConnection();  
    con.setRequestMethod("GET");  
    StringBuilder content;  
    try (BufferedReader in = new BufferedReader(new  
InputStreamReader(con.getInputStream()))) {  
        content = new StringBuilder();  
        int ch;  
        while ((ch = in.read()) != -1) {  
            content.append((char) ch);  
        }  
    }  
    con.disconnect();  
    System.out.println("MyFlagName from agent: " + content);  
}
```

Python

```
import requests
```

```

application_name = 'MyDemoApp'
environment_name = 'MyEnvironment'
config_profile_name = 'MyConfigProfile'
flag_key = 'MyFlag'

# retrieve a single flag's data by providing the "flag" query string parameter
# note: the configuration's type must be AWS.AppConfig.FeatureFlags
response = requests.get(f"http://localhost:2772/applications/{application_name}/
environments/{environment_name}/configurations/{config_profile_name}?
flag={flag_key}")
config = response.content

```

JavaScript

```

const application_name = "MyDemoApp";
const environment_name = "MyEnvironment";
const config_profile_name = "MyConfigProfile";
const flag_name = "MyFlag";

// retrieve a single flag's data by providing the "flag" query string parameter
// note: the configuration's type must be AWS.AppConfig.FeatureFlags
const url = `http://localhost:2772/applications/${application_name}/environments/
${environment_name}/configurations/${config_profile_name}?flag=${flag_name}`;
const response = await fetch(url);
const flag = await response.json(); // { "enabled": true/false }

```

Using AWS AppConfig Agent to retrieve a feature flag with variants

Each of the following samples includes comments about the actions performed by the code.

Java

```

public static void retrieveConfigFromAgentWithVariants() throws Exception {
    /*
    This sample retrieves feature flag configuration data
    containing variants from AWS AppConfig Agent.

    For more information about the agent, see How to use AWS AppConfig Agent
    */
}

```

```

    // Make a GET request to the agent's local server to retrieve the configuration
    data
    URL url = new URL("http://localhost:2772/applications/MyDemoApp/environments/
Beta/configurations/MyConfigProfile");
    HttpURLConnection con = (HttpURLConnection) url.openConnection();

    // Provide context in the 'Context' header
    // In the header value, use '=' to separate context key from context value
    // Note: Multiple context values may be passed either across
    // multiple headers or as comma-separated values in a single header
    con.setRequestProperty("Context", "country=US");

    StringBuilder content;
    try (BufferedReader in = new BufferedReader(new
InputStreamReader(con.getInputStream())) {
        content = new StringBuilder();
        int ch;
        while ((ch = in.read()) != -1) {
            content.append((char) ch);
        }
    }
    con.disconnect();
    System.out.println("Configuration from agent via HTTP: " + content);
}

```

Python

```

# This sample retrieve features flag configuration data
# containing variants from AWS AppConfig Agent.

# For more information about the agent, see How to use AWS AppConfig Agent

import requests

application_name = 'MyDemoApp'
environment_name = 'Beta'
configuration_profile_name = 'MyConfigProfile'

# make a GET request to the agent's local server to retrieve the configuration data
response = requests.get(f"http://localhost:2772/applications/{application_name}/
environments/{environment_name}/configurations/{configuration_profile_name}",
                        headers = {

```

```

        "Context": "country=US" # Provide context in the
        'Context' header
                                # In the header value, use '='
        to separate context key from context value
                                # Note: Multiple context values
        may be passed either across
                                # multiple headers or as comma-
        separated values in a single header
    }
)
print("Configuration from agent via HTTP: ", response.json())

```

JavaScript

```

// This sample retrieves feature flag configuration data
// containing variants from AWS AppConfig Agent.

// For more information about the agent, see How to use AWS AppConfig Agent

const application_name = "MyDemoApp";
const environment_name = "Beta";
const configuration_profile_name = "MyConfigProfile";

const url = `http://localhost:2772/applications/${application_name}/environments/
${environment_name}/configurations/${configuration_profile_name}`;

// make a GET request to the agent's local server to retrieve the configuration data
const response = await fetch(url, {
    method: 'GET',
    headers: {
        'Context': 'country=US' // Provide context in the 'Context' header
                                // In the header value, use '=' to separate context
        key from context value
                                // Note: Multiple context values may be passed
        either across
                                // multiple headers or as comma-separated values in
        a single header
    }
});

const config = await response.json();
console.log("Configuration from agent via HTTP: ", config);

```

Using the GetLatestConfiguration API action to read a freeform configuration profile

Each of the following samples includes comments about the actions performed by the code. The samples in this section call the following APIs:

- [GetLatestConfiguration](#)
- [StartConfigurationSession](#)

Java

```
/*
The example below uses two AWS AppConfig Data APIs: StartConfigurationSession and
GetLatestConfiguration.
For more information about these APIs, see AWS AppConfig Data.

This class is meant to be used as a singleton to retrieve the latest configuration
data from AWS AppConfig.
This class maintains a cache of the latest configuration data in addition to the
configuration token to be
passed to the next GetLatestConfiguration API call.
*/
public class AppConfigApiRetriever {

    /*
    * Set of AppConfig invalid parameter problems that require restarting the
    configuration session.
    * If the GetLatestConfiguration API call fails with any of these problems (e.g.
    token is EXPIRED or CORRUPTED),
    * we need to call StartConfigurationSession again to obtain a new configuration
    token before retrying.
    */
    private final Set<InvalidParameterProblem> SESSION_RESTART_REQUIRED =
        Stream.of(InvalidParameterProblem.EXPIRED,
InvalidParameterProblem.CORRUPTED)
            .collect(Collectors.toSet());

    /*
    AWS AppConfig Data SDK client used to interact with the AWS AppConfig Data
    service.
    */
}
```

```

private final AppConfigDataClient appConfigData;

/**
 * The configuration token to be passed to the next GetLatestConfiguration API
 * call.
 */
private String configurationToken;

/**
 * The cached configuration data to be returned when there is no new configuration
 * data available.
 */
private SdkBytes configuration;

public AppConfigApiRetriever() {
    this.appConfigData = AppConfigDataClient.create();
}

/**
 * Returns the latest configuration data stored in AWS AppConfig.
 */
public SdkBytes getConfig() {
    /**
     * If there is no configuration token yet, get one by starting a new session
     * with the StartConfigurationSession API.
     * Note that this API does not return configuration data. Rather, it returns an
     * initial configuration token that is
     * subsequently passed to the GetLatestConfiguration API.
     */
    if (this.configurationToken == null) {
        startNewSession();
    }

    GetLatestConfigurationResponse response = null;

    try {
        /**
         * Retrieve the configuration from the GetLatestConfiguration API,
         * providing the current configuration token.
         * If this caller does not yet have the latest configuration (e.g. this is
         * the first call to GetLatestConfiguration
         * or new configuration data has been deployed since the first call), the
         * latest configuration data will be returned.

```

```

        Otherwise, the GetLatestConfiguration API will not return any data since
        the caller already has the latest.
        */
        response = appConfigData.getLatestConfiguration(
            GetLatestConfigurationRequest.builder()

.configurationToken(this.configurationToken)

                                .build());

        } catch (ResourceNotFoundException e) {
            // Handle resource not found by refreshing the session
            System.err.println("Resource not found – refreshing session and
retrying...");
            startNewSession();
            response = appConfigData.getLatestConfiguration(
                GetLatestConfigurationRequest.builder()

.configurationToken(this.configurationToken)

                                .build());

        } catch (BadRequestException e) {
            // Handle expired or corrupted token by refreshing the session
            boolean needsNewSession = Optional.ofNullable(e.details())
                                                .map(details ->
details.invalidParameters()

                                                                .values()
                                                                .stream()
                                                                .anyMatch(val -
> SESSION_RESTART_REQUIRED.contains(val.problem()))
                                                                .orElse(false));

            if (needsNewSession) {
                System.err.println("Configuration token expired or corrupted –
refreshing session and retrying...");
                startNewSession();
                response = appConfigData.getLatestConfiguration(
                    GetLatestConfigurationRequest.builder()

.configurationToken(this.configurationToken)

                                .build());

            } else {
                throw e; // rethrow if it's another kind of bad request
            }
        }

        if (response == null) {

```

```

        // Should not happen, but return cached config if no response
        return this.configuration;
    }

    /**
     * Save the returned configuration token so that it can be passed to the next
     * GetLatestConfiguration API call.
     * Warning: Not persisting this token for use in the next
     * GetLatestConfiguration API call may result in higher
     * than expected usage costs.
     */
    this.configurationToken = response.nextPollConfigurationToken();

    /**
     * If the GetLatestConfiguration API returned configuration data, update the
     * cached configuration with the returned data.
     * Otherwise, assume the configuration has not changed, and return the cached
     * configuration.
     */
    SdkBytes configFromApi = response.configuration();
    if (configFromApi != null && configFromApi.asByteArray().length != 0) {
        this.configuration = configFromApi;
        System.out.println("Configuration contents have changed since the last
        GetLatestConfiguration call, new contents = "
            + this.configuration.asUtf8String());
    } else {
        System.out.println("GetLatestConfiguration returned an empty response
        because we already have the latest configuration");
    }

    return this.configuration;
}

/**
 * Starts a new session with AppConfig and retrieves an initial configuration
 * token.
 */
private void startNewSession() {
    StartConfigurationSessionResponse session =
    appConfigData.startConfigurationSession(req -> req
        .applicationIdentifier("MyDemoApp")
        .configurationProfileIdentifier("MyConfig")
        .environmentIdentifier("Beta"));
    this.configurationToken = session.initialConfigurationToken();
}

```

```
}
}
```

Python

```
# The example below uses two AWS AppConfig Data APIs: StartConfigurationSession and
  GetLatestConfiguration.
# For more information about these APIs, see AWS AppConfig Data.
#
# This class is meant to be used as a singleton to retrieve the latest configuration
  data from AWS AppConfig.
# This class maintains a cache of the latest configuration data in addition to the
  configuration token to be
# passed to the next GetLatestConfiguration API call.
class AppConfigApiRetriever:
    # Set of AppConfig invalid parameter problems that require restarting the
    configuration session.
    # If the GetLatestConfiguration API call fails with any of these problems (e.g.
    token is EXPIRED or CORRUPTED),
    # we need to call StartConfigurationSession again to obtain a new configuration
    token before retrying.
    SESSION_RESTART_REQUIRED = {"EXPIRED", "CORRUPTED"}

    def __init__(self):
        # AWS AppConfig Data SDK client used to interact with the AWS AppConfig Data
        service.
        self.appconfigdata = boto3.client('appconfigdata')

        # The configuration token to be passed to the next GetLatestConfiguration
        API call.
        self.configuration_token = None

        # The cached configuration data to be returned when there is no new
        configuration data available.
        self.configuration = None

    # Returns the latest configuration data stored in AWS AppConfig.
    def get_config(self):
        # If there is no configuration token yet, get one by starting a new session
        with the StartConfigurationSession API.
        # Note that this API does not return configuration data. Rather, it returns
        an initial configuration token that is
        # subsequently passed to the GetLatestConfiguration API.
```

```

    if not self.configuration_token:
        self._start_new_session()

    response = None
    try:
        # Retrieve the configuration from the GetLatestConfiguration API,
        # providing the current configuration token.
        # If this caller does not yet have the latest configuration (e.g. this
        # is the first call to GetLatestConfiguration
        # or new configuration data has been deployed since the first call), the
        # latest configuration data will be returned.
        # Otherwise, the GetLatestConfiguration API will not return any data
        # since the caller already has the latest.
        response = self.appconfigdata.get_latest_configuration(
            ConfigurationToken=self.configuration_token
        )
    except ClientError as e:
        error_code = e.response.get("Error", {}).get("Code")
        # ResourceNotFoundException – usually means the token/session is invalid
        # or expired
        if error_code == "ResourceNotFoundException":
            print("Resource not found – refreshing session and retrying...")
            self._start_new_session()
            response = self.appconfigdata.get_latest_configuration(
                ConfigurationToken=self.configuration_token
            )
        # BadRequestException – check if it's expired or corrupted token
        elif error_code == "BadRequestException":
            details = e.response.get("Error", {}).get("Details", {}) or {}
            invalid_params = details.get("InvalidParameters", {}) or {}
            needs_new_session = any(
                param.get("Problem") in self.SESSION_RESTART_REQUIRED
                for param in invalid_params.values()
            )
            if needs_new_session:
                print("Configuration token expired or corrupted – refreshing
                session and retrying...")
                self._start_new_session()
                response = self.appconfigdata.get_latest_configuration(
                    ConfigurationToken=self.configuration_token
                )
            else:
                raise
        else:

```

```

        raise

    if response is None:
        # Should not happen, but return cached config if no response
        return self.configuration

    # Save the returned configuration token so that it can be passed to the next
    GetLatestConfiguration API call.
    # Warning: Not persisting this token for use in the next
    GetLatestConfiguration API call may result in higher
    # than expected usage costs.
    self.configuration_token = response['NextPollConfigurationToken']

    # If the GetLatestConfiguration API returned configuration data, update the
    cached configuration with the returned data.
    # Otherwise, assume the configuration has not changed, and return the cached
    configuration.
    config_stream = response.get('Configuration')
    if config_stream:
        config_from_api = config_stream.read()
        if config_from_api:
            self.configuration = config_from_api
            print(
                'Configuration contents have changed since the last
    GetLatestConfiguration call, new contents = '
                + self.configuration.decode('utf-8', errors='ignore')
            )
        else:
            print('GetLatestConfiguration returned an empty response because we
    already have the latest configuration')

    return self.configuration

# Starts a new session with AppConfig and retrieves an initial configuration
token.
def _start_new_session(self):
    session = self.appconfigdata.start_configuration_session(
        ApplicationIdentifier='MyDemoApp',
        ConfigurationProfileIdentifier='MyConfig',
        EnvironmentIdentifier='Beta'
    )
    self.configuration_token = session['InitialConfigurationToken']

```

JavaScript

```
/*
The example below uses two AWS AppConfig Data APIs: StartConfigurationSession and
GetLatestConfiguration.
For more information about these APIs, see AWS AppConfig Data
.

This class is meant to be used as a singleton to retrieve the latest configuration
data from AWS AppConfig.
This class maintains a cache of the latest configuration data in addition to the
configuration token to be
passed to the next GetLatestConfiguration API call.
*/
class AppConfigApiRetriever {
  constructor() {
    /*
    AWS AppConfig Data SDK client used to interact with the AWS AppConfig Data
    service.
    */
    this.appconfigdata = new AppConfigDataClient();

    /*
    The configuration token to be passed to the next GetLatestConfiguration API
    call.
    */
    this.configurationToken = null;

    /*
    The cached configuration data to be returned when there is no new configuration
    data available.
    */
    this.configuration = null;
  }

  async startSession() {
    /*
    Starts a new session with the StartConfigurationSession API to get an initial
    configuration token.
    */
    const session = await this.appconfigdata.send(
      new StartConfigurationSessionCommand({
        ApplicationIdentifier: "MyDemoApp",
        ConfigurationProfileIdentifier: "MyConfig",
```

```

        EnvironmentIdentifier: "Beta"
    })
);
this.configurationToken = session.InitialConfigurationToken;
}

/*
Returns the latest configuration data stored in AWS AppConfig.
*/
async getConfig() {
    /*
    If there is no configuration token yet, get one by starting a new session with
    the StartConfigurationSession API.
    Note that this API does not return configuration data. Rather, it returns an
    initial configuration token that is
    subsequently passed to the GetLatestConfiguration API.
    */
    if (!this.configurationToken) {
        await this.startSession();
    }

    let response;
    try {
        /*
        Retrieve the configuration from the GetLatestConfiguration API, providing the
        current configuration token.
        If this caller does not yet have the latest configuration (e.g. this is the
        first call to GetLatestConfiguration
        or new configuration data has been deployed since the first call), the latest
        configuration data will be returned.
        Otherwise, the GetLatestConfiguration API will not return any data since the
        caller already has the latest.
        */
        response = await this.appconfigdata.send(
            new GetLatestConfigurationCommand({
                ConfigurationToken: this.configurationToken
            })
        );
    } catch (err) {
        /*
        Add session restart logic – if the token is invalid or expired, restart the
        session and try once more.
        */

```

```

    if (err.name === "ResourceNotFoundException" || err.name ===
"BadRequestException") {
        console.warn(
            "Configuration token invalid or expired. Restarting session..."
        );
        await this.startSession();
        response = await this.appconfigdata.send(
            new GetLatestConfigurationCommand({
                ConfigurationToken: this.configurationToken
            })
        );
    } else {
        throw err;
    }
}

/*
    Save the returned configuration token so that it can be passed to the next
    GetLatestConfiguration API call.
    Warning: Not persisting this token for use in the next GetLatestConfiguration
    API call may result in higher
    than expected usage costs.
*/
this.configurationToken = response.NextPollConfigurationToken;

/*
    If the GetLatestConfiguration API returned configuration data, update the cached
    configuration with the returned data.
    Otherwise, assume the configuration has not changed, and return the cached
    configuration.
*/
const configFromApi = response.Configuration
    ? await response.Configuration.transformToString()
    : null;

if (configFromApi) {
    this.configuration = configFromApi;
    console.log(
        "Configuration contents have changed since the last GetLatestConfiguration
call, new contents = " +
        this.configuration
    );
} else {
    console.log(

```

```
        "GetLatestConfiguration returned an empty response because we already have
the latest configuration"
    );
}

return this.configuration;
}
}
```

Cleaning up your environment

If you ran one or more of the code samples in this section, we recommend you use one of the following samples to locate and delete the AWS AppConfig resources created by those code samples. The samples in this section call the following APIs:

- [ListApplications](#)
- [DeleteApplication](#)
- [ListEnvironments](#)
- [DeleteEnvironments](#)
- [ListConfigurationProfiles](#)
- [DeleteConfigurationProfile](#)
- [ListHostedConfigurationVersions](#)
- [DeleteHostedConfigurationVersion](#)

Java

```
/*
    This sample provides cleanup code that deletes all the AWS AppConfig resources
    created in the samples above.

    WARNING: this code will permanently delete the given application and all of its
    sub-resources, including
    configuration profiles, hosted configuration versions, and environments. DO NOT
    run this code against
    an application that you may need in the future.
*/

public void cleanUpDemoResources() {
```

```

AppConfigClient appconfig = AppConfigClient.create();

// The name of the application to delete
// IMPORTANT: verify this name corresponds to the application you wish to
delete
String applicationToDelete = "MyDemoApp";

appconfig.listApplicationsPaginator(ListApplicationsRequest.builder().build()).items().forEach(
-> {
    if (app.name().equals(applicationToDelete)) {
        System.out.println("Deleting App: " + app);
        appconfig.listConfigurationProfilesPaginator(req -> req
req.applicationId(app.id())).items().forEach(cp -> {
            System.out.println("Deleting Profile: " + cp);
            appconfig
                .listHostedConfigurationVersionsPaginator(req -> req
                    .applicationId(app.id())
                    .configurationProfileId(cp.id()))
                .items()
                .forEach(hcv -> {
                    System.out.println("Deleting HCV: " + hcv);
                    appconfig.deleteHostedConfigurationVersion(req -> req
                        .applicationId(app.id())
                        .configurationProfileId(cp.id())
                        .versionNumber(hcv.versionNumber()));
                });
            appconfig.deleteConfigurationProfile(req -> req
                .applicationId(app.id())
                .configurationProfileId(cp.id()));
        });

        appconfig.listEnvironmentsPaginator(req -
>req.applicationId(app.id())).items().forEach(env -> {
            System.out.println("Deleting Environment: " + env);
            appconfig.deleteEnvironment(req -
>req.applicationId(app.id()).environmentId(env.id()));
        });

        appconfig.deleteApplication(req -> req.applicationId(app.id()));
    }
});
}

```

Python

```
# this sample provides cleanup code that deletes all the AWS AppConfig resources
# created in the samples above.
#
# WARNING: this code will permanently delete the given application and all of its
# sub-resources, including
#   configuration profiles, hosted configuration versions, and environments. DO NOT
#   run this code against
#   an application that you may need in the future.
#

import boto3

# the name of the application to delete
# IMPORTANT: verify this name corresponds to the application you wish to delete
application_name = 'MyDemoApp'

# create and iterate over a list paginator such that we end up with a list of pages,
# which are themselves lists of applications
# e.g. [ [{'Name': 'MyApp1', ...}, {'Name': 'MyApp2', ...}], [{'Name': 'MyApp3', ...}] ]
list_of_app_lists = [page['Items'] for page in
    appconfig.get_paginator('list_applications').paginate()]
# retrieve the target application from the list of lists
application = [app for apps in list_of_app_lists for app in apps if app['Name'] ==
    application_name][0]
print(f"deleting application {application['Name']} (id={application['Id']})")

# delete all configuration profiles
list_of_config_lists = [page['Items'] for page in
    appconfig.get_paginator('list_configuration_profiles').paginate(ApplicationId=application['Id'])]
for config_profile in [config for configs in list_of_config_lists for config in
    configs]:
    print(f"\tdeleting configuration profile {config_profile['Name']}
    (Id={config_profile['Id']})")

    # delete all hosted configuration versions
    list_of_hcv_lists = [page['Items'] for page in
        appconfig.get_paginator('list_hosted_configuration_versions').paginate(ApplicationId=application['Id'],
        ConfigurationProfileId=config_profile['Id'])]
    for hcv in [hcv for hcvs in list_of_hcv_lists for hcv in hcvs]:

appconfig.delete_hosted_configuration_version(ApplicationId=application['Id'],
    ConfigurationProfileId=config_profile['Id'], VersionNumber=hcv['VersionNumber'])
```

```

        print(f"\t\tdelisted hosted configuration version {hcv['VersionNumber']}")

    # delete the config profile itself
    appconfig.delete_configuration_profile(ApplicationId=application['Id'],
    ConfigurationProfileId=config_profile['Id'])
    print(f"\t\tdelisted configuration profile {config_profile['Name']}
    (Id={config_profile['Id']})")

# delete all environments
list_of_env_lists = [page['Items'] for page in
    appconfig.get_paginator('list_environments').paginate(ApplicationId=application['Id'])]
for environment in [env for envs in list_of_env_lists for env in envs]:
    appconfig.delete_environment(ApplicationId=application['Id'],
    EnvironmentId=environment['Id'])
    print(f"\t\tdelisted environment {environment['Name']} (Id={environment['Id']})")

# delete the application itself
appconfig.delete_application(ApplicationId=application['Id'])
print(f"deleted application {application['Name']} (id={application['Id']})")

```

JavaScript

```

// this sample provides cleanup code that deletes all the AWS AppConfig resources
// created in the samples above.

// WARNING: this code will permanently delete the given application and all of its
// sub-resources, including
// configuration profiles, hosted configuration versions, and environments. DO NOT
// run this code against
// an application that you may need in the future.

import {
    AppConfigClient,
    paginateListApplications,
    DeleteApplicationCommand,
    paginateListConfigurationProfiles,
    DeleteConfigurationProfileCommand,
    paginateListHostedConfigurationVersions,
    DeleteHostedConfigurationVersionCommand,
    paginateListEnvironments,
    DeleteEnvironmentCommand,
} from "@aws-sdk/client-appconfig";

```

```

const client = new AppConfigClient();

// the name of the application to delete
// IMPORTANT: verify this name corresponds to the application you wish to delete
const application_name = "MyDemoApp";

// iterate over all applications, deleting ones that have the name defined above
for await (const app_page of paginateListApplications({ client }, {})) {
  for (const application of app_page.Items) {

    // skip applications that dont have the name thats set
    if (application.Name !== application_name) continue;

    console.log( `deleting application ${application.Name} (id=${application.Id})`);

    // delete all configuration profiles
    for await (const config_page of paginateListConfigurationProfiles({ client },
    { ApplicationId: application.Id }))) {
      for (const config_profile of config_page.Items) {
        console.log(`\tdeleting configuration profile ${config_profile.Name} (Id=
${config_profile.Id})`);

        // delete all hosted configuration versions
        for await (const hosted_page of
        paginateListHostedConfigurationVersions({ client },
        { ApplicationId: application.Id, ConfigurationProfileId:
        config_profile.Id }
        )) {
          for (const hosted_config_version of hosted_page.Items) {
            await client.send(
              new DeleteHostedConfigurationVersionCommand({
                ApplicationId: application.Id,
                ConfigurationProfileId: config_profile.Id,
                VersionNumber: hosted_config_version.VersionNumber,
              })
            );
            console.log(`\t\tdeleted hosted configuration version
${hosted_config_version.VersionNumber}`);
          }
        }

        // delete the config profile itself
        await client.send(
          new DeleteConfigurationProfileCommand({

```

```
        ApplicationId: application.Id,
        ConfigurationProfileId: config_profile.Id,
    })
    );
    console.log(`\tdeleted configuration profile ${config_profile.Name} (Id=
${config_profile.Id})`)
    }

    // delete all environments
    for await (const env_page of paginateListEnvironments({ client },
    { ApplicationId: application.Id }))) {
        for (const environment of env_page.Items) {
            await client.send(
                new DeleteEnvironmentCommand({
                    ApplicationId: application.Id,
                    EnvironmentId: environment.Id,
                })
            );
            console.log(`\tdeleted environment ${environment.Name} (Id=
${environment.Id})`)
        }
    }
}

// delete the application itself
await client.send(
    new DeleteApplicationCommand({ ApplicationId: application.Id })
);
console.log(`deleted application ${application.Name} (id=${application.Id})`)
}
}
```

Configuring AWS AppConfig deletion protection

AWS AppConfig provides an account setting to help prevent users from unintentionally deleting actively-used environments and configuration profiles. AWS AppConfig monitors calls to [GetLatestConfiguration](#) and [GetConfiguration](#) and tracks which configuration profiles and environments have been included in these calls within a 60-minute interval (the default setting). Any configuration profile or environment that was accessed within that interval will be considered active. If you attempt to delete an active configuration profile or environment, AWS AppConfig returns an error. If needed, you can bypass this error by using the `DeletionProtectionCheck` parameter. For more information, see [Bypassing or forcing a deletion protection check](#).

Configure deletion protection using the console

Use the following procedure to configure deletion protection by using the AWS Systems Manager console.

To configure deletion protection (console)

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/appconfig/>.
2. In the navigation pane, choose **Settings**.
3. Use the toggle to enable or disable deletion protection.
4. For **Protection period**, set the definition of an active resource to be between 15 and 1440 minutes.
5. Click **Apply**.

Configure deletion protection using the AWS CLI

Use the following procedure to configure deletion protection by using the AWS CLI. Replace *value* in the following commands with the value you want to use in your environment.

Note

Before you begin, we recommend you update to the latest version of the AWS CLI. For more information, see [Install or update to the latest version of the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

To configure deletion protection (CLI)

1. Run the following command to view the current deletion protection settings.

```
aws appconfig get-account-settings
```

2. Run the following command to enable or disable deletion protection. Specify `false` to disable deletion protection or `true` to enable it.

```
aws appconfig update-account-settings --deletion-protection Enabled=value
```

3. You can increase the default interval to a maximum of 24 hours. Run the following command to specify a new interval.

```
aws appconfig update-account-settings --deletion-protection  
Enabled=true,ProtectionPeriodInMinutes=a number between 15 and 1440
```

Bypassing or forcing a deletion protection check

To help you manage deletion protection, the [DeleteEnvironment](#) and [DeleteConfigurationProfile](#) APIs include a parameter called `DeletionProtectionCheck`. This parameter supports the following values:

- **BYPASS:** Instructs AWS AppConfig to bypass the deletion protection check and delete a configuration profile even if deletion protection would have otherwise prevented it.
- **APPLY:** Instructs the deletion protection check to run, even if deletion protection is disabled at the account level. **APPLY** also forces the deletion protection check to run against resources created in the past hour, which are normally excluded from deletion protection checks.
- **ACCOUNT_DEFAULT:** The default setting, which instructs AWS AppConfig to implement the deletion protection value specified in the `UpdateAccountSettings` API.

Note

By default, `DeletionProtectionCheck` skips configuration profiles and environments created in the past hour. The default configuration is intended to prevent deletion protection from interfering with tests and demos that create short-lived resources. You

can override this behavior by passing `DeletionProtectionCheck=APPLY` when calling `DeleteEnvironment` or `DeleteConfigurationProfile`.

The following CLI walkthrough uses sample commands to illustrate how to use the `DeletionProtectionCheck` parameter. Replace *ID* in the following commands with the ID for your AWS AppConfig artifacts.

1. Call [GetLatestConfiguration](#) on a deployed configuration.

```
aws appconfigdata get-latest-configuration --configuration-token $(aws
  appconfigdata start-configuration-session --application-identifier ID --
  environment-identifier ID --configuration-profile-identifier ID --query
  InitialConfigurationToken) outfile.txt
```

2. Wait 60 seconds for AWS AppConfig to register that the configuration is active.
3. Run the following command to call [DeleteEnvironment](#) and apply deletion protection on the environment.

```
aws appconfig delete-environment --environment-id ID --application-id ID --
  deletion-protection-check APPLY
```

The command should return the following error.

```
An error occurred (BadRequestException) when calling the DeleteEnvironment
  operation: Environment Beta is actively being used in your application and cannot
  be deleted.
```

4. Run the following command to bypass deletion protection and delete the environment.

```
aws appconfig delete-environment --environment-id ID --application-id ID --
  deletion-protection-check BYPASS
```

Security in AWS AppConfig

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that are 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 Systems Manager, 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.

AWS AppConfig is a tool in AWS Systems Manager. To understand how to apply the shared responsibility model when using AWS AppConfig, see [Security in AWS Systems Manager](#). That section describes how to configure Systems Manager to meet the security and compliance objectives for AWS AppConfig.

Implement least privilege access

As a security best practice, grant the minimum required permissions that identities require to perform specific actions on specific resources under specific conditions. AWS AppConfig Agent offers two features that enable the agent to access the filesystem of an instance or container: *backup* and *write to disk*. If you enable these features, verify that only the AWS AppConfig Agent has permissions to write to the designated configuration files on the filesystem. Also verify that only the processes required to read from these configuration files have the ability to do so. Implementing least privilege access is fundamental in reducing security risk and the impact that could result from errors or malicious intent.

For more information about implementing least privilege access, see [SEC03-BP02 Grant least privilege access](#) in the *AWS Well-Architected Tool User Guide*. For more information about the AWS

AppConfig Agent features mentioned in this section, see [Using a manifest to enable additional retrieval features](#).

Data encryption at rest for AWS AppConfig

AWS AppConfig provides encryption by default to protect customer data at rest using AWS owned keys.

AWS owned keys — AWS AppConfig uses these keys by default to automatically encrypt data deployed by the service and hosted in the AWS AppConfig data store. You can't view, manage, or use AWS owned keys, or audit their use. However, you don't have to take any action or change any programs to protect the keys that encrypt your data. For more information, see [AWS owned keys](#) in the *AWS Key Management Service Developer Guide*.

While you can't disable this layer of encryption or select an alternate encryption type, you can specify a customer managed key to be used when you save configuration data hosted in the AWS AppConfig data store and when you deploy your configuration data.

Customer managed keys — AWS AppConfig supports the use of a symmetric customer managed key that you create, own, and manage to add a second layer of encryption over the existing AWS owned key. Because you have full control of this layer of encryption, you can perform such tasks as:

- Establishing and maintaining key policies and grants
- Establishing and maintaining IAM policies
- Enabling and disabling key policies
- Rotating key cryptographic material
- Adding tags
- Creating key aliases
- Scheduling keys for deletion

For more information, see [Customer managed key](#) in the *AWS Key Management Service Developer Guide*.

AWS AppConfig supports customer managed keys

AWS AppConfig offers support for customer managed key encryption for configuration data. For configuration versions saved to the AWS AppConfig hosted data store, customers can set

a `KmsKeyIdentifier` on the corresponding configuration profile. Each time a new version of configuration data is created using the `CreateHostedConfigurationVersion` API operation, AWS AppConfig generates an AWS KMS data key from the `KmsKeyIdentifier` to encrypt the data before storing it. When the data is later accessed, either during the `GetHostedConfigurationVersion` or `StartDeployment` API operations, AWS AppConfig decrypts the configuration data using information about the generated data key.

AWS AppConfig also offers support for customer managed key encryption for deployed configuration data. To encrypt configuration data, customers can provide a `KmsKeyIdentifier` to their deployment. AWS AppConfig generates the AWS KMS data key with this `KmsKeyIdentifier` to encrypt data on the `StartDeployment` API operation.

AWS AppConfig encryption access

When creating a customer managed key, use the following key policy to ensure that the key can be used.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Allow use of the key",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::111122223333:role/role_name"
      },
      "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKey"
      ],
      "Resource": "*"
    }
  ]
}
```

To encrypt hosted configuration data with a customer managed key, the identity calling `CreateHostedConfigurationVersion` needs the following policy statement which can be assigned to a user, group, or role:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "kms:GenerateDataKey",
      "Resource": "arn:aws:kms:us-east-1:111122223333:key/key-ID"
    }
  ]
}
```

If you are using a Secrets Manager secret or any other configuration data encrypted with a customer managed key, your `retrievalRoleArn` will need `kms:Decrypt` to decrypt and retrieve the data.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "kms:Decrypt",
      "Resource": "arn:aws:kms:us-east-1:111122223333:key/key-ID"
    }
  ]
}
```

When calling the AWS AppConfig [StartDeployment](#) API operation, the identity calling `StartDeployment` needs the following IAM policy which can be assigned to a user, group, or role:

JSON

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

```
{
  "Effect": "Allow",
  "Action": [
    "kms:GenerateDataKey*"
  ],
  "Resource": "arn:aws:kms:us-east-1:111122223333:key/key-ID"
}
```

When calling the AWS AppConfig [GetLatestConfiguration](#) API operation, the identity calling `GetLatestConfiguration` needs the following policy which can be assigned to a user, group, or role:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "kms:Decrypt",
      "Resource": "arn:aws:kms:us-east-1:111122223333:key/key-ID"
    }
  ]
}
```

Encryption context

An [encryption context](#) is an optional set of key-value pairs that contain additional contextual information about the data.

AWS KMS uses the encryption context as additional authenticated data to support authenticated encryption. When you include an encryption context in a request to encrypt data, AWS KMS binds the encryption context to the encrypted data. To decrypt data, you include the same encryption context in the request.

AWS AppConfig encryption context: AWS AppConfig uses an encryption context in all AWS KMS cryptographic operations for encrypted hosted configuration data and deployments. The context contains a key corresponding to the type of data and a value that identifies the specific data item.

Monitoring your encryption keys for AWS

When you use an AWS KMS customer managed keys with AWS AppConfig, you can use AWS CloudTrail or Amazon CloudWatch Logs to track requests that AWS AppConfig sends to AWS KMS.

The following example is a CloudTrail event for Decrypt to monitor AWS KMS operations called by AWS AppConfig to access data encrypted by your customer managed key:

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AWSService",
    "invokedBy": "appconfig.amazonaws.com"
  },
  "eventTime": "2023-01-03T02:22:28z",
  "eventSource": "kms.amazonaws.com",
  "eventName": "Decrypt",
  "awsRegion": "Region",
  "sourceIPAddress": "172.12.34.56",
  "userAgent": "ExampleDesktop/1.0 (V1; OS)",
  "requestParameters": {
    "encryptionContext": {
      "aws:appconfig:deployment:arn":
"arn:aws:appconfig:Region:account_ID:application/application_ID/
environment/environment_ID/deployment/deployment_ID"
    },
    "keyId": "arn:aws:kms:Region:account_ID:key/key_ID",
    "encryptionAlgorithm": "SYMMETRIC_DEFAULT"
  },
  "responseElements": null,
  "requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
  "eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
  "readOnly": true,
  "resources": [
    {
      "accountId": "account_ID",
      "type": "AWS::KMS::Key",
      "ARN": "arn:aws:kms:Region:account_ID:key_ID"
    }
  ]
}
```

```
],
"eventType": "AwsApiCall",
"managementEvent": true,
"eventCategory": "Management",
"recipientAccountId": "account_ID",
"sharedEventID": "dc129381-1d94-49bd-b522-f56a3482d088"
}
```

Access AWS AppConfig using an interface endpoint (AWS PrivateLink)

You can use AWS PrivateLink to create a private connection between your VPC and AWS AppConfig. You can access AWS AppConfig as if it were in your VPC, without the use of an internet gateway, NAT device, VPN connection, or Direct Connect connection. Instances in your VPC don't need public IP addresses to access AWS AppConfig.

You establish this private connection by creating an *interface endpoint*, powered by AWS PrivateLink. We create an endpoint network interface in each subnet that you enable for the interface endpoint. These are requester-managed network interfaces that serve as the entry point for traffic destined for AWS AppConfig.

For more information, see [Access AWS services through AWS PrivateLink](#) in the *AWS PrivateLink Guide*.

Considerations for AWS AppConfig

Before you set up an interface endpoint for AWS AppConfig, review [Considerations](#) in the *AWS PrivateLink Guide*.

AWS AppConfig supports making calls to the [appconfig](#) and [appconfigdata](#) services through the interface endpoint.

Create an interface endpoint for AWS AppConfig

You can create an interface endpoint for AWS AppConfig using either the Amazon VPC console or the AWS Command Line Interface (AWS CLI). For more information, see [Create an interface endpoint](#) in the *AWS PrivateLink Guide*.

Create an interface endpoint for AWS AppConfig using the following service names:

```
com.amazonaws.region.appconfig
```

```
com.amazonaws.region.appconfigdata
```

If you enable private DNS for the interface endpoint, you can make API requests to AWS AppConfig using its default Regional DNS name. For example, `appconfig.us-east-1.amazonaws.com` and `appconfigdata.us-east-1.amazonaws.com`.

Create an endpoint policy for your interface endpoint

An endpoint policy is an IAM resource that you can attach to an interface endpoint. The default endpoint policy allows full access to AWS AppConfig through the interface endpoint. To control the access allowed to AWS AppConfig from your VPC, attach a custom endpoint policy to the interface endpoint.

An endpoint policy specifies the following information:

- The principals that can perform actions (AWS accounts, IAM users, and IAM roles).
- The actions that can be performed.
- The resources on which the actions can be performed.

For more information, see [Control access to services using endpoint policies](#) in the *AWS PrivateLink Guide*.

Example: VPC endpoint policy for AWS AppConfig actions

The following is an example of a custom endpoint policy. When you attach this policy to your interface endpoint, it grants access to the listed AWS AppConfig actions for all principals on all resources.

```
{
  "Statement": [
    {
      "Principal": "*",
      "Effect": "Allow",
      "Action": [
        "appconfig:CreateApplication",
        "appconfig:CreateEnvironment",
        "appconfig:CreateConfigurationProfile",
```

```
        "appconfig:StartDeployment",
        "appconfig:GetLatestConfiguration"
        "appconfig:StartConfigurationSession"
    ],
    "Resource": "*"
}
]
```

Secrets Manager key rotation

This section describes important security information about AWS AppConfig integration with Secrets Manager. For information about Secrets Manager, see [What is AWS Secrets Manager?](#) in the *AWS Secrets Manager User Guide*.

Setting up automatic rotation of Secrets Manager secrets deployed by AWS AppConfig

Rotation is the process of periodically updating a secret stored in Secrets Manager. When you rotate a secret, you update the credentials in both the secret and the database or service. You can configure automatic secrets rotation in Secrets Manager by using an AWS Lambda function to update the secret and the database. For more information, see [Rotate AWS Secrets Manager secrets](#) in the *AWS Secrets Manager User Guide*.

To enable key rotation of Secrets Manager secrets deployed by AWS AppConfig, update your rotation Lambda function and deploy the rotated secret.

Note

Deploy your AWS AppConfig configuration profile after your secret has been rotated and fully updated to the new version. You can determine if the secret rotated because the status of VersionStage changes from AWSPENDING to AWSCURRENT. Secret rotation completion occurs within the Secrets Manager Rotation Templates `finish_secret` function.

Here is an example function that starts an AWS AppConfig deployment after a secret is rotated.

```
import time
```

```

import boto3
client = boto3.client('appconfig')

def finish_secret(service_client, arn, new_version):
    """Finish the rotation by marking the pending secret as current
    This method finishes the secret rotation by staging the secret staged AWSPENDING
    with the AWSCURRENT stage.
    Args:
        service_client (client): The secrets manager service client
        arn (string): The secret ARN or other identifier
        new_version (string): The new version to be associated with the secret
    """
    # First describe the secret to get the current version
    metadata = service_client.describe_secret(SecretId=arn)
    current_version = None
    for version in metadata["VersionIdsToStages"]:
        if "AWSCURRENT" in metadata["VersionIdsToStages"][version]:
            if version == new_version:
                # The correct version is already marked as current, return
                logger.info("finishSecret: Version %s already marked as AWSCURRENT for
%s" % (version, arn))
                return
            current_version = version
            break

    # Finalize by staging the secret version current
    service_client.update_secret_version_stage(SecretId=arn, VersionStage="AWSCURRENT",
MoveToVersionId=new_version, RemoveFromVersionId=current_version)

    # Deploy rotated secret
    response = client.start_deployment(
        ApplicationId='TestApp',
        EnvironmentId='TestEnvironment',
        DeploymentStrategyId='TestStrategy',
        ConfigurationProfileId='ConfigurationProfileId',
        ConfigurationVersion=new_version,
        KmsKeyIdentifier=key,
        Description='Deploy secret rotated at ' + str(time.time())
    )

    logger.info("finishSecret: Successfully set AWSCURRENT stage to version %s for
secret %s." % (new_version, arn))

```

Monitoring AWS AppConfig

Monitoring is an important part of maintaining the reliability, availability, and performance of AWS AppConfig and your other AWS solutions. AWS provides the following monitoring tools to watch AWS AppConfig, report when something is wrong, and take automatic actions when appropriate:

- *Amazon CloudWatch* monitors your AWS resources and the applications 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 EC2 instances and automatically launch new instances when needed. For more information, see the [Amazon CloudWatch 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](#).
- *Amazon CloudWatch Logs* enables you to monitor, store, and access your log files from Amazon EC2 instances, CloudTrail, and other sources. CloudWatch Logs can monitor information in the log files and notify you when certain thresholds are met. You can also archive your log data in highly durable storage. For more information, see the [Amazon CloudWatch Logs User Guide](#).
- *Amazon EventBridge* can be used to automate your AWS services and respond automatically to system events, such as application availability issues or resource changes. Events from AWS services are delivered to EventBridge in near real time. You can write simple rules to indicate which events are of interest to you and which automated actions to take when an event matches a rule. For more information, see [Amazon EventBridge User Guide](#).

Topics

- [Logging AWS AppConfig API calls using AWS CloudTrail](#)
- [Logging metrics for AWS AppConfig data plane calls](#)
- [Monitoring deployments for automatic rollback](#)

Logging AWS AppConfig API calls using AWS CloudTrail

AWS AppConfig is integrated with [AWS CloudTrail](#), a service that provides a record of actions taken by a user, role, or an AWS service. CloudTrail captures all API calls for AWS AppConfig as events. The calls captured include calls from the AWS AppConfig console and code calls to the AWS AppConfig API operations. Using the information collected by CloudTrail, you can determine the request that was made to AWS AppConfig, the IP address from which the request was made, when it was made, and additional details.

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 user or user credentials.
- Whether the request was made on behalf of an IAM Identity Center user.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

CloudTrail is active in your AWS account when you create the account and you automatically have access to the CloudTrail **Event history**. The CloudTrail **Event history** provides a viewable, searchable, downloadable, and immutable record of the past 90 days of recorded management events in an AWS Region. For more information, see [Working with CloudTrail Event history](#) in the *AWS CloudTrail User Guide*. There are no CloudTrail charges for viewing the **Event history**.

For an ongoing record of events in your AWS account past 90 days, create a trail or a [CloudTrail Lake](#) event data store.

CloudTrail trails

A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. All trails created using the AWS Management Console are multi-Region. You can create a single-Region or a multi-Region trail by using the AWS CLI. Creating a multi-Region trail is recommended because you capture activity in all AWS Regions in your account. If you create a single-Region trail, you can view only the events logged in the trail's AWS Region. For more information about trails, see [Creating a trail for your AWS account](#) and [Creating a trail for an organization](#) in the *AWS CloudTrail User Guide*.

You can deliver one copy of your ongoing management events to your Amazon S3 bucket at no charge from CloudTrail by creating a trail, however, there are Amazon S3 storage charges. For

more information about CloudTrail pricing, see [AWS CloudTrail Pricing](#). For information about Amazon S3 pricing, see [Amazon S3 Pricing](#).

CloudTrail Lake event data stores

CloudTrail Lake lets you run SQL-based queries on your events. CloudTrail Lake converts existing events in row-based JSON format to [Apache ORC](#) format. ORC is a columnar storage format that is optimized for fast retrieval of data. Events are aggregated into *event data stores*, which are immutable collections of events based on criteria that you select by applying [advanced event selectors](#). The selectors that you apply to an event data store control which events persist and are available for you to query. For more information about CloudTrail Lake, see [Working with AWS CloudTrail Lake](#) in the *AWS CloudTrail User Guide*.

CloudTrail Lake event data stores and queries incur costs. When you create an event data store, you choose the [pricing option](#) you want to use for the event data store. The pricing option determines the cost for ingesting and storing events, and the default and maximum retention period for the event data store. For more information about CloudTrail pricing, see [AWS CloudTrail Pricing](#).

AWS AppConfig data events in CloudTrail

[Data events](#) provide information about the resource operations performed on or in a resource (for example, retrieving the latest deployed configuration by calling `GetLatestConfiguration`). These are also known as data plane operations. Data events are often high-volume activities. By default, CloudTrail doesn't log data events. The CloudTrail **Event history** doesn't record data events.

Additional charges apply for data events. For more information about CloudTrail pricing, see [AWS CloudTrail Pricing](#).

You can log data events for the AWS AppConfig resource types by using the CloudTrail console, AWS CLI, or CloudTrail API operations. The [table](#) in this section shows the resource types available for AWS AppConfig.

- To log data events using the CloudTrail console, create a [trail](#) or [event data store](#) to log data events, or [update an existing trail or event data store](#) to log data events.
 1. Choose **Data events** to log data events.
 2. From the **Data event type** list, choose **AWS AppConfig**.

3. Choose the log selector template you want to use. You can log all data events for the resource type, log all `readOnly` events, log all `writeOnly` events, or create a custom log selector template to filter on the `readOnly`, `eventName`, and `resources.ARN` fields.
 4. For **Selector name**, enter **AppConfigDataEvents**. For information about enabling Amazon CloudWatch Logs for your data event trail, see [Logging metrics for AWS AppConfig data plane calls](#).
- To log data events using the AWS CLI, configure the `--advanced-event-selectors` parameter to set the `eventCategory` field equal to `Data` and the `resources.type` field equal to the resource type value (see [table](#)). You can add conditions to filter on the values of the `readOnly`, `eventName`, and `resources.ARN` fields.
 - To configure a trail to log data events, run the [put-event-selectors](#) command. For more information, see [Logging data events for trails with the AWS CLI](#).
 - To configure an event data store to log data events, run the [create-event-data-store](#) command to create a new event data store to log data events, or run the [update-event-data-store](#) command to update an existing event data store. For more information, see [Logging data events for event data stores with the AWS CLI](#).

The following table lists the AWS AppConfig resource types. The **Data event type (console)** column shows the value to choose from the **Data event type** list on the CloudTrail console. The **resources.type value** column shows the `resources.type` value, which you would specify when configuring advanced event selectors using the AWS CLI or CloudTrail APIs. The **Data APIs logged to CloudTrail** column shows the API calls logged to CloudTrail for the resource type.

Data event type (console)	resources.type value	Data APIs logged to CloudTrail*
AWS AppConfig	<code>AWS::AppConfig::Configuration</code>	• GetLatestConfiguration • StartConfigurationSession

*You can configure advanced event selectors to filter on the `eventName`, `readOnly`, and `resources.ARN` fields to log only those events that are important to you. For more information about these fields, see [AdvancedFieldSelector](#).

AWS AppConfig management events in CloudTrail

AWS AppConfig logs all AWS AppConfig control plane operations as management events. For a list of the AWS AppConfig control plane operations that AWS AppConfig logs to CloudTrail, see the [AWS AppConfig API Reference](#).

AWS AppConfig event examples

An event represents a single request from any source and includes information about the requested API 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 events don't appear in any specific order.

The following example shows a CloudTrail event that demonstrates the [StartConfigurationSession](#) operation.

```
{
  "eventVersion": "1.09",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:iam::123456789012:user/Administrator",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "sessionIssuer": {},
      "attributes": {
        "creationDate": "2024-01-11T14:37:02Z",
        "mfaAuthenticated": "false"
      }
    }
  },
  "eventTime": "2024-01-11T14:45:15Z",
  "eventSource": "appconfig.amazonaws.com",
  "eventName": "StartConfigurationSession",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "Boto3/1.34.11 md/Botocore#1.34.11 ua/2.0 os/macos#22.6.0
md/arch#x86_64 lang/python#3.11.4 md/pyimpl#CPython cfg/retry-mode#legacy
Botocore/1.34.11",
  "requestParameters": {
    "applicationIdentifier": "rrfexample",
```

```

    "environmentIdentifier": "mexampleqe0",
    "configurationProfileIdentifier": "3eexampleu1"
  },
  "responseElements": null,
  "requestID": "a1b2c3d4-5678-90ab-cdef-aaaaaEXAMPLE",
  "eventID": "a1b2c3d4-5678-90ab-cdef-bbbbbbEXAMPLE",
  "readOnly": false,
  "resources": [
    {
      "accountId": "123456789012",
      "type": "AWS::AppConfig::Configuration",
      "ARN": "arn:aws:appconfig:us-east-1:123456789012:application/rrfexample/
environment/mexampleqe0/configuration/3eexampleu1"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": false,
  "recipientAccountId": "123456789012",
  "eventCategory": "Data",
  "tlsDetails": {
    "tlsVersion": "TLSv1.3",
    "cipherSuite": "TLS_AES_128_GCM_SHA256",
    "clientProvidedHostHeader": "appconfigdata.us-east-1.amazonaws.com"
  }
}

```

For information about CloudTrail record contents, see [CloudTrail record contents](#) in the *AWS CloudTrail User Guide*.

Logging metrics for AWS AppConfig data plane calls

If you configured AWS CloudTrail to log AWS AppConfig data events, you can enable Amazon CloudWatch Logs to log metrics for calls to the AWS AppConfig data plane. You can then search and filter log data in CloudWatch Logs by creating one or more metric filters. Metric filters define the terms and patterns to look for in log data as it is sent to CloudWatch Logs. CloudWatch Logs uses metric filters to turn log data into numerical CloudWatch metrics. You can graph metrics or configure them with an alarm.

Before you begin

Enable logging of AWS AppConfig data events in AWS CloudTrail. The following procedure describes how to enable metric logging for an *existing AWS AppConfig trail* in CloudTrail. For

information about how to enable CloudTrail logging for AWS AppConfig data plan calls, see [AWS AppConfig data events in CloudTrail](#).

Use the following procedure to enable CloudWatch Logs to log metrics for calls to the AWS AppConfig data plane.

To enable CloudWatch Logs to log metrics for calls to the AWS AppConfig data plane

1. Open the CloudTrail console at <https://console.aws.amazon.com/cloudtrail/>.
2. On the dashboard, choose your AWS AppConfig trail.
3. In the **CloudWatch Logs** section, choose **Edit**.
4. Choose **Enabled**.
5. For **Log group name**, either leave the default name or enter a name. Make a note of the name. You will choose the log group in the CloudWatch Logs console later.
6. For **Role name**, enter a name.
7. Choose **Save changes**.

Use the following procedure to create a metric and a metric filter for AWS AppConfig in CloudWatch Logs. The procedure describes how to create a metric filter for calls by operation and (optionally) calls by operation and Amazon Resource Name (ARN).

To create a metric and a metric filter for AWS AppConfig in CloudWatch Logs

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Logs**, and then choose **Log groups**.
3. Choose the checkbox beside the AWS AppConfig log group.
4. Choose **Actions**, and then choose **Create metric filter**.
5. For **Filter name**, enter a name.
6. For **Filter pattern**, enter the following:

```
{ $.eventSource = "appconfig.amazonaws.com" }
```

7. (Optional) In the **Test pattern** section, choose your log group from the **Select log data to test** list. If CloudTrail hasn't logged any calls, you can skip this step.
8. Choose **Next**.

9. For **Metric namespace**, enter **AWS AppConfig**.
10. For **Metric name**, enter **Calls**.
11. For **Metric value**, enter **1**.
12. Skip **Default value** and **Unit**.
13. For **Dimension name**, enter **operation**.
14. For **Dimension value**, enter **\$.eventName**.

(Optional) You can enter a second dimension that includes the Amazon Resource Name (ARN) making the call. To add a second dimension, for **Dimension name**, enter **resource**. For **Dimension value**, enter **\$.resources[0].ARN**.

Choose **Next**.

15. Review the details of the filter and **Create metric filter**.

(Optional) You can repeat this procedure to create a new metric filter for a specific error code like *AccessDenied*. If you do, enter the following details:

1. For **Filter name**, enter a name.
2. For **Filter pattern**, enter the following:

```
{ $.errorCode = "codename" }
```

For example

```
{ $.errorCode = "AccessDenied" }
```

3. For **Metric namespace**, enter **AWS AppConfig**.
4. For **Metric name**, enter **Errors**.
5. For **Metric value**, enter **1**.
6. For **Default value**, enter a zero (0).
7. Skip **Unit**, **Dimensions**, and **Alarms**.

After CloudTrail logs API calls, you can view metrics in CloudWatch. For more information, see [Viewing your metrics and logs in the console](#) in the *Amazon CloudWatch User Guide*. For information about how to locate a metric you created, see [Search for available metrics](#).

Note

If you set up the error metric with no dimension, as described here, you can view those metrics on the **Metrics with no dimension** page.

Creating an alarm for a CloudWatch metric

After you create metrics, you can create metric alarms in CloudWatch. For example, you can create an alarm for the *AWS AppConfig calls* metric you created in the previous procedure. Specifically, you can create an alarm for calls to the AWS AppConfig `StartConfigurationSession` API action that surpass a threshold. For information about how to create an alarm for a metric, see [Create a CloudWatch alarm based on a static threshold](#) in the *Amazon CloudWatch User Guide*. For information about default limits for calls to the AWS AppConfig data plane, see [Data plane default limits](#) in the *Amazon Web Services General Reference*.

Monitoring deployments for automatic rollback

During a deployment, you can mitigate situations where malformed or incorrect configuration data causes errors in your application by using a combination of AWS AppConfig [deployment strategies](#) and automatic rollbacks based on Amazon CloudWatch alarms. Once configured, if one or more CloudWatch alarms go into the `ALARM` or `INSUFFICIENT_DATA` state during a deployment, AWS AppConfig automatically rolls back your configuration data to the previous version, thereby preventing application outages or errors.

Note

A deployment doesn't automatically roll back if actions have been disabled in an associated CloudWatch alarm.

You can disable and enable alarms by using the [DisableAlarmActions](#) and [EnableAlarmActions](#) API actions, or the [disable-alarm-actions](#) and [enable-alarm-actions](#) commands in the AWS CLI.

You can also roll back a configuration by calling the [StopDeployment](#) API operation while a deployment is still in progress.

Important

For deployments that successfully complete, AWS AppConfig also supports reverting configuration data to a previous version by using the `AllowRevert` parameter with the [StopDeployment](#) API operation. For some customers, reverting to a previous configuration after a successful deployment guarantees the data will be the same as it was before the deployment. Reverting also ignores alarm monitors, which may prevent a roll forward from progressing during an application emergency. For more information, see [Reverting a configuration](#).

To configure automatic rollbacks, you specify the Amazon Resource Name (ARN) of one or more CloudWatch metrics in the **CloudWatch alarms** field when you create (or edit) an AWS AppConfig environment. For more information, see [Creating environments for your application in AWS AppConfig](#).

Note

If you use a third-party monitoring solution (for example, Datadog or New Relic), you can create an AWS AppConfig extension that checks for alarms at the `AT_DEPLOYMENT_TICK` action point and, as a safety guardrail, rolls back the deployment if it triggered an alarm. For more information, see the following Datadog and New Relic integration examples on GitHub:

- [Datadog](#)
- [New Relic](#)

For more information about AWS AppConfig extensions, see the following topics:

- [Extending AWS AppConfig workflows using extensions](#)
- [Walkthrough: Creating custom AWS AppConfig extensions](#)

Recommended metrics to monitor for automatic rollback

The metrics you choose to monitor will depend on the hardware and software used by your applications. AWS AppConfig customers often monitor the following metrics. For a complete

list of recommended metrics grouped by AWS service, see [Recommended alarms](#) in the *Amazon CloudWatch User Guide*.

After you determine the metrics you want to monitor, use CloudWatch to configure alarms. For more information, see [Using Amazon CloudWatch alarms](#).

Service	Metric	Details
Amazon API Gateway	4XXError	This alarm detects a high rate of client-side errors. This can indicate an issue in the authorization or client request parameters. It could also mean that a resource was removed or a client is requesting one that doesn't exist. Consider enabling Amazon CloudWatch Logs and checking for any errors that may be causing the 4XX errors. Moreover, consider enabling detailed CloudWatch metrics to view this metric per resource and method and narrow down the source of the errors. Errors could also be caused by exceeding the configured throttling limit.
Amazon API Gateway	5XXError	This alarm helps to detect a high rate of server-side errors. This can indicate that there is something wrong on the API backend, the network, or the integration between the API gateway and the backend API.

Service	Metric	Details
Amazon API Gateway	Latency	This alarm detects high latency in a stage. Find the <code>IntegrationLatency</code> metric value to check the API backend latency. If the two metrics are mostly aligned, the API backend is the source of higher latency and you should investigate there for issues. Consider also enabling CloudWatch Logs and checking for errors that might be causing the high latency.
Amazon EC2 Auto Scaling	<code>GroupInServiceCapacity</code>	This alarm helps to detect when the capacity in the group is below the desired capacity required for your workload. To troubleshoot, check your scaling activities for launch failures and confirm that your desired capacity configuration is correct.

Service	Metric	Details
Amazon EC2	CPUUtilization	This alarm helps to monitor the CPU utilization of an EC2 instance. Depending on the application, consistently high utilization levels might be normal. But if performance is degraded, and the application is not constrained by disk I/O, memory, or network resources, then a maxed-out CPU might indicate a resource bottleneck or application performance problems.
Amazon ECS	CPUReservation	This alarm helps you detect a high CPU reservation of the ECS cluster. High CPU reservation might indicate that the cluster is running out of registered CPUs for the task.
Amazon ECS	HTTPCode_Target_5XX_Count	This alarm helps you detect a high server-side error count for the ECS service. This can indicate that there are errors that cause the server to be unable to serve requests.

Service	Metric	Details
Amazon EKS with Container insights	node_cpu_utilization	This alarm helps to detect high CPU utilization in worker nodes of the Amazon EKS cluster. If the utilization is consistently high, it might indicate a need for replacing your worker nodes with instances that have greater CPU or a need to scale the system horizontally.
Amazon EKS with Container insights	node_memory_utilization	This alarm helps in detecting high memory utilization in worker nodes of the Amazon EKS cluster. If the utilization is consistently high, it might indicate a need to scale the number of pod replicas, or optimize your application.
Amazon EKS with Container insights	pod_cpu_utilization_over_pod_limit	This alarm helps in detecting high CPU utilization in pods of the Amazon EKS cluster. If the utilization is consistently high, it might indicate a need to increase the CPU limit for the affected pod.
Amazon EKS with Container insights	pod_memory_utilization_over_pod_limit	This alarm helps in detecting high CPU utilization in pods of the Amazon EKS cluster. If the utilization is consistently high, it might indicate a need to increase the CPU limit for the affected pod.

Service	Metric	Details
AWS Lambda	Errors	This alarm detects high error counts. Errors includes the exceptions thrown by the code as well as exceptions thrown by the Lambda runtime.
AWS Lambda	Throttles	This alarm detects a high number of throttled invocation requests. Throttling occurs when there is no concurrency is available for scale up.
Lambda Insights	memory_utilization	This alarm is used to detect if the memory utilization of a lambda function is approaching the configured limit.
Amazon S3	4xxErrors	This alarm helps us report the total number of 4xx error status codes that are made in response to client requests. 403 error codes might indicate an incorrect IAM policy, and 404 error codes might indicate misbehaving client application, for example.

Service	Metric	Details
Amazon S3	5xxErrors	This alarm helps you detect a high number of server-side errors. These errors indicate that a client made a request that the server couldn't complete. This can help you correlate the issue your application is facing because of S3.

AWS AppConfig User Guide document history

The following table describes the important changes to the documentation since the last release of AWS AppConfig.

Current API version: 2019-10-09

Change	Description	Date
New version of AWS AppConfig Agent Lambda extension	The agent has been updated with minor enhancements and bug fixes. To view the new Amazon Resource Names (ARNs) for the extension, see Available versions of the AWS AppConfig Agent Lambda extension .	August 24, 2026
AWS AppConfig experimentation assignment logging and treatment retrieval	AWS AppConfig Agent can now log experiment treatment assignment records to stderr or to disk, and emits aggregated treatment traffic as CloudWatch metrics. New documentation describes the EXPERIMENT_ASSIGNMENT_LOG_DESTINATION agent option, how to retrieve an assigned treatment by using the Entity-Id and Context headers, and the minimum agent versions that support experimentation. For more information, see About data collection .	August 21, 2026

[New version of AWS AppConfig Agent Lambda extension](#)

The agent has been updated with minor enhancements and bug fixes. To view the new Amazon Resource Names (ARNs) for the extension, see [Available versions of the AWS AppConfig Agent Lambda extension](#).

June 22, 2026

[Launch of AWS AppConfig experimentation](#)

June 12, 2026

Application experimentation is the systematic process of testing hypotheses on software applications using A/B testing, feature flags, and analytics to measure user behavior, optimize performance, and drive data-driven product decisions. Experimentation enables teams to validate changes, such as new features or UI changes, before a full rollout, thereby reducing risk and improving engagement.

Built with the expertise of more than 20 years of software experimentation on Amazon.com, AWS AppConfig experimentation integrates seamlessly with your existing AWS AppConfig workflows, providing a comprehensive solution for feature management and experimentation in cloud-native applications. AWS AppConfig experimentation offers the following features:

- Support for multiple feature and test types – A/B tests, multivariate tests, and gradual feature rollouts.

- Easy experiment setup and management through the AWS Management Console.
- Powerful and flexible segmentation options for experiment participants.
- Data independence – use your existing data warehouses and analytics tools for experiments or use an AWS-native service like [Amazon CloudWatch RUM](#).

Additionally, AWS AppConfig experimentation includes automatic rollback capabilities that can revert to a safe configuration if predefined thresholds are exceeded. For example, you can monitor key metrics like page load, conversion rate, or page errors during an experiment. If a threshold is breached, the experiment will be stopped. You can set these safety monitoring capabilities using Amazon CloudWatch, Datadog, New Relic, Dynatrace, and more. For more information, see [AWS AppConfig experimentation](#).

[Entity-based targeting](#)

AWS AppConfig Agent (version 2.0.136060 or later) supports deploying feature flag or free-form configuration data to specific segments or individual users during a gradual rollout. Entity-based gradual deployments ensure that once a user or segment receives a configuration version, they continue to receive that same version throughout the deployment period, regardless of which compute resource serves their requests. For more information, see [Using AWS AppConfig Agent for user- or entity-based gradual deployments](#).

March 26, 2026

[New version of AWS AppConfig Agent Lambda extension](#)

The agent has been updated with minor enhancements and bug fixes. To view the new Amazon Resource Names (ARNs) for the extension, see [Available versions of the AWS AppConfig Agent Lambda extension](#).

February 20, 2026

[Deprecated: AWS AppConfig extension for CloudWatch Evidently](#)

The AWS AppConfig extension for CloudWatch Evidently is no longer supported.

February 20, 2026

[IPv6 support](#)

All AWS AppConfig APIs now fully support IPv4 and IPv6 calls. For more information, see [Understanding IPv6 support](#).

April 23, 2025

[New topic: Saving a previous feature flag version to a new version](#)

When you update a feature flag, AWS AppConfig automatically saves your changes to a new version. If you want to use a previous feature flag version, you must copy it to a draft version and then save it. You can't edit and save changes to a previous flag version without saving it to a new version. For more information, see [Saving a previous feature flag version to a new version](#).

April 15, 2025

[New topic: Feature flag samples for AWS AppConfig Agent local development mode](#)

AWS AppConfig Agent supports a [local development mode](#). If you enable local development mode, the agent reads configuration data from a specified directory on disk. It doesn't retrieve configuration data from AWS AppConfig. To help you better understand how to use local development mode, this guide now contains a topic with feature flag samples. For more information, see [Feature flag samples for AWS AppConfig Agent local development mode](#).

February 18, 2025

[New topic: Creating a configuration profile for non-native data sources](#)

The topic describes the high level process for using an AWS AppConfig extension to retrieve configuration data from sources that are not natively supported, including other AWS services like Amazon RDS and Amazon DynamoDB, as well as 3rd-party sources like GitHub, GitLab, or a local repo. For more information, see [Creating a configuration profile for non-native data sources](#)

December 19, 2024

[Updated topic: Fixed regex in feature flags type reference](#)

The json schema in the feature flag type reference previously showed the following regex pattern in various places: `"^[a-z][a-zA-Z\\d\\d-]{0,63}$"`. The correct regex pattern is `"^[a-z][a-zA-Z\\d_-]{0,63}$"`. The hyphen is listed after the underscore. For more information, see [Understanding the type reference for AWS.AppConfig.FeatureFlags](#)

December 18, 2024

[Updated topics: Environment variable samples added](#)

The tables describing environment variable in the following topics were updated to include samples:

December 12, 2024

- [\(Optional\) Using environment variables to configure AWS AppConfig Agent for Amazon ECS and Amazon EKS](#)
- [\(Optional\) Using environment variables to configure AWS AppConfig Agent for Amazon EC2](#)
- [Configuring the AWS AppConfig Agent Lambda extension](#)

[New section: Understanding the split operator](#)

A new section uses examples to explain how the `split` operator works for a multi-variant feature flag rule. For more information, see [Understanding multi-variant feature flag rules](#).

November 22, 2024

[New extension action point: AT_DEPLOYMENT_TICK](#)

November 22, 2024

AWS AppConfig launched a new action point for users who create custom extensions. The `AT_DEPLOYMENT_TICK` action point supports third-party monitoring integration. `AT_DEPLOYMENT_TICK` is invoked during configuration deployment processing orchestration. If you use a third-party monitoring solution (for example, Datadog), you can create an AWS AppConfig extension that checks for alarms at the `AT_DEPLOYMENT_TICK` action point and, as a safety guardrail, rolls back the deployment if it triggered an alarm. For more information about AWS AppConfig extensions, see [Extending AWS AppConfig workflows using extensions](#). For more information about custom extensions, see [Walkthrough: Creating custom AWS AppConfig extensions](#). To view a code sample of an AWS AppConfig extension that uses the `AT_DEPLOYMENT_TICK` action point to integrate with Datadog, see [aws-samples / aws-appco](#)

[nfig-tick-extn-for-datadog](#) on GitHub.

[New topic: AWS AppConfig mobile use considerations](#)

A new topic in this guide describes important considerations for using AWS AppConfig feature flags with mobile devices. For more information, see [AWS AppConfig mobile use considerations](#).

November 21, 2024

[New feature: AWS AppConfig Deletion Protection](#)

AWS AppConfig now provides an account setting to help prevent users from unintentionally deleting actively-used environments and configuration profiles. For more information, see [Configuring AWS AppConfig deletion protection](#).

August 28, 2024

[New version of AWS AppConfig Agent Lambda extension](#)

The agent has been updated with minor enhancements and bug fixes. To view the new Amazon Resource Names (ARNs) for the extension, see [Available versions of the AWS AppConfig Agent Lambda extension](#).

August 9, 2024

[New code samples for retrieving flag variants](#)

For more information, see [Using AWS AppConfig Agent to retrieve a feature flag with variants](#).

August 9, 2024

[New version of AWS AppConfig Agent Lambda extension](#)

The agent has been updated to support feature flag targets, variants, and splits. To view the new Amazon Resource Names (ARNs) for the extension, see [Available versions of the AWS AppConfig Agent Lambda extension](#).

July 23, 2024

[New feature: Multi-variant feature flags](#)

Multi-variant feature flags enable you to define a set of possible flag values to return for a request. You can also configure different statuses (enabled or disabled) for multi-variant flags. When requesting a flag configured with variants, your application provides context that AWS AppConfig evaluates against a set of user-defined rules. Depending on the context specified in the request and the rules defined for the variant, AWS AppConfig returns different flag values to the application. For more information, see [Creating multi-variant feature flags](#).

July 23, 2024

[New version of AWS AppConfig Agent Lambda extension](#)

The agent has been updated with minor enhancements and bug fixes. To view the new Amazon Resource Names (ARNs) for the extension, see [Available versions of the AWS AppConfig Agent Lambda extension](#).

February 28, 2024

[AWS AppConfig custom extension samples](#)

The [Walkthrough: Creating custom AWS AppConfig extensions](#) topic now includes links to the following sample extensions on GitHub:

February 28, 2024

- [Sample extension that prevents deployments with a blocked day moratorium calendar using Systems Manager Change Calendar](#)
- [Sample extension that prevents secrets from leaking into configuration data using git-secrets](#)
- [Sample extension that prevents personally identifiable information \(PII\) from leaking into configuration data using Amazon Comprehend](#)

[New topic: Logging AWS AppConfig API calls using AWS CloudTrail](#)

AWS AppConfig is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in AWS AppConfig. CloudTrail captures all API calls for AWS AppConfig as events. This new topic provides AWS AppConfig-specific content rather than linking to the corresponding content in the *AWS Systems Manager User Guide*. For more information, see [Logging AWS AppConfig API calls using AWS CloudTrail](#).

January 18, 2024

[AWS AppConfig now supports AWS PrivateLink](#)

You can use AWS PrivateLink to create a private connection between your VPC and AWS AppConfig. You can access AWS AppConfig as if it were in your VPC, without the use of an internet gateway, NAT device, VPN connection, or Direct Connect connection. Instances in your VPC don't need public IP addresses to access AWS AppConfig. For more information, see [Access AWS AppConfig using an interface endpoint \(AWS PrivateLink\)](#).

December 6, 2023

[Additional AWS AppConfig Agent retrieval features and a new local development mode](#)

AWS AppConfig Agent offers the following additional features to help you retrieve configurations for your applications.

December 1, 2023

[Additional retrieval features](#)

- *Multi-account retrieval:* Use AWS AppConfig Agent from a primary or *retrieval* AWS account to retrieve configuration data from multiple vendor accounts.
- *Write configuration copy to disk:* Use AWS AppConfig Agent to write configuration data to disk. This feature enables customers with applications that read configuration data from disk to integrate with AWS AppConfig.

Note

Write configuration to disk is not designed as a configuration backup feature. AWS AppConfig Agent doesn't read from the configuration files copied to disk. If you want to back up configura

tions to disk, see the `BACKUP_DIRECTORY` and `PRELOAD_BACKUP` environment variables for [Using AWS AppConfig Agent with Amazon EC2](#) or [Using AWS AppConfig Agent with Amazon ECS and Amazon EKS](#).

Local development mode

AWS AppConfig Agent supports a *local development mode*. If you enable local development mode, the agent reads configuration data from a specified directory on disk. It doesn't retrieve configuration data from AWS AppConfig. You can simulate configuration deployments by updating files in the specified directory. We recommend local development mode for the following use cases:

- Test different configuration versions before deploying them using AWS AppConfig.
- Test different configuration options for a new feature before committing changes to your code repository.

- Test different configuration scenarios to verify they work as expected.

[New code samples topic](#)

Added a new [code samples](#) topic to this guide. The topic includes examples in Java, Python, and JavaScript for programmatically performing six common AWS AppConfig actions.

November 17, 2023

[Revised table of contents to better reflect AWS AppConfig workflow](#)

Content in this user guide is now grouped under the headings Creating, Deploying, Retrieving, and Extending workflows. This organization better reflects the workflow for using AWS AppConfig and aims to help make content more discoverable.

November 7, 2023

[Payload reference added](#)

The [Creating a Lambda function for a custom AWS AppConfig extension](#) topic now includes a request and response payload reference.

November 7, 2023

[New AWS predefined deployment strategy](#)

AWS AppConfig now offers and recommends the `AppConfig.Linear20PercentEvery6Minutes` predefined deployment strategy. For more information, see [Predefined deployment strategies](#).

August 11, 2023

[AWS AppConfig integration with Amazon EC2](#)

You can integrate AWS AppConfig with applications running on your Amazon Elastic Compute Cloud (Amazon EC2) Linux instances by using AWS AppConfig Agent. The agent supports x86_64 and ARM64 architectures for Amazon EC2. For more information, see [AWS AppConfig integration with Amazon EC2](#).

July 20, 2023

[CloudFormation support for new AWS AppConfig resources and a feature flag example](#)

AWS CloudFormation now supports the [AWS::AppConfig::Extension](#) and [AWS::AppConfig::ExtensionAssociation](#) resources to help you get started with AWS AppConfig extensions.

April 12, 2023

The [AWS::AppConfig::ConfigurationProfile](#) and [AWS::AppConfig::HostedConfigurationVersion](#) resources now include an example for creating a feature flag configuration profile in the AWS AppConfig hosted configuration store.

[AWS AppConfig integration with AWS Secrets Manager](#)

February 2, 2023

AWS AppConfig integrates with AWS Secrets Manager. Secrets Manager helps you to securely encrypt, store, and retrieve credentials for your databases and other services. Instead of hardcoding credentials in your apps, you can make calls to Secrets Manager to retrieve your credentials whenever needed. Secrets Manager helps you protect access to your IT resources and data by enabling you to rotate and manage access to your secrets.

When you create a freeform configuration profile, you can choose Secrets Manager as the source of your configuration data. You must onboard with Secrets Manager and create a secret before you create the configuration profile. For more information about Secrets Manager, see [What is AWS Secrets Manager?](#) in the *AWS Secrets Manager User Guide*. For information about creating a configuration profile, see [Creating a freeform configuration profile](#).

[AWS AppConfig integration with Amazon ECS and Amazon EKS](#)

December 2, 2022

You can integrate AWS AppConfig with Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS) by using the AWS AppConfig agent. The agent functions as a sidecar container running alongside your Amazon ECS and Amazon EKS container applications. The agent enhances containerized application processing and management in the following ways:

- The agent calls AWS AppConfig on your behalf by using an AWS Identity and Access Management (IAM) role and managing a local cache of configuration data. By pulling configuration data from the local cache, your application requires fewer code updates to manage configuration data, retrieves configuration data in milliseconds, and isn't affected by network issues that can disrupt calls for such data.
- The agent offers a native experience for retrieving

g and resolving AWS AppConfig feature flags.

- Out of the box, the agent provides best practices for caching strategies, polling intervals, and local configuration data availability while tracking the configuration tokens needed for subsequent service calls.
- While running in the background, the agent periodically polls the AWS AppConfig data plane for configuration data updates. Your containerized application can retrieve the data by connecting to localhost on port 2772 (a customizable default port value) and calling HTTP GET to retrieve the data.
- The AWS AppConfig agent updates configuration data in your containers without having to restart or recycle those containers.

For more information, see [AWS AppConfig integration with Amazon ECS and Amazon EKS](#).

[New extension: AWS
AppConfig extension for
CloudWatch Evidently](#)

September 13, 2022

You can use Amazon CloudWatch Evidently to safely validate new features by serving them to a specified percentage of your users while you roll out the feature. You can monitor the performance of the new feature to help you decide when to ramp up traffic to your users. This helps you reduce risk and identify unintended consequences before you fully launch the feature. You can also conduct A/B experiments to make feature design decisions based on evidence and data.

The AWS AppConfig extension for CloudWatch Evidently allows your application to assign variations to user sessions locally instead of by calling the [EvaluateFeature](#) operation. A local session mitigates the latency and availability risks that come with an API call. For information about how to configure and use the extension, see [Perform launches and A/B experiments with CloudWatch Evidently](#) in the *Amazon CloudWatch User Guide*.

[Deprecation of the GetConfiguration API action](#)

On Nov 18, 2021, AWS AppConfig released a new data plane service. This service replaces the previous process of retrieving configuration data by using the `GetConfiguration` API action. The data plane service uses two new API actions, [StartConfigurationSession](#) and [GetLatestConfiguration](#). The data plane service also uses [new endpoints](#).

September 13, 2022

For more information, see [About the AWS AppConfig data plane service](#).

[New version of the AWS AppConfig Agent Lambda extension](#)

Version 2.0.122 of the AWS AppConfig Agent Lambda extension is now available . The new extension uses different Amazon Resource Names (ARNs). For more information, see [AWS AppConfig Agent Lambda extension release notes](#).

August 23, 2022

[Launch of AWS AppConfig extensions](#)

An extension augments your ability to inject logic or behavior at different points during the AWS AppConfig workflow of creating or deploying a configuration. You can use AWS-authored extensions or create your own. For more information, see [Working with AWS AppConfig extensions](#).

July 12, 2022

[New version of the AWS AppConfig Agent Lambda extension](#)

Version 2.0.58 of the AWS AppConfig Agent Lambda extension is now available. The new extension uses different Amazon Resource Names (ARNs). For more information, see [Available versions of the AWS AppConfig Lambda extension](#).

May 3, 2022

[AWS AppConfig integration with Atlassian Jira](#)

Integrating with Atlassian Jira allows AWS AppConfig to create and update issues in the Atlassian console whenever you make changes to a [feature flag](#) in your AWS account for the specified AWS Region. Each Jira issue includes the flag name, application ID, configuration profile ID, and flag values. After you update, save, and deploy your flag changes, Jira updates the existing issues with the details of the change. For more information, see [AWS AppConfig integration with Atlassian Jira](#).

April 7, 2022

[General availability of feature flags and Lambda extension support for ARM64 \(Graviton 2\) processors](#)

With AWS AppConfig feature flags, you can develop a new feature and deploy it to production while hiding the feature from users. You start by adding the flag to AWS AppConfig as configuration data. Once the feature is ready to be released, you can update the flag configuration data without deploying any code. This feature improves the safety of your dev-ops environment because you don't need to deploy new code to release the feature. For more information, see [Creating a feature flag configuration profile](#).

March 15, 2022

General availability of feature flags in AWS AppConfig includes the following enhancements:

- The console includes an option to designate a flag as a *short term* flag. You can filter and sort the list of flags on short-term flags.
- For customers using feature flags in AWS Lambda, the new Lambda extension allows you to call individual feature flags by using an HTTP endpoint. For more information, see see

[Retrieving one or more flags from a feature flag configuration.](#)

This update also provides support for AWS Lambda extensions developed for ARM64 (Graviton2) processors. For more information, see [Available versions of the AWS AppConfig Lambda extension](#).

[The GetConfiguration API action is deprecated](#)

The GetConfiguration API action is deprecated. Calls to receive configuration data should use the StartConfigurationSession and GetLatestConfiguration APIs instead. For more information about these APIs and how to use them, see [Retrieving the configuration](#).

January 28, 2022

[New region ARN for AWS AppConfig Lambda extension](#)

AWS AppConfig Lambda extension is available in the new Asia Pacific (Osaka) region. The Amazon Resource Name (ARN) is required to create a Lambda in the region. For more information about the Asia Pacific (Osaka) region ARN, see [Adding the AWS AppConfig Lambda extension](#).

March 4, 2021

[AWS AppConfig Lambda extension](#)

If you use AWS AppConfig to manage configurations for a Lambda function, then we recommend that you add the AWS AppConfig Lambda extension. This extension includes best practices that simplify using AWS AppConfig while reducing costs. Reduced costs result from fewer API calls to the AWS AppConfig service and, separately, reduced costs from shorter Lambda function processing times. For more information, see [AWS AppConfig integration with Lambda extensions](#).

October 8, 2020

[New section](#)

Added a new section that provides instructions for setting up AWS AppConfig. For more information, see [Setting up AWS AppConfig](#).

September 30, 2020

[Added commandline procedures](#)

Procedures in this user guide now include commandline steps for the AWS Command Line Interface (AWS CLI) and Tools for Windows PowerShell. For more information, see [Working with AWS AppConfig](#).

September 30, 2020

[Launch of AWS AppConfig user guide](#)

Use AWS AppConfig, a tool in AWS Systems Manager, to create, manage, and quickly deploy application configurations. AWS AppConfig supports controlled deployments to applications of any size and includes built-in validation checks and monitoring. You can use AWS AppConfig with applications hosted on EC2 instances, AWS Lambda, containers, mobile applications, or IoT devices.

July 31, 2020