

# Amazon SDKs and Tools



# Amazon SDKs and Tools: Reference Guide

# Table of Contents

<b>Amazon SDKs and Tools Reference Guide .....</b>	<b>1</b>
Developer resources .....	2
<b>Configuration .....</b>	<b>3</b>
Shared config and credentials files .....	4
Profiles .....	4
Format of the config file .....	5
Format of the credentials file .....	8
Location of the shared files .....	9
Change the default location of these files .....	10
Environment variables .....	11
How to set environment variables .....	11
Serverless environment variable setup .....	12
JVM system properties .....	13
How to set JVM system properties .....	13
<b>Authentication and access .....</b>	<b>15</b>
Amazon Builder ID .....	16
IAM Identity Center authentication .....	17
Configure programmatic access using IAM Identity Center .....	17
Understand IAM Identity Center authentication .....	21
IAM Roles Anywhere .....	24
Step 1: Configure IAM Roles Anywhere .....	24
Step 2: Use IAM Roles Anywhere .....	25
Assume a role .....	26
Assume an IAM role .....	27
Federate with web identity or OpenID Connect .....	28
Amazon access keys .....	29
Use short-term credentials .....	29
Use long-term credentials .....	30
Short-term credentials .....	31
Long-term credentials .....	32
IAM roles for Amazon EC2 instances .....	35
Create an IAM role .....	36
Launch an Amazon EC2 instance and specify your IAM role .....	36
Connect to the EC2 instance .....	36

Run the sample application on the EC2 instance .....	37
<b>Settings reference .....</b>	<b>38</b>
Creating service clients .....	38
Precedence of settings .....	38
Config file settings list .....	39
Credentials file settings list .....	43
Environment variables list .....	43
JVM system properties list .....	47
Standardized credential providers .....	50
Credential provider chain .....	50
Amazon access keys .....	52
Assume role provider .....	55
Container provider .....	60
IAM Identity Center provider .....	64
IMDS provider .....	70
Process provider .....	74
Standardized features .....	78
Amazon EC2 instance metadata .....	79
Amazon S3 access points .....	81
Amazon S3 Multi-Region Access Points .....	83
Amazon Web Services Region .....	85
Amazon STS Regionalized endpoints .....	88
Dual-stack and FIPS endpoints .....	90
Endpoint discovery .....	92
General configuration .....	94
IMDS client .....	98
Retry behavior .....	100
Request compression .....	105
Service-specific endpoints .....	107
Smart configuration defaults .....	151
<b>Common Runtime .....</b>	<b>157</b>
CRT dependencies .....	158
<b>Maintenance and support .....</b>	<b>159</b>
Maintenance policy .....	159
Overview .....	159
Versioning .....	159

---

SDK major version lifecycle .....	160
Dependency lifecycle .....	160
Communication methods .....	161
Version support matrix .....	162
IDE Toolkits .....	164
Telemetry Notification .....	165
<b>Document history .....</b>	<b>166</b>
<b>Amazon Glossary .....</b>	<b>168</b>

# Amazon SDKs and Tools Reference Guide

Many SDKs and tools share some common functionality, either through shared design specifications or through a shared library.

This guide includes information regarding:

- [Configuration](#) – How to use the shared `config` and `credentials` files or environment variables to configure your Amazon SDKs and tools.
- [Authentication and access](#) – Establish how your code or tool authenticates with Amazon when you develop with Amazon Web Services.
- [Settings reference](#) – Reference for all standardized settings available for authentication and configuration.
- [Amazon Common Runtime \(CRT\) libraries](#) – Overview of the shared Amazon Common Runtime (CRT) libraries that are available to almost all SDKs.
- [Amazon SDKs and Tools maintenance policy](#) covers the maintenance policy and versioning for Amazon Software Development Kits (SDKs) and tools, including Mobile and Internet of Things (IoT) SDKs, and their underlying dependencies.

This Amazon SDKs and Tools Reference Guide is intended to be a base of information that is applicable to multiple SDKs and tools. The specific guide for the SDK or tool that you are using should be used in addition to any information presented here. The following are the SDK and tools which have relevant sections of material in this guide:

If you are using:	This guide's relevant sections for you are:
• Any SDK or tool	<a href="#">Amazon SDKs and Tools maintenance policy</a>
• <a href="#">Amazon Cloud Development Kit (Amazon CDK) Developer Guide</a> • <a href="#">Amazon Serverless Application Model Developer Guide</a> • <a href="#">Amazon Toolkit for Eclipse User Guide</a> • <a href="#">Amazon Toolkit for JetBrains User Guide</a> • <a href="#">Amazon Toolkit for Visual Studio User Guide</a>	<a href="#">Configuration</a> <a href="#">Authentication and access</a> <a href="#">Amazon SDKs and Tools maintenance policy</a>

If you are using:	This guide's relevant sections for you are:
<ul style="list-style-type: none"><li>• <a href="#">Amazon Toolkit for Visual Studio Code User Guide</a></li></ul>	
<ul style="list-style-type: none"><li>• <a href="#">Amazon Command Line Interface User Guide</a></li><li>• <a href="#">Amazon SDK for C++ Developer Guide</a></li><li>• <a href="#">Amazon SDK for Go Developer Guide</a></li><li>• <a href="#">Amazon SDK for Java Developer Guide</a></li><li>• <a href="#">Amazon SDK for JavaScript Developer Guide</a></li><li>• <a href="#">Amazon SDK for Kotlin</a></li><li>• <a href="#">Amazon SDK for .NET Developer Guide</a></li><li>• <a href="#">Amazon SDK for PHP Developer Guide</a></li><li>• <a href="#">Amazon SDK for Python (Boto3) Getting Started</a></li><li>• <a href="#">Amazon SDK for Ruby Developer Guide</a></li><li>• <a href="#">Amazon SDK for Rust</a></li><li>• <a href="#">Amazon SDK for Swift</a></li><li>• <a href="#">Amazon Tools for Windows PowerShell User Guide</a></li></ul>	<ul style="list-style-type: none"><li>• <a href="#">Configuration</a></li><li>• <a href="#">Authentication and access</a></li><li>• <a href="#">Settings reference</a></li><li>• <a href="#">Amazon Common Runtime (CRT) libraries</a></li><li>• <a href="#">Amazon SDKs and Tools maintenance policy</a></li></ul>

## Developer resources

Amazon CodeWhisperer is a machine learning (ML)–powered service that helps improve developer productivity by generating code recommendations based on both code comments and code in the integrated development environment (IDE). To learn more about which languages and IDEs are supported, as well as how to sign up for free preview, see [Amazon CodeWhisperer](#).

# Configuration

With Amazon SDKs and other Amazon developer tools, such as the Amazon Command Line Interface (Amazon CLI), you can interact with Amazon service APIs. Before attempting that, however, you must configure the SDK or tool with the information that it needs to perform the requested operation.

This information includes the following items:

- **Credentials information** that identifies who is calling the API. The credentials are used to encrypt the request to the Amazon servers. Using this information, Amazon confirms your identity and can retrieve permissions policies associated with it. Then it can determine what actions you're allowed to perform.
- **Other configuration details** that you use to tell the Amazon CLI or SDK how to process the request, where to send the request (to which Amazon service endpoint), and how to interpret or display the response.

Each SDK or tool supports multiple sources that you can use to supply the required credential and configuration information. Some sources are unique to the SDK or tool, and you must refer to the documentation for that tool or SDK for the details on how to use that method.

However, most of the Amazon SDKs and tools support common settings from two primary sources (beyond the code itself):

- [Shared Amazon config and credentials files](#) – The shared config and credentials files are the most common way that you can specify authentication and configuration to an Amazon SDK or tool. Use these files to store settings that your tools and applications can use. Settings within the shared config and credentials files are associated with a specific profile. With multiple profiles, you can create different settings configurations to apply in different scenarios. When you use an Amazon tool to invoke a command or use an SDK to invoke an Amazon API, you can specify which profile, and thus which configuration settings, to use for that action. One of the profiles is designated as the default profile and is used automatically when you don't explicitly specify a profile to use. The settings that you can store in these files are documented in this reference guide.
- [Environment variables](#) – Some of the settings can alternatively be stored in the environment variables of your operating system. Although you can have only one set of environment



variables in effect at a time, they are easily modified dynamically as your program runs and your requirements change.

### Additional topics in this section

- [Shared config and credentials files](#)
- [Location of the shared config and credentials files](#)
- [Environment variables support](#)
- [JVM system properties support](#)

## Shared config and credentials files

The shared Amazon `config` and `credentials` files contain a set of profiles. A profile is a set of configuration values that can be referenced from the SDK/tool using its profile name. Configuration values are attached to a profile in order to configure some aspect of the SDK/tool when that profile is used. These files are "shared" in that the values take affect for any applications, processes, or SDKs on the local environment for a user.

As a general rule, any value that you can place in the shared `credentials` file can alternatively be placed in the shared `config` file. The reverse isn't true; only a few settings can be placed in the `credentials` file. However, as a security best practice, we recommend that you keep any sensitive values, such as access key IDs and secret keys, in the separate `credentials` file. This way, you can provide separate permissions for each file, if necessary.

Both the shared `config` and `credentials` files are plaintext files that contain only ASCII characters (UTF-8 encoded). They take the form of what are generally referred to as [INI files](#).

## Profiles

Settings within the shared `config` and `credentials` files are associated with a specific profile. With multiple profiles, you can create different settings configurations to apply in different scenarios.

The `[default]` profile contains the values that are used by an SDK or tool operation if a specific named profile is not specified. You can also create separate profiles that you can explicitly reference by name. Each named profile can have a different group of settings.

[default] is simply an unnamed profile. This profile is named default because it is the default profile used by the SDK if the user does not specify a profile. It does not provide inherited default values to other profiles. For example, if you set something in the [default] profile and you don't set it in a named profile, then the value isn't set when you use the named profile.

## Set a named profile

Instead of using the [default] profile, you can set a named profile.

Configure this functionality by setting one of the following:

### AWS\_PROFILE - environment variable

All SDK code and Amazon CLI commands use the settings in this named profile.

Linux/macOS example of setting environment variables via command line:

```
export AWS_PROFILE="my_default_profile_name";
```

Windows example of setting environment variables via command line:

```
setx AWS_PROFILE "my_default_profile_name"
```

### aws.profile - JVM system property

For SDKs that run on a JVM, you can [set the aws.profile system property](#). When the SDK creates a service client, it uses the settings in the named profile unless the setting is overridden in code.

## Format of the config file

The config file is organized into sections. A section is a named collection of settings, and continues until another section definition line is encountered.

The config file is a plaintext file that uses the following format:

- All entries in a section take the general form of `setting-name=value`.
- Lines can be commented out by starting the line with a hashtag character (#).

## Section types

A section definition is a line that applies a name to a collection of settings. Section definition lines start and end with square brackets ([ ]). Inside the brackets, there is a section type identifier and a custom name for the section. You can use letters, numbers, hyphens ( - ), and underscores ( \_ ), but no spaces.

### Section type: profile

Example section definition line: [profile *dev*]

The `profile` section definition line names a configuration grouping that you can apply in different scenarios. [default] is the only profile that does not require the `profile` section identifier. To better understand named profiles, see the preceding section on Profiles.

The following example shows a basic config file with a [default] profile. It sets the [region](#) setting.

```
[default]
#Full line comment, this text is ignored.
region = us-east-2
```

The following example shows a config file with a `profile` section definition line. It uses the identifier `profile` followed by a unique name for the profile. All settings that follow this line, up until another section definition is encountered, will be included with this named profile.

```
[profile developers]
...settings...
```

Some settings have their own nested group of subsettings, such as the `s3` setting and subsettings in the following example. Associate the subsettings with the group by indenting them by one or more spaces.

```
[profile testers]
region = us-west-2
s3 =
    max_concurrent_requests=10
    max_queue_size=1000
```

## Section type: sso-session

Example section definition line: `[sso-session my-sso]`

The `sso-session` section definition line names a group of settings that you use to configure a profile to resolve Amazon credentials using Amazon IAM Identity Center. For more information on configuring single sign-on authentication, see [IAM Identity Center authentication](#). A profile is linked to a `sso-session` section by a key-value pair where `sso-session` is the key and the name of your `sso-session` section is the value, such as `sso-session = <name-of-sso-session-section>`.

The following example configures a profile that will get short-term Amazon credentials for the "SampleRole" IAM role in the "111122223333" account using a token from the "my-sso". The "my-sso" `sso-session` section is referenced in the profile section by name using the `sso-session` key.

```
[profile dev]  
sso_session = my-sso  
sso_account_id = 111122223333  
sso_role_name = SampleRole  
  
[sso-session my-sso]  
sso_region = us-east-1  
sso_start_url = https://my-sso-portal.awsapps.com/start
```

## Section type: services

Example section definition line: `[services dev]`

### Note

The `services` section supports service-specific endpoint customizations and is only available in SDKs and tools that include this feature. To see if this feature is available for your SDK, see [Compatibility with Amazon SDKs](#) for service-specific endpoints.

The `services` section definition line names a group of settings that configures custom endpoints for Amazon Web Service requests. A profile is linked to a `services` section by a key-value pair where `services` is the key and the name of your `services` section is the value, such as `services = <name-of-services-section>`.

The `services` section is further separated into subsections by `<SERVICE> =` lines, where `<SERVICE>` is the Amazon Web Service identifier key. The Amazon Web Service identifier is based on the API model's `serviceId` by replacing all spaces with underscores and lowercasing all letters. For a list of all service identifier keys to use in the `services` section, see [Identifiers for service-specific endpoints](#). The service identifier key is followed by nested settings with each on its own line and indented by two spaces.

The following example uses a `services` definition to configure the endpoint to use for requests made only to the Amazon DynamoDB service. The `"local-dynamodb"` `services` section is referenced in the `profile` section by name using the `services` key. The Amazon Web Service identifier key is `dynamodb`. The Amazon DynamoDB service subsection begins on the line `dynamodb =`. Any immediately following lines that are indented are included in that subsection and apply to that service.

```
[profile dev]
services = local-dynamodb

[services local-dynamodb]
dynamodb =
  endpoint_url = http://localhost:8000
```

For more information on custom endpoint configuration, see [Service-specific endpoints](#).

## Format of the credentials file

The rules for the `credentials` file are generally identical to those for the `config` file, except that profile sections don't begin with the word `profile`. Use only the unique profile name itself between square brackets.

```
[dev]
...settings...
```

You can store only a small subset of settings and values in the `credentials` file. Generally, it's only those with values that would be considered "secrets" or sensitive, such as access key IDs and secret keys. The page for each setting in this guide states whether it can be stored in the `credentials` file or only in the `config` file.

The following example shows a basic `credentials` file with a `[default]` profile. It sets the [aws\\_access\\_key\\_id](#) and [aws\\_secret\\_access\\_key](#) global settings.



## Change the default location of these files

### Use environment variables

The following environment variables can be set to change the location or name of these files from the default to a custom value:

- config file environment variable: **AWS\_CONFIG\_FILE**
- credentials file environment variable: **AWS\_SHARED\_CREDENTIALS\_FILE**

#### Linux/macOS

You can specify an alternate location by running the following [export](#) commands on Linux or macOS.

```
$ export AWS_CONFIG_FILE=/some/file/path/on/the/system/config-file-name
$ export AWS_SHARED_CREDENTIALS_FILE=/some/other/file/path/on/the/system/
credentials-file-name
```

#### Windows

You can specify an alternate location by running the following [setx](#) commands on Windows.

```
C:\> setx AWS_CONFIG_FILE c:\some\file\path\on\the\system\config-file-name
C:\> setx AWS_SHARED_CREDENTIALS_FILE c:\some\other\file\path\on\the\system
\credentials-file-name
```

### Use JVM system properties

You can set the following JVM system properties to change the location or name of these files from the default to a custom value:

- config file JVM system property: **aws.configFile**
- credentials file environment variable: **aws.sharedCredentialsFile**

For instructions on how to set JVM system properties see [the section called “How to set JVM system properties”](#).

# Environment variables support

Environment variables provide another way to specify configuration options and credentials, and can be useful for scripting or temporarily setting a named profile as the default. For the list of environment variables supported by most SDKs, see [Environment variables list](#).

## Precedence of options

- If you specify a setting by using its environment variable, it overrides any value loaded from a profile in the shared Amazon config and credentials files.
- If you specify a setting by using a parameter on the Amazon CLI command line, it overrides any value from either the corresponding environment variable or a profile in the configuration file.

## How to set environment variables

The following examples show how you can configure environment variables for the default user.

Linux, OS X, or Unix

```
$ export AWS_ACCESS_KEY_ID=AKIAIOSFODNN7EXAMPLE
$ export AWS_SECRET_ACCESS_KEY=wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
$ export
  AWS_SESSION_TOKEN=AQoEXAMPLEH4aoAH0gNCAPy...truncated...zrkuWJ0gQs8IZZaIv2BXIa2R40lgk
$ export AWS_REGION=us-west-2
```

Setting the environment variable changes the value used until the end of your shell session, or until you set the variable to a different value. You can make the variables persistent across future sessions by setting them in your shell's startup script.

Windows Command Prompt

```
C:\> setx AWS_ACCESS_KEY_ID AKIAIOSFODNN7EXAMPLE
C:\> setx AWS_SECRET_ACCESS_KEY wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
C:\> setx
  AWS_SESSION_TOKEN AQoEXAMPLEH4aoAH0gNCAPy...truncated...zrkuWJ0gQs8IZZaIv2BXIa2R40lgk
C:\> setx AWS_REGION us-west-2
```

Using [set](#) to set an environment variable changes the value used until the end of the current Command Prompt session, or until you set the variable to a different value. Using [setx](#) to set



an environment variable changes the value used in both the current Command Prompt session and all Command Prompt sessions that you create after running the command. It does *not* affect other command shells that are already running at the time you run the command.

## PowerShell

```
PS C:\> $Env:AWS_ACCESS_KEY_ID="AKIAIOSFODNN7EXAMPLE"
PS C:\> $Env:AWS_SECRET_ACCESS_KEY="wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY"
PS C:\>
\> $Env:AWS_SESSION_TOKEN="AQoEXAMPLEH4aoAH0gNCAPy...truncated...zrkuWJ0gQs8IZZaIv2BXIa2R40L"
PS C:\> $Env:AWS_REGION="us-west-2"
```

If you set an environment variable at the PowerShell prompt as shown in the previous examples, it saves the value for only the duration of the current session. To make the environment variable setting persistent across all PowerShell and Command Prompt sessions, store it by using the **System** application in **Control Panel**. Alternatively, you can set the variable for all future PowerShell sessions by adding it to your PowerShell profile. See the [PowerShell documentation](#) for more information about storing environment variables or persisting them across sessions.

## Serverless environment variable setup

If you use a serverless architecture for development, you have other options for setting environment variables. Depending on your container, you can use different strategies for code running in those containers to see and access environment variables, similar to non-cloud environments.

For example, with Amazon Lambda, you can directly set environment variables. For details, see [Using Amazon Lambda environment variables](#) in the *Amazon Lambda Developer Guide*.

In Serverless Framework, you can often set SDK environment variables in the `serverless.yml` file under the provider key under the environment setting. For information on the `serverless.yml` file, see [General function settings](#) in the Serverless Framework documentation.

Regardless of which mechanism you use to set container environment variables, there are some that are reserved by the container, such as those documented for Lambda at [Defined runtime environment variables](#). Always consult the official documentation for the container that you're using to determine how environment variables are treated and whether there are any restrictions.

# JVM system properties support

[JVM system properties](#) provide another way to specify configuration options and credentials for SDKs that run on the JVM such as the Amazon SDK for Java and the Amazon SDK for Kotlin. For a list of JVM system properties supported by SDKs, see [Settings reference](#).

## Precedence of options

- If you specify a setting by using its JVM system property, it overrides any value found in environment variables or loaded from a profile in the shared AWS config and credentials files.
- If you specify a setting by using its environment variable, it overrides any value loaded from a profile in the shared AWS config and credentials files.

## How to set JVM system properties

You can set JVM system properties several ways.

### On the command line

Set JVM system properties on the command-line when invoking the `java` command by using the `-D` switch. The following command configures the Amazon Web Services Region globally for all service clients unless you explicitly override the value in code.

```
java -Daws.region=us-east-1 -jar <your_application.jar> <other_arguments>
```

If you need to set multiple JVM system properties, specify the `-D` switch multiple times.

### With an environment variable

If you can't access the command line to invoke the JVM to run your application, you can use the `JAVA_TOOL_OPTIONS` environment variable to configure command-line options. This approach is useful in situations such as running an Amazon Lambda function on the Java runtime or running code in an embedded JVM.

The following example configures the Amazon Web Services Region globally for all service clients unless you explicitly override the value in code.

## Linux, OS X, or Unix

```
$ export JAVA_TOOL_OPTIONS="-Daws.region=us-east-1"
```

Setting the environment variable changes the value used until the end of your shell session, or until you set the variable to a different value. You can make the variables persistent across future sessions by setting them in your shell's startup script.

## Windows Command Prompt

```
C:\> setx JAVA_TOOL_OPTIONS -Daws.region=us-east-1
```

Using [set](#) to set an environment variable changes the value used until the end of the current Command Prompt session, or until you set the variable to a different value. Using [setx](#) to set an environment variable changes the value used in both the current Command Prompt session and all Command Prompt sessions that you create after running the command. It does **not** affect other command shells that are already running at the time you run the command.

## At runtime

You can also set JVM system properties at runtime in code by using the `System.setProperty` method as shown in the following example.

```
System.setProperty("aws.region", "us-east-1");
```

### Important

Set any JVM system properties *before* you initialize SDK service clients, otherwise service clients may use other values.

# Authentication and access

You must establish how your code authenticates with Amazon when you develop with Amazon Web Services. You can configure programmatic access to Amazon resources in different ways, depending on the environment and the Amazon access available to you.

## Authentication options for code running locally (not in Amazon)

- [IAM Identity Center authentication](#) – As a security best practice, we recommend using Amazon Organizations with IAM Identity Center to manage access across all your Amazon Web Services accounts. You can create users in Amazon IAM Identity Center, use Microsoft Active Directory, use a SAML 2.0 identity provider (IdP), or individually federate your IdP to Amazon Web Services accounts. To check if your Region supports IAM Identity Center, see [Amazon IAM Identity Center endpoints and quotas](#) in the *Amazon Web Services General Reference*.
- [IAM Roles Anywhere](#) – You can use IAM Roles Anywhere to obtain temporary security credentials in IAM for workloads such as servers, containers, and applications that run outside of Amazon. To use IAM Roles Anywhere, your workloads must use X.509 certificates.
- [Assume a role](#) – You can assume an IAM role to temporarily access Amazon resources that you might not have access to otherwise.
- [Amazon access keys](#) – Other options that might be less convenient or might increase the security risk to your Amazon resources.

## Authentication options for code running within an Amazon environment

- [Using IAM roles for Amazon EC2 instances](#) – Use IAM roles to securely run your application on an Amazon EC2 instance.
- You can programmatically interact with Amazon using IAM Identity Center in the following ways:
  - Use [Amazon CloudShell](#) to run Amazon CLI commands from the console.
  - Use [Amazon Cloud9](#) to start programming on Amazon using an integrated development environment (IDE) with Amazon resources.
  - To try cloud-based collaboration space for software development teams, consider using [Amazon CodeCatalyst](#).

## Authentication through a web-based identity provider - Mobile or client-based web applications

If you are creating mobile applications or client-based web applications that require access to Amazon, build your app so that it requests temporary Amazon security credentials dynamically by using web identity federation.

With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, app users can sign in using a well-known external identity provider (IdP), such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP. They can receive an authentication token, and then exchange that token for temporary security credentials in Amazon that map to an IAM role with permissions to use the resources in your Amazon Web Services account.

To learn how to configure this for your SDK or tool, see [Federate with web identity or OpenID Connect](#).

For mobile applications, consider using Amazon Cognito. Amazon Cognito acts as an identity broker and does much of the federation work for you. For more information, see [Using Amazon Cognito for mobile apps](#) in the *IAM User Guide*.

### More information about access management

The *IAM User Guide* has the following information about securely controlling access to Amazon resources:

- [IAM Identities \(users, user groups, and roles\)](#) – Understand the basics of identities in Amazon.
- [Security best practices in IAM](#) – Security recommendations to follow when developing Amazon applications according to the [shared-responsibility model](#).

The *Amazon Web Services General Reference* has foundational basics on the following:

- [Understanding and getting your Amazon credentials](#) – Access key options and management practices for both console and programmatic access.

## Amazon Builder ID

Your Amazon Builder ID complements any Amazon Web Services accounts you might already own or want to create. While an Amazon Web Services account acts as a container for Amazon resources you create and provides a security boundary for those resources, your Amazon Builder ID represents you as an individual. You can sign in with your Amazon Builder ID to access developer tools and services such as Amazon CodeWhisperer and Amazon CodeCatalyst.

- [Sign in with Amazon Builder ID](#) in the *Amazon Sign-In User Guide* – Learn how to create and use an Amazon Builder ID and learn what the Builder ID provides.
- [Authenticating with CodeWhisperer and Amazon Toolkit - Builder ID](#) in the *CodeWhisperer User Guide* – Learn how CodeWhisperer uses an Amazon Builder ID.
- [CodeCatalyst concepts - Amazon Builder ID](#) in the *Amazon CodeCatalyst User Guide* – Learn how CodeCatalyst uses an Amazon Builder ID.

## IAM Identity Center authentication

Amazon IAM Identity Center is the recommended method of providing Amazon credentials when developing on a non-Amazon compute service. For example, this would be something like your local development environment. If you are developing on an Amazon resource, such as Amazon Elastic Compute Cloud (Amazon EC2) or Amazon Cloud9, we recommend getting credentials from that service instead.

In this tutorial, you establish IAM Identity Center access and will configure it for your SDK or tool by using the Amazon access portal and the Amazon CLI.

- The Amazon access portal is the web location where you manually sign in to the IAM Identity Center. The format of the URL is `d-xxxxxxxxxx.awsapps.com/start` or `your_subdomain.awsapps.com/start`. When signed in to the Amazon access portal, you can view Amazon Web Services accounts and roles that have been configured for that user. This procedure uses the Amazon access portal to get configuration values you need for the SDK/tool authentication process.
- The Amazon CLI is used to configure your SDK or tool to use IAM Identity Center authentication for API calls made by your code. This one-time process updates your shared Amazon config file, that is then used by your SDK or tool when you run your code.

## Configure programmatic access using IAM Identity Center

### Step 1: Establish access and select appropriate permission set

If you haven't enabled IAM Identity Center yet, see [Enabling IAM Identity Center](#) in the *Amazon IAM Identity Center User Guide*.

Choose one of the following methods to access your Amazon credentials.

## I do not have established access through IAM Identity Center

1. Add a user and add administrative permissions by following the [Configure user access with the default IAM Identity Center directory](#) procedure in the *Amazon IAM Identity Center User Guide*.
2. The `AdministratorAccess` permission set should not be used for regular development. Instead, we recommend using the predefined `PowerUserAccess` permission set, unless your employer has created a custom permission set for this purpose.

Follow the same [Configure user access with the default IAM Identity Center directory](#) procedure again, but this time:

- Instead of creating the *Admin team* group, create a *Dev team* group, and substitute this thereafter in the instructions.
- You can use the existing user, but the user must be added to the new *Dev team* group.
- Instead of creating the *AdministratorAccess* permission set, create a *PowerUserAccess* permission set, and substitute this thereafter in the instructions.

When you are done, you should have the following:

- A *Dev team* group.
  - An attached `PowerUserAccess` permission set to the *Dev team* group.
  - Your user added to the *Dev team* group.
3. Exit the portal and sign in again to see your Amazon Web Services accounts and options for `Administrator` or `PowerUserAccess`. Select `PowerUserAccess` when working with your tool/SDK.

## I already have access to Amazon through a federated identity provider managed by my employer (such as Microsoft Entra or Okta)

Sign in to Amazon through your identity provider's portal. If your Cloud Administrator has granted you `PowerUserAccess` (developer) permissions, you see the Amazon Web Services accounts that you have access to and your permission set. Next to the name of your permission set, you see options to access the accounts manually or programmatically using that permission set.

Custom implementations might result in different experiences, such as different permission set names. If you're not sure which permission set to use, contact your IT team for help.

## I already have access to Amazon through the Amazon access portal managed by my employer

Sign in to Amazon through the Amazon access portal. If your Cloud Administrator has granted you `PowerUserAccess` (developer) permissions, you see the Amazon Web Services accounts that you have access to and your permission set. Next to the name of your permission set, you see options to access the accounts manually or programmatically using that permission set.

## I already have access to Amazon through a federated custom identity provider managed by my employer

Contact your IT team for help.

## Step 2: Configure SDKs and tools to use IAM Identity Center

1. On your development machine, install the latest Amazon CLI.
  - a. See [Installing or updating the latest version of the Amazon CLI](#) in the *Amazon Command Line Interface User Guide*.
  - b. (Optional) To verify that the Amazon CLI is working, open a command prompt and run the `aws --version` command.
2. Sign in to the Amazon access portal. Your employer may provide this URL or you may get it in an email following **Step 1: Establish access**. If not, you can find your Amazon access portal URL on the Dashboard of <https://console.amazonaws.cn/singlesignon/>.
  - a. In the Amazon access portal, select the appropriate permission set, then select its **Command line or programmatic access** link. Use the predefined `PowerUserAccess` permission set, or whichever permission set you or your employer has created to apply least-privilege permissions for development.
  - b. In the **Get credentials** dialog box, choose either **MacOS and Linux** or **Windows**, depending on your operating system.
  - c. Choose the **IAM Identity Center credentials** method to get the `SSO_Start_URL` and `SSO_Region` values that you need for the next step.
3. In the Amazon CLI command prompt, run the `aws configure sso` command. When prompted, enter the configuration values that you collected in the previous step. For details on this Amazon CLI command, see [Configure your profile with the `aws configure sso` wizard](#).



- For **CLI profile name**, we recommend entering *default* when you are getting started. For information about how to set non-default (named) profiles and their associated environment variable, see [Profiles](#).
4. (Optional) In the Amazon CLI command prompt, confirm the active session identity by running the `aws sts get-caller-identity` command. The response should show the IAM Identity Center permission set that you configured.
  5. If you are using an Amazon SDK, create an application for your SDK in your development environment.
    - a. For some SDKs, additional packages such as SSO and SS00IDC must be added to your application before you can use IAM Identity Center authentication. For details, see your specific SDK.
    - b. If you previously configured access to Amazon, review your shared Amazon `credentials` file for any [Amazon access keys](#). You must remove any static credentials before the SDK or tool will use the IAM Identity Center credentials because of the [Credential provider chain](#) precedence.

For a deep dive into how the SDKs and tools use and refresh credentials using this configuration, see [Understand IAM Identity Center authentication](#).

Depending on your configured session lengths, your access will eventually expire and the SDK or tool will encounter an authentication error. To refresh the access portal session again when needed, use the Amazon CLI to run the `aws sso login` command.

You can extend both the IAM Identity Center access portal session duration and the permission set session duration. This lengthens the amount of time that you can run code before you need to manually sign in again with the Amazon CLI. For more information, see the following topics in the *Amazon IAM Identity Center User Guide*:

- **IAM Identity Center session duration** – [Configure the duration of your users' Amazon access portal sessions](#)
- **Permission set session duration** – [Set session duration](#)

For details on all IAM Identity Center provider settings for SDKs and tools, see [IAM Identity Center credential provider](#) in this guide.

# Understand IAM Identity Center authentication

## Relevant IAM Identity Center terms

The following terms help you understand the process and configuration behind Amazon IAM Identity Center. The documentation for Amazon SDK APIs uses different names than IAM Identity Center for some of these authentication concepts. It's helpful to know both names.

The following table shows how alternative names relate to each other.

IAM Identity Center name	SDK API name	Description
Identity Center	sso	Although Amazon Single Sign-On is renamed, the sso API namespaces will keep their original name for backward compatibility purposes. For more information, see <a href="#">IAM Identity Center rename</a> in the <i>Amazon IAM Identity Center User Guide</i> .
IAM Identity Center console Administrative console		The console you use to configure single sign-on.
Amazon access portal URL		A URL unique to your IAM Identity Center account, like <code>https://xxx.awsapps.com/start</code> . You sign in to this portal using your IAM Identity Center sign-in credentials.
IAM Identity Center Access Portal session	Authentication session	Provides a bearer access token to the caller.
Permission set session		The IAM session that the SDK uses internally to make the

IAM Identity Center name	SDK API name	Description
		Amazon Web Service calls. In informal discussions, you might see this incorrectly referred to as "role session."
Permission set credentials	Amazon credentials sigv4 credentials	The credentials the SDK actually uses for most Amazon Web Service calls (specifically, all sigv4 Amazon Web Service calls). In informal discussions, you might see this incorrectly referred to as "role credentials."
IAM Identity Center credential provider	SSO credential provider	How you get the credentials, such as the class or module providing the functionality.

## Understand SDK credential resolution for Amazon Web Services

The IAM Identity Center API exchanges bearer token credentials for sigv4 credentials. Most Amazon Web Services are sigv4 APIs, with a few exceptions like Amazon CodeWhisperer and Amazon CodeCatalyst. The following describes the credential resolution process for supporting most Amazon Web Service calls for your application code through Amazon IAM Identity Center.

### Start an Amazon access portal session

- Start the process by signing in to the session with your credentials.
  - Use the `aws sso login` command in the Amazon Command Line Interface (Amazon CLI). This starts a new IAM Identity Center session if you don't already have an active session.
- When you start a new session, you receive a refresh token and access token from IAM Identity Center. The Amazon CLI also updates an SSO cache JSON file with a new access token and refresh token and makes it available for use by SDKs.
- If you already have an active session, the Amazon CLI command reuses the existing session and will expire whenever the existing session expires. To learn how to set the length of an IAM

Identity Center session, see [Configure the duration of your users' Amazon access portal sessions](#) in the *Amazon IAM Identity Center User Guide*.

- The maximum session length has been extended to 90 days to reduce the need for frequent sign-ins.

## How the SDK gets credentials for Amazon Web Service calls

SDKs provide access to Amazon Web Services when you instantiate a client object per service. When the selected profile of the shared Amazon config file is configured for IAM Identity Center credential resolution, IAM Identity Center is used to resolve credentials for your application.

- The [credential resolution process](#) is completed during runtime when a client is created.

To retrieve credentials for sigv4 APIs using IAM Identity Center single sign-on, the SDK uses the IAM Identity Center access token to get an IAM session. This IAM session is called a permission set session, and it provides Amazon access to the SDK by assuming an IAM role.

- The permission set session duration is set independently from the IAM Identity Center session duration.
  - To learn how to set the permission set session duration, see [Set session duration](#) in the *Amazon IAM Identity Center User Guide*.
- Be aware that the permission set credentials are also referred to as *Amazon credentials* and *sigv4 credentials* in most Amazon SDK API documentation.

The permission set credentials are returned from a call to [getRoleCredentials](#) of the IAM Identity Center API to the SDK. The SDK's client object uses that assumed IAM role to make calls to the Amazon Web Service, such as asking Amazon S3 to list the buckets in your account. The client object can continue to operate using those permission set credentials until the permission set session expires.

## Session expiration and refresh

When using the [SSO token provider configuration](#), the hourly access token obtained from IAM Identity Center is automatically refreshed using the refresh token.

- If the access token is expired when the SDK tries to use it, the SDK uses the refresh token to try to get a new access token. The IAM Identity Center compares the refresh token to your IAM

Identity Center access portal session duration. If the refresh token is not expired, the IAM Identity Center responds with another access token.

- This access token can be used to either refresh the permission set session of existing clients, or to resolve credentials for new clients.

However, if the IAM Identity Center access portal session is expired, then no new access token is granted. Therefore, the permission set duration cannot be renewed. It will expire (and access will be lost) whenever the cached permission set session length times out for existing clients.

Any code that creates a new client will fail authentication as soon as the IAM Identity Center session expires. This is because the permission set credentials are not cached. Your code won't be able to create a new client and complete the credential resolution process until you have a valid access token.

To recap, when the SDK needs new permission set credentials, the SDK first checks for any valid, existing credentials and uses those. This applies whether the credentials are for a new client or for an existing client with expired credentials. If credentials aren't found or they're not valid, then the SDK calls the IAM Identity Center API to get new credentials. To call the API, it needs the access token. If the access token is expired, the SDK uses the refresh token to try to get a new access token from the IAM Identity Center service. This token is granted if your IAM Identity Center access portal session is not expired.


## IAM Roles Anywhere

You can use IAM Roles Anywhere to get temporary security credentials in IAM for workloads such as servers, containers, and applications that run outside of Amazon. To use IAM Roles Anywhere, your workloads must use X.509 certificates. Your Cloud Administrator should provide the certificate and private key needed to configure IAM Roles Anywhere as your credential provider.

### Step 1: Configure IAM Roles Anywhere

IAM Roles Anywhere provides a way to get temporary credentials for a workload or process that runs outside of Amazon. A trust anchor is established with the certificate authority to get temporary credentials for the associated IAM role. The role sets the permissions your workload will have when your code authenticates with IAM Roles Anywhere.

For steps to set up the trust anchor, IAM role, and IAM Roles Anywhere profile, see [Creating a trust anchor and profile in Amazon Identity and Access Management Roles Anywhere](#) in the *IAM Roles Anywhere User Guide*.

 **Note**

A *profile* in the *IAM Roles Anywhere User Guide* refers to a unique concept within the IAM Roles Anywhere service. It's not related to the profiles within the shared Amazon config file.

## Step 2: Use IAM Roles Anywhere

To get temporary security credentials from IAM Roles Anywhere, use the credential helper tool provided by IAM Roles Anywhere. The credential tool implements the signing process for IAM Roles Anywhere.

For instructions to download the credential helper tool, see [Obtaining temporary security credentials from Amazon Identity and Access Management Roles Anywhere](#) in the *IAM Roles Anywhere User Guide*.

To use temporary security credentials from IAM Roles Anywhere with Amazon SDKs and the Amazon CLI, you can configure `credential_process` setting in the shared Amazon config file. The SDKs and Amazon CLI support a process credential provider that uses `credential_process` to authenticate. The following shows the general structure to set `credential_process`.

```
credential_process = [path to helper tool] [command] [--parameter1 value] [--parameter2 value] [...]
```

The `credential-process` command of the helper tool returns temporary credentials in a standard JSON format that is compatible with the `credential_process` setting. Note that the command name contains a hyphen but the setting name contains an underscore. The command requires the following parameters:

- `private-key` – The path to the private key that signed the request.
- `certificate` – The path to the certificate.
- `role-arn` – The ARN of the role to get temporary credentials for.

- `profile-arn` – The ARN of the profile that provides a mapping for the specified role.
- `trust-anchor-arn` – The ARN of the trust anchor used to authenticate.

Your Cloud Administrator should provide the certificate and private key. All three ARN values can be copied from the Amazon Web Services Management Console. The following example shows a shared config file that configures retrieving temporary credentials from the helper tool.

```
[profile dev]
credential_process = ./aws_signing_helper credential-process --certificate /
path/to/certificate --private-key /path/to/private-key --trust-anchor-
arn arn:aws:rolesanywhere:region:account:trust-anchor/TA_ID --profile-
arn arn:aws:rolesanywhere:region:account:profile/PROFILE_ID --role-
arn arn:aws:iam::account:role/ROLE_ID
```

For optional parameters and additional helper tool details, see [IAM Roles Anywhere Credential Helper](#) on GitHub.

For details on the SDK configuration setting itself and the process credential provider, see [Process credential provider](#) in this guide.

## Assume a role

Assuming a role involves using a set of temporary security credentials to access Amazon resources that you might not have access to otherwise. These temporary credentials consist of an access key ID, a secret access key, and a security token. To learn more about Amazon Security Token Service (Amazon STS) API requests, see [Actions](#) in the *Amazon Security Token Service API Reference*.

To set up your SDK or tool to assume a role, you must first create or identify a specific *role* to assume. IAM roles are uniquely identified by a role Amazon Resource Name ([ARN](#)). Roles establish trust relationships with another entity. The trusted entity that uses the role might be an Amazon Web Service, another Amazon Web Services account, a web identity provider or OIDC, or SAML federation. To learn more about IAM roles, see [Using IAM roles](#) in the *IAM User Guide*.

After the IAM role is identified, if you are trusted by that role, you can configure your SDK or tool to use the permissions that are granted by the role. To do this, either [Assume an IAM role](#) or [Federate with web identity or OpenID Connect](#).

## Assume an IAM role

When assuming a role, Amazon STS returns a set of temporary security credentials. These credentials are sourced from another profile or from the instance or container that your code is running in. Other examples of assuming a role include managing multiple Amazon Web Services accounts from Amazon EC2, using Amazon CodeCommit across Amazon Web Services accounts, or accessing another account from Amazon CodeBuild.

### Step 1: Set up an IAM role

To set up your SDK or tool to assume a role, you must first create or identify a specific role to assume. IAM roles are uniquely identified using a role [ARN](#). Roles establish trust relationships with another entity, typically within your account or for cross-account access. To set this up, see [Creating IAM roles](#) in the *IAM User Guide*.

### Step 2: Configure the SDK or tool

Configure the SDK or tool to source credentials from `credential_source` or `source_profile`.

Use `credential_source` to source credentials from an Amazon ECS container, an Amazon EC2 instance, or from environment variables.

Use `source_profile` to source credentials from another profile. `source_profile` also supports role chaining, which is hierarchies of profiles where an assumed role is then used to assume another role.

When you specify this in a profile, the SDK or tool automatically makes the corresponding Amazon STS [AssumeRole](#) API call for you. To retrieve and use temporary credentials by assuming a role, specify the following configuration values in the shared Amazon config file. For more details on each of these settings, see the [Assume role credential provider settings](#) section.

- `role_arn` - From the IAM role you created in Step 1
- Configure either `source_profile` or `credential_source`
- (Optional) `duration_seconds`
- (Optional) `external_id`
- (Optional) `mfa_serial`
- (Optional) `role_session_name`



The following examples show the configuration of both assume role options in a shared config file:

```
role_arn = arn:aws-cn:iam::123456789012:role/my-role-name
source_profile = profile-name-with-user-that-can-assume-role
```

```
role_arn = arn:aws-cn:iam::123456789012:role/my-role-name
credential_source = Ec2InstanceMetadata
```

For details on all assume role credential provider settings, see [Assume role credential provider](#) in this guide.

## Federate with web identity or OpenID Connect

When creating mobile applications or client-based web applications that require access to Amazon, Amazon STS returns a set of temporary security credentials for federated users who are authenticated through a public identity provider (IdP). Examples of public identity providers include Login with Amazon, Facebook, Google, or any OpenID Connect (OIDC)-compatible identity provider. With this method, your users don't need their own Amazon or IAM identities.

If you are using Amazon Elastic Kubernetes Service, this feature provides the ability to specify different IAM roles for each one of your containers. Kubernetes provides capability to distribute OIDC tokens to your containers which are used by this credential provider to obtain temporary credentials. For more information on this Amazon EKS configuration, see [IAM roles for service accounts](#) in the Amazon EKS User Guide. However, for a simpler option, we recommend you use [Amazon EKS Pod Identities](#) instead if your [SDK supports it](#).

### Step 1: Set up an identity provider and IAM role

To configure federation with an external IdP, use an IAM identity provider to inform Amazon about the external IdP and its configuration. This establishes *trust* between your Amazon Web Services account and the external IdP. Before configuring the SDK to use the web identity token for authentication, you must first set up the identity provider (IdP) and the IAM role used to access it. To set these up, see [Creating a role for web identity or OpenID Connect Federation \(console\)](#) in the *IAM User Guide*.

### Step 2: Configure the SDK or tool

Configure the SDK or tool to use a web identity token from Amazon STS for authentication.

When you specify this in a profile, the SDK or tool automatically makes the corresponding Amazon STS [AssumeRoleWithWebIdentity](#) API call for you. To retrieve and use temporary credentials using web identity federation, specify the following configuration values in the shared Amazon config file. For more details on each of these settings, see the [Assume role credential provider settings](#) section.

- `role_arn` - From the IAM role you created in Step 1
- `web_identity_token_file` - From the external IdP
- (Optional) `duration_seconds`
- (Optional) `role_session_name`

The following is an example of a shared config file configuration to assume a role with web identity:

```
[profile web-identity]  
role_arn=arn:aws-cn:iam::123456789012:role/my-role-name  
web_identity_token_file=/path/to/a/token
```

### Note

For mobile applications, consider using Amazon Cognito. Amazon Cognito acts as an identity broker and does much of the federation work for you. However, the Amazon Cognito identity provider isn't included in the SDKs and tools core libraries like other identity providers. To access the Amazon Cognito API, include the Amazon Cognito service client in the build or libraries for your SDK or tool. For usage with Amazon SDKs, see [Code Examples](#) in the *Amazon Cognito Developer Guide*.

For details on all assume role credential provider settings, see [Assume role credential provider](#) in this guide.

## Amazon access keys

### Use short-term credentials

We recommend configuring your SDK or tool to use [IAM Identity Center authentication](#) to use extended session duration options.

However, to set up the SDK or tool's temporary credentials directly, see [Authenticate using short-term credentials](#).

## Use long-term credentials

### Warning

To avoid security risks, don't use IAM users for authentication when developing purpose-built software or working with real data. Instead, use federation with an identity provider such as [Amazon IAM Identity Center](#).

## Manage access across Amazon Web Services accounts

As a security best practice, we recommend using Amazon Organizations with IAM Identity Center to manage access across all your Amazon Web Services accounts. For more information, see [Security best practices in IAM](#) in the *IAM User Guide*.

You can create users in IAM Identity Center, use Microsoft Active Directory, use a SAML 2.0 identity provider (IdP), or individually federate your IdP to Amazon Web Services accounts. Using one of these approaches, you can provide a single sign-on experience for your users. You can also enforce multi-factor authentication (MFA) and use temporary credentials for Amazon Web Services account access. This differs from an IAM user, which is a long-term credential that can be shared and which might increase the security risk to your Amazon resources.

## Create IAM users for sandbox environments only

If you're new to Amazon, you might create a test IAM user and then use it to run tutorials and explore what Amazon has to offer. It's okay to use this type of credential when you're learning, but we recommend that you avoid using it outside of a sandbox environment.

For the following use cases, it might make sense to get started with IAM users in Amazon:

- Getting started with your Amazon SDK or tool and exploring Amazon Web Services in a sandbox environment.
- Running scheduled scripts, jobs, and other automated processes that don't support a human-attended sign-in process as part of your learning.

If you're using IAM users outside of these use cases, then transition to IAM Identity Center or federate your identity provider to Amazon Web Services accounts as soon as possible. For more information, see [Identity federation in Amazon](#).

## Secure IAM user access keys

You should rotate IAM user access keys regularly. Follow the guidance in [Rotating access keys](#) in the *IAM User Guide*. If you believe that you have accidentally shared your IAM user access keys, then rotate your access keys.

IAM user access keys should be stored in the shared `Amazon credentials` file on the local machine. Don't store the IAM user access keys in your code. Don't include configuration files that contain your IAM user access keys inside of any source code management software. External tools, such as the open source project [git-secrets](#), can help you from inadvertently committing sensitive information to a Git repository. For more information, see [IAM Identities \(users, user groups, and roles\)](#) in the *IAM User Guide*.

To set up an IAM user to get started, see [Authenticate using long-term credentials](#).

## Authenticate using short-term credentials

We recommend configuring your SDK or tool to use [IAM Identity Center authentication](#) with extended session duration options. However, you can copy and use temporary credentials that are available in the Amazon access portal. New credentials will need to be copied when these expire. You can use the temporary credentials in a profile or use them as values for system properties and environment variables.

### Set up a credentials file using short-term credentials retrieved from Amazon access portal

1. [Create a shared credentials file](#).
2. In the credentials file, paste the following placeholder text until you paste in working temporary credentials.

```
[default]
aws_access_key_id=<value from Amazon access portal>
aws_secret_access_key=<value from Amazon access portal>
aws_session_token=<value from Amazon access portal>
```



Review the [Security best practices in IAM](#) topic and transition to IAM Identity Center or other temporary credentials as soon as possible.

## Important warnings and guidance for credentials

### Warnings for credentials

- **Do NOT** use your account's root credentials to access Amazon resources. These credentials provide unrestricted account access and are difficult to revoke.
- **Do NOT** put literal access keys or credential information in your application files. If you do, you create a risk of accidentally exposing your credentials if, for example, you upload the project to a public repository.
- **Do NOT** include files that contain credentials in your project area.
- Be aware that any credentials stored in the shared `Amazon credentials` file are stored in plaintext.

### Additional guidance for securely managing credentials

For a general discussion of how to securely manage Amazon credentials, see [Best practices for managing Amazon access keys](#) in the [Amazon Web Services General Reference](#). In addition to that discussion, consider the following:

- Use [IAM roles for tasks](#) for Amazon Elastic Container Service (Amazon ECS) tasks.
- Use [IAM roles](#) for applications that are running on Amazon EC2 instances.

### Prerequisites: Create an Amazon account

To use an IAM user to access Amazon services, you need an Amazon account and Amazon credentials.

#### 1. Create an account.

To create an Amazon account, see [Getting started: Are you a first-time Amazon user?](#) in the *Amazon Account Management Reference Guide*.

#### 2. Create an administrative user.

Avoid using your root user account (the initial account you create) to access the management console and services. Instead, create an administrative user account, as explained in [Create an administrative user](#) in the *IAM User Guide*.

After you create the administrative user account and record the login details, **be sure to sign out of your root user account** and sign back in using the administrative account.

Neither of these accounts are appropriate for doing development on Amazon or for running applications on Amazon. As a best practice, you need to create users, permission sets, or service roles that are appropriate for these tasks. For more information, see [Apply least-privilege permissions](#) in the *IAM User Guide*.

## Step 1: Create your IAM user

- Create your IAM user by following the [Creating IAM users \(console\)](#) procedure in the *IAM User Guide*.
  - For **Permission options**, choose **Attach policies directly** for how you want to assign permissions to this user.
    - Most "Getting Started" SDK tutorials use the Amazon S3 service as an example. To provide your application with full access to Amazon S3, select the AmazonS3FullAccess policy to attach to this user.
  - You can ignore optional steps of that procedure.

## Step 2: Get your access keys

1. In the navigation pane of the IAM console, select **Users** and then select the **User name** of the user that you created previously.
2. On the user's page, select the **Security credentials** page. Then, under **Access keys**, select **Create access key**.
3. For **Create access key Step 1**, choose either **Command Line Interface (CLI)** or **Local code**. Both options generate the same type of key to use with both the Amazon CLI and the SDKs.
4. For **Create access key Step 2**, enter an optional tag and select **Next**.
5. For **Create access key Step 3**, select **Download .csv file** to save a .csv file with your IAM user's access key and secret access key. You need this information for later.

**⚠ Warning**

Use appropriate security measures to keep these credentials safe.

6. Select **Done**.

### Step 3: Update the shared credentials file

1. Create or open the shared Amazon credentials file. This file is `~/.aws/credentials` on Linux and macOS systems, and `%USERPROFILE%\.aws\credentials` on Windows. For more information, see [Location of Credentials Files](#).
2. Add the following text to the shared credentials file. Replace the example ID value and example key value with the values in the `.csv` file that you downloaded earlier.

```
[default]
aws_access_key_id = AKIAIOSFODNN7EXAMPLE
aws_secret_access_key = wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
```

3. Save the file.

The shared `credentials` file is the most common way to store credentials. These can also be set as environment variables, see [Amazon access keys](#) for environment variable names. This is a way to get you started, but we recommend you transition to IAM Identity Center or other temporary credentials as soon as possible. After you transition away from using long-term credentials, remember to delete these credentials from the shared `credentials` file.

## Using IAM roles for Amazon EC2 instances

This example covers setting up an Amazon Identity and Access Management role with Amazon S3 access to use in your application deployed to an Amazon EC2 instance.

For an Amazon Elastic Compute Cloud instance, create an IAM role, and then give your Amazon EC2 instance access to that role. For more information, see [IAM Roles for Amazon EC2](#) in the *Amazon EC2 User Guide for Linux Instances* or [IAM Roles for Amazon EC2](#) in the *Amazon EC2 User Guide for Windows Instances*.



## Create an IAM role

Create an IAM role that grants read-only access to Amazon S3.

1. Sign in to the Amazon Web Services Management Console and open the IAM console at <https://console.amazonaws.cn/iam/>.
2. In the navigation pane, select **Roles**, then select **Create role**.
3. For **Select trusted entity**, under **Trusted entity type**, choose **Amazon Web Service**.
4. Under **Use case**, choose **Amazon EC2**, then select **Next**.
5. For **Add permissions**, select the checkbox for **Amazon S3 Read Only Access** from the policy list, then select **Next**.
6. Enter a name for the role, then select **Create role**. *Remember this name because you'll need it when you launch your Amazon EC2 instance.*

## Launch an Amazon EC2 instance and specify your IAM role

You can launch an Amazon EC2 instance with an IAM role using the Amazon EC2 console.

Follow the directions to launch an instance in the [Amazon EC2 User Guide for Linux Instances](#) or the [Amazon EC2 User Guide for Windows Instances](#).

When you reach the **Review Instance Launch** page, select **Edit instance details**. In **IAM role**, choose the IAM role that you created previously. Complete the procedure as directed.

### Note

You need to create or use an existing security group and key pair to connect to the instance.

With this IAM and Amazon EC2 setup, you can deploy your application to the Amazon EC2 instance and it will have read access to the Amazon S3 service.

## Connect to the EC2 instance

Connect to the EC2 instance so that you can transfer the sample application to it and then run the application. You'll need the file that contains the private portion of the key pair you used to launch the instance; that is, the PEM file.

You can do this by following the connect procedure in the [Amazon EC2 User Guide for Linux Instances](#) or the [Amazon EC2 User Guide for Windows Instances](#). When you connect, do so in such a way that you can transfer files from your development machine to your instance.

If you're using an Amazon Toolkit, you can often also connect to the instance by using the Toolkit. For more information, see the specific user guide for the Toolkit you use.

## Run the sample application on the EC2 instance

1. Copy the application files from your local drive to your instance.

For information about how to transfer files to your instance see the [Amazon EC2 User Guide for Linux Instances](#) or the [Amazon EC2 User Guide for Windows Instances](#).

2. Start the application and verify that it runs with the same results as on your development machine.
3. (Optional) Verify that the application uses the credentials provided by the IAM role.
  - a. Sign in to the Amazon Web Services Management Console and open the Amazon EC2 console at <https://console.amazonaws.cn/ec2/>.
  - b. Select the instance and detach the IAM role through **Actions, Instance Settings, Attach/Replace IAM Role**.
  - c. Run the application again and confirm that it returns an authorization error.

## Settings reference

SDKs provide language-specific APIs for Amazon Web Services. They take care of some of the heavy lifting necessary in successfully making API calls, including authentication, retry behavior, and more. To do this, the SDKs have flexible strategies to obtain credentials to use for your requests, to maintain settings to use with each service, and to obtain values to use for global settings.

You can find detailed information about configuration settings in the following sections:

- [Standardized credential providers](#) – Common credential providers standardized across multiple SDKs.
- [Standardized features](#) – Common features standardized across multiple SDKs.

## Creating service clients

To programmatically access Amazon Web Services, SDKs use a client class/object for each Amazon Web Service. For example, if your application needs to access Amazon EC2, your application creates an Amazon EC2 client object to interface with that service. You then use the service client to make requests to that Amazon Web Service. In most SDKs, a service client object is immutable, so you must create a new client for each service to which you make requests and for making requests to the same service using a different configuration.

## Precedence of settings

Global settings configure features, credential providers, and other functionality that are supported by most SDKs and have a broad impact across Amazon Web Services. All SDKs have a series of places (or sources) that they check in order to find a value for global settings. The following is the setting lookup precedence:

1. Any explicit setting set in the code or on a service client itself takes precedence over anything else.
  - Some settings can be set on a per-operation basis, and can be changed as needed for each operation that you invoke. For the Amazon CLI or Amazon Tools for PowerShell, these take the form of per-operation parameters that you enter on the command line. For an SDK, explicit

assignments can take the form of a parameter that you set when you instantiate an Amazon Web Service client or configuration object, or sometimes when you call an individual API.

2. Java/Kotlin only: The JVM system property for the setting is checked. If it's set, that value is used to configure the client.
3. The environment variable is checked. If it's set, that value is used to configure the client.
4. The SDK checks the shared `credentials` file for the setting. If it's set, the client uses it.
5. The shared `config` file for the setting. If the setting is present, the SDK uses it.
  - The `AWS_PROFILE` environment variable or the `aws.profile` JVM system property can be used to specify which profile that the SDK loads.
6. Any default value provided by the SDK source code itself is used last.

### Note

Some SDKs and tools might check in a different order. Also, some SDKs and tools support other methods of storing and retrieving parameters. For example, the Amazon SDK for .NET supports an additional source called the [SDK Store](#). For more information about providers that are unique to a SDK or tool, see the specific guide for the SDK or tool that you are using.

The order determines which methods take precedence and override others. For example, if you set up a profile in the shared `config` file, it's only found and used after the SDK or tool checks the other places first. This means that if you put a setting in the `credentials` file, it is used instead of one found in the `config` file. If you configure an environment variable with a setting and value, it would override that setting in both the `credentials` and `config` files. And finally, a setting on the individual operation (Amazon CLI command-line parameter or API parameter) or in code would override all other values for that one command.

## Config file settings list

The settings listed in the following table can be assigned in the shared Amazon `config` file. They are global and affect all Amazon Web Services.

Setting name	Details
api_versions	<a href="#">General configuration settings</a>
aws_access_key_id	<a href="#">Amazon access keys</a>
aws_secret_access_key	<a href="#">Amazon access keys</a>
aws_session_token	<a href="#">Amazon access keys</a>
ca_bundle	<a href="#">General configuration settings</a>
credential_process	<a href="#">Process credential provider</a>
credential_source	<a href="#">Assume role credential provider</a>
defaults_mode	<a href="#">Smart configuration defaults</a>
disable_request_compression	<a href="#">Request compression</a>
duration_seconds	<a href="#">Assume role credential provider</a>
ec2_metadata_service_endpoint	<a href="#">IMDS credential provider</a>
ec2_metadata_service_endpoint_mode	<a href="#">IMDS credential provider</a>

Setting name	Details
ec2_metadata_v1_disabled	<a href="#">IMDS credential provider</a>
endpoint_discovery_enabled	<a href="#">Endpoint discovery</a>
endpoint_url	<a href="#">Service-specific endpoints</a>
external_id	<a href="#">Assume role credential provider</a>
ignore_configured_endpoint_urls	<a href="#">Service-specific endpoints</a>
max_attempts	<a href="#">Retry behavior</a>
metadata_service_num_attempts	<a href="#">Amazon EC2 instance metadata</a>
metadata_service_timeout	<a href="#">Amazon EC2 instance metadata</a>
mfa_serial	<a href="#">Assume role credential provider</a>
output	<a href="#">General configuration settings</a>
parameter_validation	<a href="#">General configuration settings</a>
region	<a href="#">Amazon Web Services Region</a>

Setting name	Details
request_m in_compre ssion_siz e_bytes	<a href="#">Request compression</a>
retry_mode	<a href="#">Retry behavior</a>
role_arn	<a href="#">Assume role credential provider</a>
role_sess ion_name	<a href="#">Assume role credential provider</a>
s3_disabl e_multire gion_acce ss_points	<a href="#">Amazon S3 Multi-Region Access Points</a>
s3_use_ar n_region	<a href="#">Amazon S3 access points</a>
source_profile	<a href="#">Assume role credential provider</a>
sso_account_id	<a href="#">IAM Identity Center credential provider</a>
sso_region	<a href="#">IAM Identity Center credential provider</a>
sso_regis tration_scopes	<a href="#">IAM Identity Center credential provider</a>
sso_role_name	<a href="#">IAM Identity Center credential provider</a>
sso_start_url	<a href="#">IAM Identity Center credential provider</a>
sts_regio nal_endpoints	<a href="#">Amazon STS Regionalized endpoints</a>
use_duals tack_endpoint	<a href="#">Dual-stack and FIPS endpoints</a>

Setting name	Details
use_fips_endpoint	<a href="#">Dual-stack and FIPS endpoints</a>
web_identity_token_file	<a href="#">Assume role credential provider</a>

## Credentials file settings list

The settings listed in the following table can be assigned in the shared Amazon credentials file. They are global and affect all Amazon Web Services.

Setting name	Details
aws_access_key_id	<a href="#">Amazon access keys</a>
aws_secret_access_key	<a href="#">Amazon access keys</a>
aws_session_token	<a href="#">Amazon access keys</a>

## Environment variables list

Environment variables supported by most SDKs are listed in the following table. They are global and affect all Amazon Web Services.

Setting name	Details
AWS_ACCESS_KEY_ID	<a href="#">Amazon access keys</a>
AWS_CA_BUNDLE	<a href="#">General configuration settings</a>



Setting name	Details
AWS_CONFIG_FILE	<a href="#">Location of the shared config and credentials files</a>
AWS_CONTAINER_AUTHORIZATION_TOKEN	<a href="#">Container credential provider</a>
AWS_CONTAINER_AUTHORIZATION_TOKEN_FILE	<a href="#">Container credential provider</a>
AWS_CONTAINER_CREDENTIALS_FULL_URI	<a href="#">Container credential provider</a>
AWS_CONTAINER_CREDENTIALS_RELATIVE_URI	<a href="#">Container credential provider</a>
AWS_DEFAULTS_MODE	<a href="#">Smart configuration defaults</a>
AWS_DISABLE_REQUEST_COMPRESSION	<a href="#">Request compression</a>
AWS_EC2_METADATA_DISABLED	<a href="#">IMDS credential provider</a>

Setting name	Details
AWS_EC2_M ETADATA_S ERVICE_EN DPOINT	<a href="#">IMDS credential provider</a>
AWS_EC2_M ETADATA_S ERVICE_EN DPOINT_MODE	<a href="#">IMDS credential provider</a>
AWS_EC2_M ETADATA_V 1_DISABLED	<a href="#">IMDS credential provider</a>
AWS_ENABL E_ENDPOIN T_DISCOVERY	<a href="#">Endpoint discovery</a>
AWS_ENDPO INT_URL	<a href="#">Service-specific endpoints</a>
AWS_ENDPO INT_URL_< SERVICE>	<a href="#">Service-specific endpoints</a>
AWS_IAM_R OLE_ARN	<a href="#">Assume role credential provider</a>
AWS_IAM_R OLE_SESSI ON_NAME	<a href="#">Assume role credential provider</a>
AWS_IGNORE CONFIGU RED_ENDPO INT_URLS	<a href="#">Service-specific endpoints</a>

Setting name	Details
AWS_MAX_A TTEMPTS	<a href="#">Retry behavior</a>
AWS_METAD ATA_SERVI CE_NUM_AT TEMPTS	<a href="#">Amazon EC2 instance metadata</a>
AWS_METAD ATA_SERVI CE_TIMEOUT	<a href="#">Amazon EC2 instance metadata</a>
AWS_PROFILE	<a href="#">Shared config and credentials files</a>
AWS_REGION	<a href="#">Amazon Web Services Region</a>
AWS_REQUE ST_MIN_CO MPRESSION _SIZE_BYTES	<a href="#">Request compression</a>
AWS_RETRY_MODE	<a href="#">Retry behavior</a>
AWS_S3_DI SABLE_MUL TIREGION_ ACCESS_POINTS	<a href="#">Amazon S3 Multi-Region Access Points</a>
AWS_S3_US E_ARN_REGION	<a href="#">Amazon S3 access points</a>
AWS_SECRE T_ACCESS_KEY	<a href="#">Amazon access keys</a>
AWS_SESSI ON_TOKEN	<a href="#">Amazon access keys</a>

Setting name	Details
AWS_SHARED_CREDENTIALS_FILE	<a href="#">Location of the shared config and credentials files</a>
AWS_STS_REGIONAL_ENDPOINTS	<a href="#">Amazon STS Regionalized endpoints</a>
AWS_USE_DUALSTACK_ENDPOINT	<a href="#">Dual-stack and FIPS endpoints</a>
AWS_USE_FIPS_ENDPOINT	<a href="#">Dual-stack and FIPS endpoints</a>
AWS_WEB_IDENTITY_TOKEN_FILE	<a href="#">Assume role credential provider</a>

## JVM system properties list

You can use the following JVM system properties for the Amazon SDK for Java and the Amazon SDK for Kotlin (targeting the JVM). See [the section called “How to set JVM system properties”](#) for instructions on how to set JVM system properties.

Setting name	Details
<code>aws.accessKeyId</code>	<a href="#">Amazon access keys</a>
<code>aws.configFile</code>	<a href="#">Location of the shared config and credentials files</a>
<code>aws.defaultMode</code>	<a href="#">Smart configuration defaults</a>

Setting name	Details
<code>aws.disableEc2MetadataV1</code>	<a href="#">IMDS credential provider</a>
<code>aws.disableRequestCompression</code>	<a href="#">Request compression</a>
<code>aws.ec2MetadataServiceEndpoint</code>	<a href="#">IMDS credential provider</a>
<code>aws.ec2MetadataServiceEndpointMode</code>	<a href="#">IMDS credential provider</a>
<code>aws.endpointDiscoveryEnabled</code>	<a href="#">Endpoint discovery</a>
<code>aws.endpointUrl</code>	<a href="#">Service-specific endpoints</a>
<code>aws.endpointUrl&lt;ServiceName&gt;</code>	<a href="#">Service-specific endpoints</a>
<code>aws.ignoreConfiguredEndpointUrls</code>	<a href="#">Service-specific endpoints</a>
<code>aws.maxAttempts</code>	<a href="#">Retry behavior</a>
<code>aws.profile</code>	<a href="#">Shared config and credentials files</a>
<code>aws.region</code>	<a href="#">Amazon Web Services Region</a>

Setting name	Details
<code>aws.requestMinCompressionSizeBytes</code>	<a href="#">Request compression</a>
<code>aws.retryMode</code>	<a href="#">Retry behavior</a>
<code>aws.roleArn</code>	<a href="#">Assume role credential provider</a>
<code>aws.roleSessionName</code>	<a href="#">Assume role credential provider</a>
<code>aws.s3DisableMultiRegionAccessPoints</code>	<a href="#">Amazon S3 Multi-Region Access Points</a>
<code>aws.s3UseArnRegion</code>	<a href="#">Amazon S3 access points</a>
<code>aws.secretAccessKey</code>	<a href="#">Amazon access keys</a>
<code>aws.sessionToken</code>	<a href="#">Amazon access keys</a>
<code>aws.sharedCredentialsFile</code>	<a href="#">Location of the shared config and credentials files</a>
<code>aws.useDualstackEndpoint</code>	<a href="#">Dual-stack and FIPS endpoints</a>
<code>aws.useFipsEndpoint</code>	<a href="#">Dual-stack and FIPS endpoints</a>

Setting name	Details
<code>aws.webId</code> <code>entityTokenFile</code>	<a href="#">Assume role credential provider</a>

## Standardized credential providers

Many credential providers have been standardized to consistent defaults and to work the same way across many SDKs. This consistency increases productivity and clarity when coding across multiple SDKs. All settings can be overridden in code. For details, see your specific SDK API.

### Important

Not all SDKs support all providers, or even all aspects within a provider.

### Topics

- [Credential provider chain](#)
- [Amazon access keys](#)
- [Assume role credential provider](#)
- [Container credential provider](#)
- [IAM Identity Center credential provider](#)
- [IMDS credential provider](#)
- [Process credential provider](#)

## Credential provider chain

All SDKs have a series of places (or sources) that they check in order to find valid credentials to use to make a request to an Amazon Web Service. After valid credentials are found, the search is stopped. This systematic search is called the default credential provider chain.

Although the distinct chain used by each SDK varies, they most often include sources such as the following:

Credential provider	Description
<a href="#">Amazon access keys</a>	Amazon access keys for an IAM user (such as <code>AWS_ACCESS_KEY_ID</code> , and <code>AWS_SECRET_ACCESS_KEY</code> ).
<a href="#">Federate with web identity or OpenID Connect</a> - Assume role credential provider	Sign in using a well-known external identity provider (IdP), such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP. Assume the permissions of an IAM role using a web identity token from Amazon Security Token Service (Amazon STS).
<a href="#">IAM Identity Center credential provider</a>	Get credentials from Amazon IAM Identity Center.
<a href="#">Assume role credential provider</a>	Get access to other resources by assuming the permissions of an IAM role. (Retrieve and then use temporary credentials for a role).
<a href="#">Container credential provider</a>	Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS) credentials. The container credential provider fetches credentials for the customer's containerized application.
<a href="#">Process credential provider</a>	Custom credential provider. Get your credentials from an external source or process, including IAM Roles Anywhere.
<a href="#">IMDS credential provider</a>	Amazon Elastic Compute Cloud (Amazon EC2) instance profile credentials. Associate an IAM role with each of your EC2 instances. Temporary credentials for that role are made available to code running in the instance. The credentials are delivered through the Amazon EC2 metadata service.



For each step in the chain, there are multiple ways to assign setting values. Setting values that are specified in code always take precedence. However, there are also [Environment variables](#) and the [Shared config and credentials files](#). For more information, see [Precedence of settings](#).

## Amazon access keys

### Warning

To avoid security risks, don't use IAM users for authentication when developing purpose-built software or working with real data. Instead, use federation with an identity provider such as [Amazon IAM Identity Center](#).

Amazon access keys for an IAM user can be used as your Amazon credentials. The Amazon SDK automatically uses these Amazon credentials to sign API requests to Amazon, so that your workloads can access your Amazon resources and data securely and conveniently. It is recommended to always use the `aws_session_token` so that the credentials are temporary and no longer valid after they expire. Using long term credentials is not recommended.

### Note

If Amazon becomes unable to refresh these temporary credentials, Amazon may extend the validity of the credentials so that your workloads are not impacted.

The shared `Amazon credentials` file is the recommended location for storing credentials information because it is safely outside of application source directories and separate from the SDK-specific settings of the `shared config` file.

To learn more about Amazon credentials and using access keys, see [Amazon security credentials](#) and [Managing access keys for IAM users](#) in the *IAM User Guide*.

Configure this functionality by using the following:

**`aws_access_key_id` - shared Amazon config file setting, `aws_access_key_id` - shared Amazon credentials file setting (*recommended method*), `AWS_ACCESS_KEY_ID` - environment variable, `aws.accessKeyId` - JVM system property: Java/Kotlin only**

Specifies the Amazon access key used as part of the credentials to authenticate the user.

**aws\_secret\_access\_key** - shared Amazon config file setting, **aws\_secret\_access\_key** - shared Amazon credentials file setting (*recommended method*), **AWS\_SECRET\_ACCESS\_KEY** - environment variable, **aws.secretAccessKey** - JVM system property: Java/Kotlin only

Specifies the Amazon secret key used as part of the credentials to authenticate the user.

**aws\_session\_token** - shared Amazon config file setting, **aws\_session\_token** - shared Amazon credentials file setting (*recommended method*), **AWS\_SESSION\_TOKEN** - environment variable, **aws.sessionToken** - JVM system property: Java/Kotlin only

Specifies an Amazon session token used as part of the credentials to authenticate the user. You receive this value as part of the temporary credentials returned by successful requests to assume a role. A session token is required only if you manually specify temporary security credentials. However, we recommend you always use temporary security credentials instead of long-term credentials. For security recommendations, see [Security best practices in IAM](#).

For instructions on how to obtain these values, see [Authenticate using short-term credentials](#).

Example of setting these required values in the config or credentials file:

```
[default]
aws_access_key_id = AKIAIOSFODNN7EXAMPLE
aws_secret_access_key = wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
aws_session_token = AQoEXAMPLEH4aoAH0gNCAPy...truncated...zrkuWJ0gQs8IZZaIv2BXIa2R40lgk
```

Linux/macOS example of setting environment variables via command line:

```
export AWS_ACCESS_KEY_ID=AKIAIOSFODNN7EXAMPLE
export AWS_SECRET_ACCESS_KEY=wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
export
AWS_SESSION_TOKEN=AQoEXAMPLEH4aoAH0gNCAPy...truncated...zrkuWJ0gQs8IZZaIv2BXIa2R40lgk
```

Windows example of setting environment variables via command line:

```
setx AWS_ACCESS_KEY_ID AKIAIOSFODNN7EXAMPLE
setx AWS_SECRET_ACCESS_KEY wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
setx
AWS_SESSION_TOKEN AQoEXAMPLEH4aoAH0gNCAPy...truncated...zrkuWJ0gQs8IZZaIv2BXIa2R40lgk
```

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Supported	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	Yes	shared config file not supported.
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	Yes	To use shared config file settings, you must turn on loading from the config file; see <a href="#">Sessions</a> .
<a href="#">SDK for Java 2.x</a>	Yes	
<a href="#">SDK for Java 1.x</a>	Yes	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	Yes	
<a href="#">SDK for Kotlin</a>	Yes	
<a href="#">SDK for .NET 3.x</a>	Yes	Environment variables not supported.
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	Yes	
<a href="#">Tools for PowerShell</a>	Yes	Environment variables not supported.

## Assume role credential provider

Assuming a role involves using a set of temporary security credentials to access Amazon resources that you might not have access to otherwise. These temporary credentials consist of an access key ID, a secret access key, and a security token.

To set up your SDK or tool to assume a role, you must first create or identify a specific *role* to assume. IAM roles are uniquely identified by a role Amazon Resource Name ([ARN](#)). Roles establish trust relationships with another entity. The trusted entity that uses the role might be an Amazon Web Service, another Amazon Web Services account, a web identity provider or OIDC, or SAML federation.

After the IAM role is identified, if you are trusted by that role, you can configure your SDK or tool to use the permissions that are granted by the role. To do this, use the following settings.

For guidance on getting started using these settings, see [Assume a role](#) in this guide.

### Assume role credential provider settings

Configure this functionality by using the following:

#### `credential_source` - shared Amazon config file setting

Used within Amazon EC2 instances or Amazon Elastic Container Service containers to specify where the SDK or tool can find credentials that have permission to assume the role that you specify with the `role_arn` parameter.

**Default value:** None

**Valid values:**

- **Environment** – Specifies that the SDK or tool is to retrieve source credentials from the environment variables [AWS\\_ACCESS\\_KEY\\_ID](#) and [AWS\\_SECRET\\_ACCESS\\_KEY](#).
- **Ec2InstanceMetadata** – Specifies that the SDK or tool is to use the [IAM role attached to the EC2 instance profile](#) to get source credentials.
- **EcsContainer** – Specifies that the SDK or tool is to use the [IAM role attached to the ECS container](#) to get source credentials.

You cannot specify both `credential_source` and `source_profile` in the same profile.

Example of setting this in a config file to indicate that credentials should be sourced from Amazon EC2:

```
credential_source = Ec2InstanceMetadata
role_arn = arn:aws-cn:iam::123456789012:role/my-role-name
```

### **duration\_seconds** - shared Amazon config file setting

Specifies the maximum duration of the role session, in seconds.

This setting applies only when the profile specifies to assume a role.

**Default value:** 3600 seconds (one hour)

**Valid values:** The value can range from 900 seconds (15 minutes) up to the maximum session duration setting configured for the role (which can be a maximum of 43200 seconds, or 12 hours). For more information, see [View the Maximum Session Duration Setting for a Role](#) in the *IAM User Guide*.

Example of setting this in a config file:

```
duration_seconds = 43200
```

### **external\_id** - shared Amazon config file setting

Specifies a unique identifier that is used by third parties to assume a role in their customers' accounts.

This setting applies only when the profile specifies to assume a role and the trust policy for the role requires a value for ExternalId. The value maps to the ExternalId parameter that is passed to the AssumeRole operation when the profile specifies a role.

**Default value:** None.

**Valid values:** See [How to use an External ID When Granting Access to Your Amazon Resources to a Third Party](#) in the *IAM User Guide*.

Example of setting this in a config file:

```
external_id = unique_value_assigned_by_3rd_party
```

### **mfa\_serial** - shared Amazon config file setting

Specifies the identification or serial number of a multi-factor authentication (MFA) device that the user must use when assuming a role.

Required when assuming a role where the trust policy for that role includes a condition that requires MFA authentication.

**Default value:** None.

**Valid values:** The value can be either a serial number for a hardware device (such as GAHT12345678), or an Amazon Resource Name (ARN) for a virtual MFA device. For more information about MFA, see [Configuring MFA-Protected API Access](#) in the *IAM User Guide*.

Example of setting this in a config file:

```
mfa_serial = arn:aws-cn:iam::123456789012:mfa/my-user-name
```

**role\_arn** - shared Amazon config file setting, **AWS\_IAM\_ROLE\_ARN** - environment variable, **aws.roleArn** - JVM system property: Java/Kotlin only

Specifies the Amazon Resource Name (ARN) of an IAM role that you want to use to perform operations requested using this profile.

**Default value:** None.

**Valid values:** The value must be the ARN of an IAM role, formatted as follows: `arn:aws-cn:iam::account-id:role/role-name`

In addition, you must also specify **one** of the following settings:

- `source_profile` – To identify another profile to use to find credentials that have permission to assume the role in this profile.
- `credential_source` – To use either credentials identified by the current environment variables or credentials attached to an Amazon EC2 instance profile, or an Amazon ECS container instance.
- `web_identity_token_file` – To use public identity providers or any OpenID Connect (OIDC)-compatible identity provider for users who have been authenticated in a mobile or web application.

**role\_session\_name** - shared Amazon config file setting, **AWS\_IAM\_ROLE\_SESSION\_NAME** - environment variable, **aws.roleSessionName** - JVM system property: Java/Kotlin only

Specifies the name to attach to the role session. This name appears in Amazon CloudTrail logs for entries associated with this session, which can be useful when auditing.

**Default value:** An optional parameter. If you don't provide this value, a session name is generated automatically if the profile assumes a role.

**Valid values:** Provided to the `RoleSessionName` parameter when the Amazon CLI or Amazon API calls the `AssumeRole` operation (or operations such as the `AssumeRoleWithWebIdentity` operation) on your behalf. The value becomes part of the assumed role user Amazon Resource Name (ARN) that you can query, and shows up as part of the CloudTrail log entries for operations invoked by this profile.

```
arn:aws-cn:sts::123456789012:assumed-role/my-role-name/my-role_session_name.
```

Example of setting this in a config file:

```
role_session_name = my-role-session-name
```

### **source\_profile - shared Amazon config file setting**

Specifies another profile whose credentials are used to assume the role specified by the `role_arn` setting in the original profile. To understand how profiles are used in the shared Amazon config and credentials files, see [Shared config and credentials files](#).

If you specify a profile that is also an assume role profile, each role will be assumed in sequential order to fully resolve the credentials. This chain is stopped when the SDK encounters a profile with credentials. Role chaining limits your Amazon CLI or Amazon API role session to a maximum of one hour and can't be increased. For more information, see [Roles terms and concepts](#) in the *IAM User Guide*.

**Default value:** None.

**Valid values:** A text string that consists of the name of a profile defined in the config and credentials files. You must also specify a value for `role_arn` in the current profile.

You cannot specify both `credential_source` and `source_profile` in the same profile.

Example of setting this in a config file:

```
[profile A]  
source_profile = B  
role_arn = arn:aws-cn:iam::123456789012:role/RoleA  
  
[profile B]
```





SDK	Support	Notes or more information
<a href="#">SDK for C++</a>	Partial	<code>credential_source</code> not supported. <code>duration_seconds</code> not supported. <code>mfa_serial</code> not supported.
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	Yes	To use shared config file settings, you must turn on loading from the config file; see <a href="#">Sessions</a> .
<a href="#">SDK for Java 2.x</a>	Partial	<code>mfa_serial</code> not supported. Use <code>AWS_ROLE_ARN</code> instead of <code>AWS_IAM_ROLE_ARN</code> . Use <code>AWS_ROLE_SESSION_NAME</code> instead of <code>AWS_IAM_ROLE_SESSION_NAME</code> .
<a href="#">SDK for Java 1.x</a>	Partial	<code>mfa_serial</code> not supported.
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	Partial	<code>credential_source</code> not supported.
<a href="#">SDK for Kotlin</a>	Yes	Use <code>AWS_ROLE_ARN</code> instead of <code>AWS_IAM_ROLE_ARN</code> . Use <code>AWS_ROLE_SESSION_NAME</code> instead of <code>AWS_IAM_ROLE_SESSION_NAME</code> .
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	Yes	
<a href="#">Tools for PowerShell</a>	Yes	

## Container credential provider

The container credential provider fetches credentials for customer's containerized application. This credential provider is useful for Amazon Elastic Container Service (Amazon ECS) and Amazon

Elastic Kubernetes Service (Amazon EKS) customers. SDKs attempt to load credentials from the specified HTTP endpoint through a GET request.

If you use Amazon ECS, we recommend you use a task IAM Role for improved credential isolation, authorization, and auditability. When configured, Amazon ECS sets the `AWS_CONTAINER_CREDENTIALS_RELATIVE_URI` environment variable that the SDKs and tools use to obtain credentials. To configure Amazon ECS for this functionality, see [Task IAM role](#) in the *Amazon Elastic Container Service Developer Guide*.

If you use Amazon EKS, we recommend you use Amazon EKS Pod Identity for improved credential isolation, least privilege, auditability, independent operation, reusability, and scalability. Both your Pod and an IAM role are associated with a Kubernetes service account to manage credentials for your applications. To learn more on Amazon EKS Pod Identity, see [Amazon EKS Pod Identities](#) in the Amazon EKS User Guide. When configured, Amazon EKS sets the `AWS_CONTAINER_CREDENTIALS_FULL_URI` and `AWS_CONTAINER_AUTHORIZATION_TOKEN_FILE` environment variables that the SDKs and tools use to obtain credentials. For setup information, see [Setting up the Amazon EKS Pod Identity Agent](#) in the Amazon EKS User Guide or [Amazon EKS Pod Identity simplifies IAM permissions for applications on Amazon EKS clusters](#) at the Amazon Blog website.

Configure this functionality by using the following:

### **AWS\_CONTAINER\_CREDENTIALS\_FULL\_URI - environment variable**

Specifies the full HTTP URL endpoint for the SDK to use when making a request for credentials. This includes both the scheme and the host.

**Default value:** None.

**Valid values:** Valid URI.

*Note: This setting is an alternative to `AWS_CONTAINER_CREDENTIALS_RELATIVE_URI` and will only be used if `AWS_CONTAINER_CREDENTIALS_RELATIVE_URI` is not set.*

Linux/macOS example of setting environment variables via command line:

```
export AWS_CONTAINER_CREDENTIALS_FULL_URI=http://localhost/get-credentials
```

or

```
export AWS_CONTAINER_CREDENTIALS_FULL_URI=http://localhost:8080/get-credentials
```

### **AWS\_CONTAINER\_CREDENTIALS\_RELATIVE\_URI - environment variable**

Specifies the relative HTTP URL endpoint for the SDK to use when making a request for credentials. The value is appended to the default Amazon ECS hostname of 169.254.170.2.

**Default value:** None.

**Valid values:** Valid relative URI.

Linux/macOS example of setting environment variables via command line:

```
export AWS_CONTAINER_CREDENTIALS_RELATIVE_URI=/get-credentials?a=1
```

### **AWS\_CONTAINER\_AUTHORIZATION\_TOKEN - environment variable**

Specifies an authorization token in plain text. If this variable is set, the SDK will set the Authorization header on the HTTP request with the environment variable's value.

**Default value:** None.

**Valid values:** String.

*Note: This setting is an alternative to AWS\_CONTAINER\_AUTHORIZATION\_TOKEN\_FILE and will only be used if AWS\_CONTAINER\_AUTHORIZATION\_TOKEN\_FILE is not set.*

Linux/macOS example of setting environment variables via command line:

```
export AWS_CONTAINER_CREDENTIALS_FULL_URI=http://localhost/get-credential  
export AWS_CONTAINER_AUTHORIZATION_TOKEN=Basic abcd
```

### **AWS\_CONTAINER\_AUTHORIZATION\_TOKEN\_FILE - environment variable**

Specifies an absolute file path to a file that contains the authorization token in plain text.

**Default value:** None.

**Valid values:** String.

Linux/macOS example of setting environment variables via command line:

```
export AWS_CONTAINER_CREDENTIALS_FULL_URI=http://localhost/get-credential
export AWS_CONTAINER_AUTHORIZATION_TOKEN_FILE=/path/to/token
```

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Support	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	Yes	
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	Yes	
<a href="#">SDK for Java 2.x</a>	Yes	
<a href="#">SDK for Java 1.x</a>	Parti	Amazon EKS Pod Identity and AWS_CONTAINER_AUTHORIZATION_TOKEN_FILE not supported.
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	Parti	Amazon EKS Pod Identity and AWS_CONTAINER_AUTHORIZATION_TOKEN_FILE not supported.
<a href="#">SDK for Kotlin</a>	Yes	
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	

SDK	Support	Notes or more information
<a href="#">SDK for Rust</a>	Partial	Amazon EKS Pod Identity and AWS_CONTAINER_AUTHORIZATION_TOKEN_FILE not supported.
<a href="#">Tools for PowerShell</a>	Yes	

## IAM Identity Center credential provider

This authentication mechanism uses Amazon IAM Identity Center to get single sign-on (SSO) access to Amazon Web Services for your code.

### Note

In the Amazon SDK API documentation, the IAM Identity Center credential provider is called the SSO credential provider.

After you enable IAM Identity Center, you define a profile for its settings in your shared Amazon config file. This profile is used to connect to the IAM Identity Center access portal. When a user successfully authenticates with IAM Identity Center, the portal returns short-term credentials for the IAM role associated with that user. To learn how the SDK gets temporary credentials from the configuration and uses them for Amazon Web Service requests, see [Understand IAM Identity Center authentication](#).

There are two ways to configure IAM Identity Center through the config file:

- **SSO token provider configuration (recommended)** – Extended session durations.
- **Legacy non-refreshable configuration** – Uses a fixed, eight-hour session.

In both configurations, you need to sign in again when your session expires.

To set custom session durations, you must use the SSO token provider configuration.

The following two guides contain additional information about IAM Identity Center:

- [Amazon IAM Identity Center User Guide](#)
- [Amazon IAM Identity Center Portal API Reference](#)

## Prerequisites

You must first enable IAM Identity Center. For details about enabling IAM Identity Center authentication, see [Getting Started](#) in the *Amazon IAM Identity Center User Guide*.

Alternatively, follow the [IAM Identity Center authentication](#) instructions in this guide. These instructions provide complete guidance, from enabling IAM Identity Center to completing the necessary shared config file configuration that follows here.

## SSO token provider configuration

### Note

To use the Amazon CLI to create this configuration for you, see [Configure your profile with the `aws configure sso` wizard](#) in the Amazon CLI.

When you use the SSO token provider configuration, your Amazon SDK or tool automatically refreshes your session up to your extended session period. For more information on session duration and maximum duration, see [Configure the session duration of the Amazon access portal and IAM Identity Center integrated applications](#) in the *Amazon IAM Identity Center User Guide*.

The `sso-session` section of the `config` file is used to group configuration variables for acquiring SSO access tokens, which can then be used to acquire Amazon credentials. For more details about formatting sections within a `config` file, see [Format of the config file](#).

You define an `sso-session` section and associate it to a profile. `sso_region` and `sso_start_url` must be set within the `sso-session` section. Typically, `sso_account_id` and `sso_role_name` must be set in the `profile` section so that the SDK can request Amazon credentials.

### Note

For a deep dive on how the SDKs and tools use and refresh credentials using this configuration, see [Understand IAM Identity Center authentication](#).

The following example configures the SDK to request IAM Identity Center credentials. It also supports automated token refresh.

```
[profile dev]
sso_session = my-sso
sso_account_id = 111122223333
sso_role_name = SampleRole

[sso-session my-sso]
sso_region = us-east-1
sso_start_url = https://my-sso-portal.awsapps.com/start
sso_registration_scopes = sso:account:access
```

You can reuse sso-session configurations across multiple profiles.

```
[profile dev]
sso_session = my-sso
sso_account_id = 111122223333
sso_role_name = SampleRole

[profile prod]
sso_session = my-sso
sso_account_id = 111122223333
sso_role_name = SampleRole2

[sso-session my-sso]
sso_region = us-east-1
sso_start_url = https://my-sso-portal.awsapps.com/start
sso_registration_scopes = sso:account:access
```

sso\_account\_id and sso\_role\_name aren't required for all scenarios of SSO token configuration. If your application only uses Amazon Web Services that support bearer authentication, then traditional Amazon credentials are not needed. Bearer authentication is an HTTP authentication scheme that uses security tokens called bearer tokens. In this scenario, sso\_account\_id and sso\_role\_name aren't required. See the individual guide for your Amazon Web Service to determine if it supports bearer token authorization.

Registration scopes are configured as part of an sso-session. Scope is a mechanism in OAuth 2.0 to limit an application's access to a user's account. An application can request one or more scopes, and the access token issued to the application is limited to the scopes granted. These scopes define the permissions requested to be authorized for the registered OIDC client and access tokens retrieved by the client. For the supported access scope options, see [Access scopes](#) in the

*Amazon IAM Identity Center User Guide*. The following example sets `sso_registration_scopes` to provide access for listing accounts and roles.

```
[sso-session my-sso]  
sso_region = us-east-1  
sso_start_url = https://my-sso-portal.awsapps.com/start  
sso_registration_scopes = sso:account:access
```

The authentication token is cached to disk under the `~/.aws/sso/cache` directory with a file name based on the session name.

## Legacy non-refreshable configuration

Automated token refresh isn't supported using the legacy non-refreshable configuration. We recommend using the [SSO token provider configuration](#) instead.

To use the legacy non-refreshable configuration, you must specify the following settings within your profile:

- `sso_start_url`
- `sso_region`
- `sso_account_id`
- `sso_role_name`

You specify the user portal for a profile with the `sso_start_url` and `sso_region` settings. You specify permissions with the `sso_account_id` and `sso_role_name` settings.

The following example sets the four required values in the config file.

```
[profile my-sso-profile]  
sso_start_url = https://my-sso-portal.awsapps.com/start  
sso_region = us-west-2  
sso_account_id = 111122223333  
sso_role_name = SSOReadOnlyRole
```

The authentication token is cached to disk under the `~/.aws/sso/cache` directory with a file name based on the `sso_start_url`.



## IAM Identity Center credential provider settings

Configure this functionality by using the following:

### `sso_start_url` - shared Amazon config file setting

The URL that points to your organization's IAM Identity Center access portal. For more information on the IAM Identity Center access portal, see [Using the Amazon access portal](#) in the *Amazon IAM Identity Center User Guide*.

To find this value, open the [IAM Identity Center console](#), view the **Dashboard**, and find **Amazon access portal URL**.

### `sso_region` - shared Amazon config file setting

The Amazon Web Services Region that contains your IAM Identity Center portal host; that is, the Region you selected before enabling IAM Identity Center. This is independent from your default Amazon Region, and can be different.

For a complete list of the Amazon Web Services Regions and their codes, see [Regional Endpoints](#) in the *Amazon Web Services General Reference*. To find this value, open the [IAM Identity Center console](#), view the **Dashboard**, and find **Region**.

### `sso_account_id` - shared Amazon config file setting

The numeric ID of the Amazon Web Services account that was added through the Amazon Organizations service to use for authentication.

To see the list of available accounts, go to the [IAM Identity Center console](#) and open the **Amazon Web Services accounts** page. You can also see the list of available accounts using the [ListAccounts](#) API method in the *Amazon IAM Identity Center Portal API Reference*. For example, you can call the Amazon CLI method [list-accounts](#).

### `sso_role_name` - shared Amazon config file setting

The name of a permission set provisioned as an IAM role that defines the user's resulting permissions. The role must exist in the Amazon Web Services account specified by `sso_account_id`. Use the role name, not the role Amazon Resource Name (ARN).

Permission sets have IAM policies and custom permissions policies attached to them and define the level of access that users have to their assigned Amazon Web Services accounts.

To see the list of available permission sets per Amazon Web Services account, go to the [IAM Identity Center console](#) and open the **Amazon Web Services accounts** page. Choose the correct permission set name listed in the Amazon Web Services accounts table. You can also see the list of available permission sets using the [ListAccountRoles](#) API method in the *Amazon IAM Identity Center Portal API Reference*. For example, you can call the Amazon CLI method [list-account-roles](#).

### **sso\_registration\_scopes** - shared Amazon config file setting

A comma-delimited list of valid scope strings to be authorized for the `sso-session`. Scopes authorize access to IAM Identity Center bearer token authorized endpoints. A minimum scope of `sso:account:access` must be granted to get a refresh token back from the IAM Identity Center service. For the supported access scope strings, see [Access scopes](#) in the *Amazon IAM Identity Center User Guide*. This setting doesn't apply to the legacy non-refreshable configuration. Tokens issued using the legacy configuration are limited to scope `sso:account:access` implicitly.

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Supported	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	Yes	
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	Yes	To use shared config file settings, you must turn on loading from the config file; see <a href="#">Sessions</a> .
<a href="#">SDK for Java 2.x</a>	Yes	Configuration values also supported in <code>credentials</code> file.
<a href="#">SDK for Java 1.x</a>	No	
<a href="#">SDK for JavaScript 3.x</a>	Yes	

SDK	Supported	Notes or more information
<a href="#">SDK for JavaScript 2.x</a>	Yes	
<a href="#">SDK for Kotlin</a>	Yes	
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	Partial	Legacy non-refreshable configuration only.
<a href="#">Tools for PowerShell</a>	Yes	

## IMDS credential provider

Instance Metadata Service (IMDS) provides data about your instance that you can use to configure or manage the running instance. For more information about the data available, see [Instance metadata and user data](#) in the *Amazon EC2 User Guide for Linux Instances* or [Instance metadata and user data](#) in the *Amazon EC2 User Guide for Windows Instances*. Amazon EC2 provides a local endpoint available to instances that can provide various bits of information to the instance. If the instance has a role attached, it can provide a set of credentials that are valid for that role. The SDKs can use that endpoint to resolve credentials as part of their [default credential provider chain](#). Instance Metadata Service Version 2 (IMDSv2), a more secure version of IMDS that uses a session token, is used by default. If that fails due to a non-retryable condition (HTTP error codes 403, 404, 405), IMDSv1 is used as a fallback.

Configure this functionality by using the following:

### **AWS\_EC2\_METADATA\_DISABLED - environment variable**

Whether or not to attempt to use Amazon EC2 Instance Metadata Service (IMDS) to obtain credentials.

**Default value:** false.

**Valid values:**

- **true** – Do not use IMDS to obtain credentials.
- **false** – Use IMDS to obtain credentials.

**ec2\_metadata\_v1\_disabled** - shared Amazon config file setting,  
**AWS\_EC2\_METADATA\_V1\_DISABLED** - environment variable, `aws.disableEc2MetadataV1` -  
JVM system property: Java/Kotlin only

Whether or not to use Instance Metadata Service Version 1 (IMDSv1) as a fallback if IMDSv2 fails.

**Note**

New SDKs don't support IMDSv1 and, thus, don't support this setting. For details, see table [Compatibility with Amazon SDKs](#).

**Default value:** false.

**Valid values:**

- **true** – Do not use IMDSv1 as a fallback.
- **false** – Use IMDSv1 as a fallback.

**ec2\_metadata\_service\_endpoint** - shared Amazon config file setting,  
**AWS\_EC2\_METADATA\_SERVICE\_ENDPOINT** - environment variable,  
`aws.ec2MetadataServiceEndpoint` - JVM system property: Java/Kotlin only

The endpoint of IMDS.

**Default value:** If `ec2_metadata_service_endpoint_mode` equals IPv4, then default endpoint is `http://169.254.169.254`. If `ec2_metadata_service_endpoint_mode` equals IPv6, then default endpoint is `http://[fd00:ec2::254]`.

**Valid values:** Valid URI.

**ec2\_metadata\_service\_endpoint\_mode** - shared Amazon config file setting,  
**AWS\_EC2\_METADATA\_SERVICE\_ENDPOINT\_MODE** - environment variable,  
`aws.ec2MetadataServiceEndpointMode` - JVM system property: Java/Kotlin only

The endpoint mode of IMDS.

**Default value:**IPv4.

**Valid values:** IPv4, IPv6.

### Note

The IMDS credential provider is a part of the [Credential provider chain](#). However, the IMDS credential provider is only checked after several other providers that are in this series. Therefore, if you want your program use this provider's credentials, you must remove other valid credential providers from your configuration or use a different profile. Alternatively, instead of relying on the credential provider chain to automatically discover which provider returns valid credentials, specify the use of the IMDS credential provider in code. You can specify credential sources directly when you create service clients.

## Security for IMDS credentials

By default, when the Amazon SDK is not configured with valid credentials the SDK will attempt to use the Amazon EC2 Instance Metadata Service (IMDS) to retrieve credentials for an Amazon role. This behavior can be disabled by setting the `AWS_EC2_METADATA_DISABLED` environment variable to `true`. This prevents unnecessary network activity and enhances security on untrusted networks where the Amazon EC2 Instance Metadata Service may be impersonated.

### Note

Amazon SDK clients configured with valid credentials will never use IMDS to retrieve credentials, regardless of any of these settings.

## Disabling use of Amazon EC2 IMDS credentials

How you set this environment variable depends on what operating system is in use as well as whether or not you want the change to be persistent.

### Linux and macOS

Customers using Linux or macOS can set this environment variable with the following command:

```
$ export AWS_EC2_METADATA_DISABLED=true
```

If you want this setting to be persistent across multiple shell sessions and system restarts, you can add the above command to your shell profile file, such as `.bash_profile`, `.zsh_profile`, or `.profile`.

## Windows

Customers using Windows can set this environment variable with the following command:

```
$ set AWS_EC2_METADATA_DISABLED=true
```

If you want this setting to be persistent across multiple shell sessions and system restarts can use the following command instead:

```
$ setx AWS_EC2_METADATA_DISABLED=true
```

### Note

The `setx` command does not apply the value to the current shell session, so you will need to reload or reopen the shell for the change to take effect.

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Support	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	Yes	
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	Yes	To use shared config file settings, you must turn on loading from the config file; see <a href="#">Sessions</a> .
<a href="#">SDK for Java 2.x</a>	Yes	

SDK	Supported	Notes or more information
<a href="#">SDK for Java 1.x</a>	Yes	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	Yes	
<a href="#">SDK for Kotlin</a>	Yes	Does not use IMDSv1 fallback.
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	Yes	Does not use IMDSv1 fallback.
<a href="#">Tools for PowerShell</a>	Yes	You can disable IMDSv1 fallback explicitly in code using <code>[Amazon.Util.EC2InstanceMetadata]::EC2MetadataV1Disabled = \$true</code> .

## Process credential provider

SDKs provide a way to extend the credential provider chain for custom use cases.

IAM Roles Anywhere provides a way to get temporary credentials for a workload or process that runs outside of Amazon. To configure `credential_process` for this use, see [IAM Roles Anywhere](#).

### Warning

The following describes a method of sourcing credentials from an external process. This can potentially be dangerous, so proceed with caution. Other credential providers should be preferred if at all possible. If using this option, you should make sure that the config file is as locked down as possible using security best practices for your operating system. Confirm that your custom credential tool does not write any secret information to `StdErr`, because

the SDKs and Amazon CLI can capture and log such information, potentially exposing it to unauthorized users.

Configure this functionality by using the following:

### **credential\_process - shared Amazon config file setting**

Specifies an external command that the SDK or tool runs on your behalf to generate or retrieve authentication credentials to use. The setting specifies the name of a program/command that the SDK will invoke. When the SDK invokes the process, it waits for the process to write JSON data to `stdout`. The custom provider must return information in a specific format. That information contains the credentials that the SDK or tool can use to authenticate you.

#### **Note**

The process credential provider is a part of the [Credential provider chain](#). However, the process credential provider is only checked after several other providers that are in this series. Therefore, if you want your program use this provider's credentials, you must remove other valid credential providers from your configuration or use a different profile. Alternatively, instead of relying on the credential provider chain to automatically discover which provider returns valid credentials, specify the use of the process credential provider in code. You can specify credential sources directly when you create service clients.

### **Specifying the path to the credentials program**

The setting's value is a string that contains a path to a program that the SDK or development tool runs on your behalf:

- The path and file name can consist of only these characters: A-Z, a-z, 0-9, hyphen ( - ), underscore ( \_ ), period ( . ), forward slash ( / ), backslash ( \ ), and space.
- If the path or file name contains a space, surround the complete path and file name with double-quotation marks ( " ").
- If a parameter name or a parameter value contains a space, surround that element with double-quotation marks ( " "). Surround only the name or value, not the pair.



- Don't include any environment variables in the strings. For example, don't include \$HOME or %USERPROFILE%.
- Don't specify the home folder as ~. \* You must specify either the full path or a base file name. If there is a base file name, the system attempts to find the program within folders specified by the PATH environment variable.

The following example shows setting `credential_process` in the shared config file on Linux/macOS.

```
credential_process = "/path/to/credentials.sh" parameterWithoutSpaces "parameter with spaces"
```

The following example shows setting `credential_process` in the shared config file on Windows.

```
credential_process = "C:\Path\To\credentials.cmd" parameterWithoutSpaces "parameter with spaces"
```

## Valid output from the credentials program

The SDK runs the command as specified in the profile and then reads data from the standard output stream. The command you specify, whether a script or binary program, must generate JSON output on STDOUT that matches the following syntax.

```
{  
  "Version": 1,  
  "AccessKeyId": "an AWS access key",  
  "SecretAccessKey": "your AWS secret access key",  
  "SessionToken": "the AWS session token for temporary credentials",  
  "Expiration": "RFC3339 timestamp for when the credentials expire"  
}
```

### Note

As of this writing, the `Version` key must be set to 1. This might increment over time as the structure evolves.

The `Expiration` key is an RFC3339 formatted timestamp. If the `Expiration` key isn't present in the tool's output, the SDK assumes that the credentials are long-term credentials that don't refresh. Otherwise, the credentials are considered temporary credentials, and they are automatically refreshed by rerunning the `credential_process` command before the credentials expire.

### Note

The SDK does **not** cache external process credentials the way it does assume-role credentials. If caching is required, you must implement it in the external process.

The external process can return a non-zero return code to indicate that an error occurred while retrieving the credentials.

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Support	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	Yes	
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	Yes	To use shared config file settings, you must turn on loading from the config file; see <a href="#">Sessions</a> .
<a href="#">SDK for Java 2.x</a>	Yes	
<a href="#">SDK for Java 1.x</a>	Yes	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	Yes	

SDK	Support	Notes or more information
<a href="#">SDK for Kotlin</a>	Yes	
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	Yes	
<a href="#">Tools for PowerShell</a>	Yes	

## Standardized features

Many features have been standardized to consistent defaults and to work the same way across many SDKs. This consistency increases productivity and clarity when coding across multiple SDKs. All settings can be overridden in code, see your specific SDK API for details.

### Important

Not all SDKs support all features, or even all aspects within a feature.

### Topics

- [Amazon EC2 instance metadata](#)
- [Amazon S3 access points](#)
- [Amazon S3 Multi-Region Access Points](#)
- [Amazon Web Services Region](#)
- [Amazon STS Regionalized endpoints](#)
- [Dual-stack and FIPS endpoints](#)
- [Endpoint discovery](#)
- [General configuration settings](#)

- [IMDS client](#)
- [Retry behavior](#)
- [Request compression](#)
- [Service-specific endpoints](#)
- [Smart configuration defaults](#)

## Amazon EC2 instance metadata

Amazon EC2 provides a service on instances called the Instance Metadata Service (IMDS). To learn more about this service, see [Instance metadata and user data](#) in the *Amazon EC2 User Guide for Linux Instances* or [Instance metadata and user data](#) in the *Amazon EC2 User Guide for Windows Instances*. When attempting to retrieve credentials on an Amazon EC2 instance that has been configured with an IAM role, the connection to the instance metadata service is adjustable.

Configure this functionality by using the following:

**metadata\_service\_num\_attempts** - shared Amazon config file setting,  
**AWS\_METADATA\_SERVICE\_NUM\_ATTEMPTS** - environment variable

This setting specifies the number of total attempts to make before giving up when attempting to retrieve data from the instance metadata service.

**Default value:** 1

**Valid values:** Number greater than or equal to 1.

**metadata\_service\_timeout** - shared Amazon config file setting,  
**AWS\_METADATA\_SERVICE\_TIMEOUT** - environment variable

Specifies the number of seconds before timing out when attempting to retrieve data from the instance metadata service.

**Default value:** 1

**Valid values:** Number greater than or equal to 1.

Example of setting these values in the config file:

```
[default]
metadata_service_num_attempts=10
```

```
metadata_service_timeout=10
```

Linux/macOS example of setting environment variables via command line:

```
export AWS_METADATA_SERVICE_NUM_ATTEMPTS=10
export AWS_METADATA_SERVICE_TIMEOUT=10
```

Windows example of setting environment variables via command line:

```
setx AWS_METADATA_SERVICE_NUM_ATTEMPTS 10
setx AWS_METADATA_SERVICE_TIMEOUT 10
```

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Support	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	No	
<a href="#">SDK for Go V2 (1.x)</a>	No	
<a href="#">SDK for Go 1.x (V1)</a>	No	
<a href="#">SDK for Java 2.x</a>	No	
<a href="#">SDK for Java 1.x</a>	Partial	metadata_service_num_attempts not supported.
<a href="#">SDK for JavaScript 3.x</a>	No	
<a href="#">SDK for JavaScript 2.x</a>	No	
<a href="#">SDK for Kotlin</a>	No	
<a href="#">SDK for .NET 3.x</a>	No	
<a href="#">SDK for PHP 3.x</a>	Yes	

SDK	Supported	Notes or more information
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	No	
<a href="#">SDK for Rust</a>	No	
<a href="#">Tools for PowerShell</a>	No	

## Amazon S3 access points

The Amazon S3 service provides access points as an alternative way to interact with Amazon S3 buckets. Access points have unique policies and configurations that can be applied to them instead of directly to the bucket. With Amazon SDKs, you can use access point Amazon Resource Names (ARNs) in the bucket field for API operations instead of specifying the bucket name explicitly. They are used for specific operations such as using an access point ARN with [GetObject](#) to fetch an object from a bucket, or using an access point ARN with [PutObject](#) to add an object to a bucket.

To learn more about Amazon S3 access points and ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Configure this functionality by using the following:

**s3\_use\_arn\_region** - shared Amazon config file setting, **AWS\_S3\_USE\_ARN\_REGION** - environment variable, **aws.s3UseArnRegion** - JVM system property: Java/Kotlin only, To configure value directly in code, consult your specific SDK directly.

This setting controls whether the SDK uses the access point ARN Amazon Web Services Region to construct the Regional endpoint for the request. The SDK validates that the ARN Amazon Web Services Region is served by the same Amazon partition as the client's configured Amazon Web Services Region to prevent cross-partition calls that most likely will fail. If multiply defined, the code-configured setting takes precedence, followed by the environment variable setting.

**Default value:** false

**Valid values:**

- **true** – The SDK uses the ARN's Amazon Web Services Region when constructing the endpoint instead of the client's configured Amazon Web Services Region. Exception: If the client's configured Amazon Web Services Region is a FIPS Amazon Web Services Region, then it must match the ARN's Amazon Web Services Region. Otherwise, an error will result.
- **false** – The SDK uses the client's configured Amazon Web Services Region when constructing the endpoint.

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Support	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	Yes	
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	Yes	To use shared config file settings, you must turn on loading from the config file; see <a href="#">Sessions</a> .
<a href="#">SDK for Java 2.x</a>	Yes	
<a href="#">SDK for Java 1.x</a>	Yes	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	Yes	
<a href="#">SDK for Kotlin</a>	Yes	
<a href="#">SDK for .NET 3.x</a>	Yes	Doesn't follow standard precedence; shared config file value takes precedence over environment variable.
<a href="#">SDK for PHP 3.x</a>	Yes	

SDK	Supported	Notes or more information
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	No	
<a href="#">Tools for PowerShell</a>	Yes	Doesn't follow standard precedence; shared config file value takes precedence over environment variable.

## Amazon S3 Multi-Region Access Points

Amazon S3 Multi-Region Access Points provide a global endpoint that applications can use to fulfill requests from Amazon S3 buckets located in multiple Amazon Web Services Regions. You can use Multi-Region Access Points to build multi-Region applications with the same architecture used in a single Region, and then run those applications anywhere in the world.

To learn more about Multi-Region Access Points, see [Multi-Region Access Points in Amazon S3](#) in the *Amazon S3 User Guide*.

To learn more about Multi-Region Access Point Amazon Resource Names (ARNs), see [Making requests using a Multi-Region Access Point](#) in the *Amazon S3 User Guide*.

To learn more about creating Multi-Region Access Points, see [Managing Multi-Region Access Points](#) in the *Amazon S3 User Guide*.

The SigV4A algorithm is the signing implementation used to sign the global Region requests. This algorithm is obtained by the SDK through a dependency on the [Amazon Common Runtime \(CRT\) libraries](#).

Configure this functionality by using the following:



**s3\_disable\_multiregion\_access\_points** - shared Amazon config file setting, **AWS\_S3\_DISABLE\_MULTIREGION\_ACCESS\_POINTS** - environment variable, **aws.s3DisableMultiRegionAccessPoints** - JVM system property: Java/Kotlin only, To configure value directly in code, consult your specific SDK directly.

This setting controls whether the SDK potentially attempts cross-Region requests. If multiply defined, the code-configured setting takes precedence, followed by the environment variable setting.

**Default value:** `false`

**Valid values:**

- **true** – Stops the use of cross-Region requests.
- **false** – Enables cross-Region requests using Multi-Region Access Points.

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Support	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	Yes	
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	No	
<a href="#">SDK for Java 2.x</a>	Yes	
<a href="#">SDK for Java 1.x</a>	No	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	No	
<a href="#">SDK for Kotlin</a>	Yes	

SDK	Supported	Notes or more information
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	Yes	
<a href="#">Tools for PowerShell</a>	Yes	

## Amazon Web Services Region

Amazon Web Services Regions are an important concept to understand when working with Amazon Web Services.

With Amazon Web Services Regions, you can access Amazon Web Services that physically reside in a specific geographic area. This can be useful to keep your data and applications running close to where you and your users will access them. Regions provide fault tolerance, stability, and resilience, and can also reduce latency. With Regions, you can create redundant resources that remain available and unaffected by a Regional outage.

Most Amazon Web Service requests are associated with a particular geographic region. The resources that you create in one Region do not exist in any other Region unless you explicitly use a replication feature offered by an Amazon Web Service. For example, Amazon S3 and Amazon EC2 support cross-Region replication. Some services, such as IAM, do not have Regional resources.

The *Amazon Web Services General Reference* contains information on the following:

- To understand the relationship between Regions and endpoints, and to view a list of existing Regional endpoints, see [Amazon service endpoints](#).
- To view the current list of all supported Regions and endpoints for each Amazon Web Service, see [Service endpoints and quotas](#).

### Creating service clients

To programmatically access Amazon Web Services, SDKs use a client class/object for each Amazon Web Service. If your application needs to access Amazon EC2, for example, your application would create an Amazon EC2 client object to interface with that service.

If no Region is explicitly specified for the client, the client defaults to using the Region set through the following `region` setting. However, the active Region for a client can be explicitly set for any individual client object. Setting the Region in this way takes precedence over any global setting for that particular service client. The alternative Region is specified during instantiation of that client, specific to your SDK (check your specific SDK Guide or your SDK's code base).

Configure this functionality by using the following:

**region** - shared Amazon config file setting, **AWS\_REGION** - environment variable, **aws.region** - JVM system property: Java/Kotlin only

Specifies the default Amazon Web Services Region to use for Amazon requests. This Region is used for SDK service requests that aren't provided with a specific Region to use.

**Default value:** None. You must specify this value explicitly.

**Valid values:**

- Any of the Region codes available for the chosen service, as listed in [Amazon service endpoints](#) in the *Amazon General Reference*. For example, the value `us-east-1` sets the endpoint to the Amazon Web Services Region US East (N. Virginia).
- `aws-global` specifies the global endpoint for services that support a separate global endpoint in addition to Regional endpoints, such as Amazon Security Token Service (Amazon STS) and Amazon Simple Storage Service (Amazon S3).

Example of setting this value in the config file:

```
[default]
region = us-west-2
```

Linux/macOS example of setting environment variables via command line:

```
export AWS_REGION=us-west-2
```

Windows example of setting environment variables via command line:

```
setx AWS_REGION us-west-2
```

Most SDKs have a "configuration" object that is available for setting the default Region from within the application code. For details, see your specific Amazon SDK developer guide.

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Support	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	Amazon CLI v2 uses any value in <code>AWS_REGION</code> before any value in <code>AWS_DEFAULT_REGION</code> (both variables are checked).
<a href="#">Amazon CLI v1</a>	Yes	Amazon CLI v1 uses environment variable named <code>AWS_DEFAULT_REGION</code> for this purpose.
<a href="#">SDK for C++</a>	Yes	
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	Yes	To use shared config file settings, you must turn on loading from the config file; see <a href="#">Sessions</a> .
<a href="#">SDK for Java 2.x</a>	Yes	
<a href="#">SDK for Java 1.x</a>	Yes	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	Yes	
<a href="#">SDK for Kotlin</a>	Yes	
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	

SDK	Support	Notes or more information
<a href="#">SDK for Python (Boto3)</a>	Yes	This SDK uses environment variable named <code>AWS_DEFAULT_REGION</code> for this purpose.
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	Yes	
<a href="#">Tools for PowerShell</a>	Yes	

## Amazon STS Regionalized endpoints

By default, Amazon Security Token Service (Amazon STS) is available as a global service, and all Amazon STS requests go to a single endpoint at `https://sts.amazonaws.com`. Global requests map to the US East (N. Virginia) Region. Amazon recommends using Regional Amazon STS endpoints instead of the global endpoint. For more information on Amazon STS endpoints, [Endpoints](#) in the *Amazon Security Token Service API Reference*.

Configure this functionality by using the following:

**`sts_regional_endpoints` - shared Amazon config file setting,**  
**`AWS_STS_REGIONAL_ENDPOINTS` - environment variable**

This setting specifies how the SDK or tool determines the Amazon Web Service endpoint that it uses to talk to the Amazon Security Token Service (Amazon STS).

**Default value:** `legacy`

### Note

All new SDK major versions releasing after July 2022 will default to `regional`. New SDK major versions might remove this setting and use `regional` behavior. To reduce future impact regarding this change, we recommend you start using `regional` in your application when possible.

**Valid values:** *(Recommended value: `regional`)*

- **legacy** – Uses the global Amazon STS endpoint, `sts.amazonaws.com`, for the following Amazon Regions: `ap-northeast-1`, `ap-south-1`, `ap-southeast-1`, `ap-southeast-2`, `aws-global`, `ca-central-1`, `eu-central-1`, `eu-north-1`, `eu-west-1`, `eu-west-2`, `eu-west-3`, `sa-east-1`, `us-east-1`, `us-east-2`, `us-west-1`, and `us-west-2`. All other Regions automatically use their respective Regional endpoint.
- **regional** – The SDK or tool always uses the Amazon STS endpoint for the currently configured Region. For example, if the client is configured to use `us-west-2`, all calls to Amazon STS are made to the Regional endpoint `sts.us-west-2.amazonaws.com`, instead of the global `sts.amazonaws.com` endpoint. To send a request to the global endpoint while this setting is enabled, you can set the Region to `aws-global`.

Example of setting these values in the config file:

```
[default]
sts_regional_endpoints = regional
```

Linux/macOS example of setting environment variables via command line:

```
export AWS_STS_REGIONAL_ENDPOINTS=regional
```

Windows example of setting environment variables via command line:

```
setx AWS_STS_REGIONAL_ENDPOINTS regional
```

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Support	Notes or more information
<a href="#">Amazon CLI v2</a>	Partial	Default value is <code>regional</code> .
<a href="#">SDK for C++</a>	Partial	Environment variable and config file setting not supported. SDK performs with <code>regional</code> setting.

SDK	Supported	Notes or more information
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	Yes	To use shared config file settings, you must turn on loading from the config file; see <a href="#">Sessions</a> .
<a href="#">SDK for Java 2.x</a>	Yes	
<a href="#">SDK for Java 1.x</a>	Yes	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	Yes	
<a href="#">SDK for Kotlin</a>	No	
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	Yes	
<a href="#">Tools for PowerShell</a>	Yes	

## Dual-stack and FIPS endpoints

Configure this functionality by using the following:

**use\_dualstack\_endpoint** - shared Amazon config file setting,  
**AWS\_USE\_DUALSTACK\_ENDPOINT** - environment variable, **aws.useDualstackEndpoint** - JVM  
 system property: Java/Kotlin only

Turns on or off whether the SDK will send requests to dual-stack endpoints. To learn more about dual-stack endpoints, which support both IPv4 and IPv6 traffic, see [Using Amazon S3](#)

[dual-stack endpoints](#) in the *Amazon Simple Storage Service User Guide*. Dual-stack endpoints are available for some services in some regions.

**Default value:** false

**Valid values:**

- **true** – The SDK or tool will attempt to use dual-stack endpoints to make network requests. If a dual-stack endpoint does not exist for the service and/or Amazon Web Services Region, the request will fail.
- **false** – The SDK or tool will not use dual-stack endpoints to make network requests.

**use\_fips\_endpoint** - shared Amazon config file setting, **AWS\_USE\_FIPS\_ENDPOINT** - environment variable, **aws.useFipsEndpoint** - JVM system property: Java/Kotlin only

Turns on or off whether the SDK or tool will send requests to FIPS-compliant endpoints. The Federal Information Processing Standards (FIPS) are a set of US Government security requirements for data and its encryption. Government agencies, partners, and those wanting to do business with the federal government are required to adhere to FIPS guidelines. Unlike standard Amazon endpoints, FIPS endpoints use a TLS software library that complies with FIPS 140-2. If this setting is enabled and a FIPS endpoint does not exist for the service in your Amazon Web Services Region, the Amazon call may fail. [Service-specific endpoints](#) and the `--endpoint-url` option for the Amazon Command Line Interface override this setting.

To learn more about other ways to specify FIPS endpoints by Amazon Web Services Region, see [FIPS Endpoints by Service](#). For more information on Amazon Elastic Compute Cloud service endpoints, see [Dual-stack \(IPv4 and IPv6\) endpoints](#) in the *Amazon EC2 API Reference*.

**Default value:** false

**Valid values:**

- **true** – The SDK or tool will send requests to FIPS-compliant endpoints.
- **false** – The SDK or tool will not send requests to FIPS-compliant endpoints.

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.



SDK	Supported	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	Yes	
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	Yes	To use shared config file settings, you must turn on loading from the config file; see <a href="#">Sessions</a> .
<a href="#">SDK for Java 2.x</a>	Yes	
<a href="#">SDK for Java 1.x</a>	No	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	Yes	
<a href="#">SDK for Kotlin</a>	Yes	
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	Yes	
<a href="#">Tools for PowerShell</a>	Yes	

## Endpoint discovery

SDKs use endpoint discovery to access service endpoints (URLs to access various resources), while still maintaining flexibility for Amazon to alter URLs as needed. This way, your code can automatically detect new endpoints. There are no fixed endpoints for some services. Instead, you get the available endpoints during runtime by making a request to get the endpoints first. After retrieving the available endpoints, the code then uses the endpoint to access other operations.

For example, for Amazon Timestream, the SDK makes a `DescribeEndpoints` request to retrieve the available endpoints, and then uses those endpoints to complete specific operations such as `CreateDatabase` or `CreateTable`.

Endpoint discovery is required in some services and optional in others. It defaults to either `true` or `false` depending on whether the service requires endpoint discovery. For example, Timestream defaults to `true`, and Amazon DynamoDB defaults to `false`. For services where endpoint discovery is not required, endpoint discovery is not enabled. Instead, configuration options are available through environment variables, the shared Amazon config file, or SDK code constructs (for example, configuration classes). For operations where endpoint discovery is required, the SDK automatically attempts to discover an endpoint.

Configure this functionality by using the following:

**endpoint\_discovery\_enabled** - shared Amazon config file setting, **AWS\_ENABLE\_ENDPOINT\_DISCOVERY** - environment variable, **aws.endpointDiscoveryEnabled** - JVM system property: Java/Kotlin only, To configure value directly in code, consult your specific SDK directly.

Enables/disables endpoint discovery for services where endpoint discovery is optional. Endpoint discovery is required in some services.

**Default value:** `false`

**Valid values:**

- **true** – The SDK should automatically attempt to discover an endpoint for services where endpoint discovery is optional.
- **false** – The SDK should not automatically attempt to discover an endpoint for services where endpoint discovery is optional.

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Supported	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	Yes	
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	Yes	To use shared config file settings, you must turn on loading from the config file; see <a href="#">Sessions</a> .
<a href="#">SDK for Java 2.x</a>	Yes	The SDK for Java 2.x uses <code>AWS_ENDPOINT_DISCOVERY_ENABLED</code> for the environment variable name.
<a href="#">SDK for Java 1.x</a>	Yes	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	Yes	
<a href="#">SDK for Kotlin</a>	Yes	
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	Partial	Supported for Timestream only.
<a href="#">Tools for PowerShell</a>	Yes	

## General configuration settings

SDKs support some general settings that configure overall SDK behaviors.

Configure this functionality by using the following:

## api\_versions - shared Amazon config file setting

Some Amazon services maintain multiple API versions to support backward compatibility. By default, SDK and Amazon CLI operations use the latest available API version. To require a specific API version to use for your requests, include the `api_versions` setting in your profile.

**Default value:** None. (Latest API version is used by the SDK.)

**Valid values:** This is a nested setting that's followed by one or more indented lines that each identify one Amazon service and the API version to use. See the documentation for the Amazon service to understand which API versions are available.

The example sets a specific API version for two Amazon services in the config file. These API versions are used only for commands that run under the profile that contains these settings. Commands for any other service use the latest version of that service's API.

```
api_versions =  
  ec2 = 2015-03-01  
  cloudfront = 2015-09-017
```

## ca\_bundle - shared Amazon config file setting, AWS\_CA\_BUNDLE - environment variable

Specifies the path to a custom certificate bundle (a file with a `.pem` extension) to use when establishing SSL/TLS connections.

**Default value:** none

**Valid values:** Specify either the full path or a base file name. If there is a base file name, the system attempts to find the program within folders specified by the `PATH` environment variable.

Example of setting this value in the config file:

```
[default]  
ca_bundle = dev/apps/ca-certs/cabundle-2019mar05.pem
```

Linux/macOS example of setting environment variables via command line:

```
export AWS_CA_BUNDLE=/dev/apps/ca-certs/cabundle-2019mar05.pem
```

Windows example of setting environment variables via command line:

```
setx AWS_CA_BUNDLE C:\dev\apps\ca-certs\cabundle-2019mar05.pem
```

## output - shared Amazon config file setting

Specifies how results are formatted in the Amazon CLI and other Amazon SDKs and tools.

**Default value:** json

### Valid values:

- **json** – The output is formatted as a [JSON](#) string.
- **yaml** – The output is formatted as a [YAML](#) string.
- **yaml-stream** – The output is streamed and formatted as a [YAML](#) string. Streaming allows for faster handling of large data types.
- **text** – The output is formatted as multiple lines of tab-separated string values. This can be useful to pass the output to a text processor, like `grep`, `sed`, or `awk`.
- **table** – The output is formatted as a table using the characters `+|-` to form the cell borders. It typically presents the information in a "human-friendly" format that is much easier to read than the others, but not as programmatically useful.

## parameter\_validation - shared Amazon config file setting

Specifies whether the SDK or tool attempts to validate command line parameters before sending them to the Amazon service endpoint.

**Default value:** true

### Valid values:

- **true** – The default. The SDK or tool performs client-side validation of command line parameters. This helps the SDK or tool confirm that parameters are valid, and catches some errors. The SDK or tool can reject requests that aren't valid before sending requests to the Amazon service endpoint.
- **false** – The SDK or tool doesn't validate command line parameters before sending them to the Amazon service endpoint. The Amazon service endpoint is responsible for validating all requests and rejecting requests that aren't valid.

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Support	Notes or more information
<a href="#">Amazon CLI v2</a>	Partial	<code>api_versions</code> not supported.
<a href="#">SDK for C++</a>	Yes	
<a href="#">SDK for Go V2 (1.x)</a>	Partial	<code>api_versions</code> and <code>parameter_validation</code> not supported.
<a href="#">SDK for Go 1.x (V1)</a>	Partial	<code>api_versions</code> and <code>parameter_validation</code> not supported. To use shared config file settings, you must turn on loading from the config file; see <a href="#">Sessions</a> .
<a href="#">SDK for Java 2.x</a>	No	
<a href="#">SDK for Java 1.x</a>	No	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	Yes	
<a href="#">SDK for Kotlin</a>	No	
<a href="#">SDK for .NET 3.x</a>	No	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	No	
<a href="#">Tools for PowerShell</a>	No	

## IMDS client

SDKs implement an Instance Metadata Service Version 2 (IMDSv2) client using session-oriented requests. For more information on IMDSv2, see [Use IMDSv2](#) in the *Amazon EC2 User Guide for Linux Instances* or [Use IMDSv2](#) in the *Amazon EC2 User Guide for Windows Instances*. The IMDS client is configurable via a client configuration object available in the SDK code base.

Configure this functionality by using the following:

### **retries** - client configuration object member

The number of additional retry attempts for any failed request.

**Default value:** 3

**Valid values:** Number greater than 0.

### **port** - client configuration object member

The port for the endpoint.

**Default value:** 80

**Valid values:** Number.

### **token\_ttl** - client configuration object member

The TTL of the token.

**Default value:** 21,600 seconds (6 hours, the maximum time allotted).

**Valid values:** Number.

### **endpoint** - client configuration object member

The endpoint of IMDS.

**Default value:** If `endpoint_mode` equals IPv4, then default endpoint is `http://169.254.169.254`. If `endpoint_mode` equals IPv6, then default endpoint is `http://[fd00:ec2::254]`.

**Valid values:** Valid URI.

The following options are supported by most SDKs. See your specific SDK code base for details.

**endpoint\_mode - client configuration object member**

The endpoint mode of IMDS.

**Default value:** IPv4

**Valid values:** IPv4, IPv6

**http\_open\_timeout - client configuration object member (name may vary)**

The number of seconds to wait for the connection to open.

**Default value:** 1 second.

**Valid values:** Number greater than 0.

**http\_read\_timeout - client configuration object member (name may vary)**

The number of seconds for one chunk of data to be read.

**Default value:** 1 second.

**Valid values:** Number greater than 0.

**http\_debug\_output - client configuration object member (name may vary)**

Sets an output stream for debugging.

**Default value:** None.

**Valid values:** A valid I/O stream, like STDOUT.

**backoff - client configuration object member (name may vary)**

The number of seconds to sleep in between retries or a customer provided backoff function to call. This overrides the default exponential backoff strategy.

**Default value:** Varies by SDK.

**Valid values:** Varies by SDK. Can be either a numeric value or a call out to a custom function.

**Compatibility with Amazon SDKs**

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.



SDK	Supported	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	No	IMDSv2 used internally only. See <a href="#">IMDS credential provider</a> .
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	Yes	
<a href="#">SDK for Java 2.x</a>	Yes	
<a href="#">SDK for Java 1.x</a>	Yes	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	Yes	
<a href="#">SDK for Kotlin</a>	Yes	
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	Yes	
<a href="#">Tools for PowerShell</a>	Yes	

## Retry behavior

Retry behavior includes settings regarding how the SDKs attempt to recover from failures resulting from requests made to Amazon Web Services.

Configure this functionality by using the following:

**max\_attempts** - shared Amazon config file setting, **AWS\_MAX\_ATTEMPTS** - environment variable, **aws.maxAttempts** - JVM system property: Java/Kotlin only

Specifies the maximum number attempts to make on a request.

**Default value:** If this value is not specified, its default depends on the value of the `retry_mode` setting:

- If `retry_mode` is `legacy` – Uses a default value specific to your SDK (check your specific SDK guide or your SDK's code base for `max_attempts` default).
- If `retry_mode` is `standard` – Makes three attempts.
- If `retry_mode` is `adaptive` – Makes three attempts.

**Valid values:** Number greater than 0.

**retry\_mode** - shared Amazon config file setting, **AWS\_RETRY\_MODE** - environment variable, **aws.retryMode** - JVM system property: Java/Kotlin only

Specifies how the SDK or developer tool attempts retries.

**Default value:** `legacy` is the default retry strategy.

**Valid values:**

- `legacy` – Specific to your SDK (check your specific SDK guide or your SDK's code base).
- `standard` – The standard set of retry rules across Amazon SDKs. This mode includes a standard set of errors that are retried, and support for retry quotas. The default maximum number of attempts with this mode is three, unless `max_attempts` is explicitly configured.
- `adaptive` – An experimental retry mode that includes the functionality of standard mode but includes automatic client-side throttling. Because this mode is experimental, it might change behavior in the future.

Following is the high-level pseudocode for both the `standard` and `adaptive` retry modes:

```
MakeSDKRequest() {
  attempts = 0
  loop {
    GetSendToken()
    response = SendHTTPRequest()
    RequestBookkeeping(response)
```

```
    if not Retryable(response)
        return response
    attempts += 1
    if attempts >= MAX_ATTEMPTS:
        return response
    if not HasRetryQuota(response)
        return response
    delay = ExponentialBackoff(attempts)
    sleep(delay)
}
```

Following are more details about the components used in the pseudocode:

### **GetSendToken:**

Token buckets are only used in adaptive retry mode. Token buckets enforce a maximum request rate by requiring a token to be available in order to initiate a request. The SDK client is configurable to either fast fail the request or block until a token becomes available.

*Client Side Rate Limiting* is an algorithm that initially lets requests be made at any rate up to the token allowance. However, after a throttled response is detected, the client rate-of-request is then limited accordingly. The token allowance is also increased accordingly if successful responses are received.

With adaptive rate limiting, SDKs can slow down the rate at which requests are sent in order to better accommodate the capacity of Amazon Web Services.

### **SendHTTPRequest:**

Most Amazon SDKs use an HTTP library that uses connection pools so that you can reuse an existing connection when making an HTTP request. Generally, connections are reused when retrying requests due to throttling errors. Requests are not reused when retrying due to transient errors.

### **RequestBookkeeping:**

The retry quota should be updated if the request is successful. For adaptive retry mode only, the state variable `maxsendrate` is updated based on the type of response received.

### **Retryable:**

This step determines whether a response can be retried based on the following:

- The HTTP status code.
- The error code returned from the service.
- Connection errors, defined as any error received by the SDK in which an HTTP response from the service is not received.

Transient errors (HTTP status codes 400, 408, 500, 502, 503, and 504) and throttling errors (HTTP status codes 400, 403, 429, 502, 503, and 509) can all potentially be retried. SDK retry behavior is determined in combination with error codes or other data from the service.

### **MAX\_ATTEMPTS:**

Specified by the config file setting or the environment variable.

### **HasRetryQuota**

This step throttles retry requests by requiring a token to be available in the retry quota bucket. Retry quota buckets are a mechanism to prevent retries that are unlikely to succeed. These quotas are SDK-dependent, are often client-dependent, and are sometimes even dependent on service endpoints. The available retry quota tokens are removed when requests fail for various reasons, and replenished when they succeed. When no tokens remain, the retry loop is exited.

### **ExponentialBackoff**

For an error that can be retried, the retry delay is calculated using truncated exponential backoff. The SDKs use truncated binary exponential backoff with jitter. The following algorithm shows how the amount of time to sleep, in seconds, is defined for a response for request  $i$ :

$$\text{seconds\_to\_sleep\_i} = \min(b \cdot r^i, \text{MAX\_BACKOFF})$$

In the preceding algorithm, the following values apply:

$b$  = random number within the range of:  $0 \leq b \leq 1$

$r$  = 2

$\text{MAX\_BACKOFF}$  = 20 seconds for most SDKs. See your specific SDK guide or source code for confirmation.

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Support	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	Yes	
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	No	
<a href="#">SDK for Java 2.x</a>	Yes	
<a href="#">SDK for Java 1.x</a>	Yes	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	No	Supports a maximum number of retries, exponential backoff with jitter, and an option for a custom method for retry backoff.
<a href="#">SDK for Kotlin</a>	Yes	
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	Yes	
<a href="#">Tools for PowerShell</a>	Yes	

## Request compression

Amazon SDKs and tools can automatically compress payloads when sending requests to Amazon Web Services that support receiving compressed payloads. Compressing the payload on the client prior to sending it to a service may reduce the overall number of requests and bandwidth required to send data to the service, as well as reduce unsuccessful requests due to service limitations on the payload size. For compression, the SDK or tool selects an encoding algorithm that is supported by both the service and the SDK. However, the current list of possible encodings consists only of gzip, but it may expand in the future.

Request compression can be especially useful if your application is using [Amazon CloudWatch](#). CloudWatch is a monitoring and observability service that collects monitoring and operational data in the form of logs, metrics, and events. One example of a service operation that supports compression is CloudWatch's [PutMetricDataAPI](#) method.

Configure this functionality by using the following:

**disable\_request\_compression** - shared Amazon config file setting, **AWS\_DISABLE\_REQUEST\_COMPRESSION** - environment variable, **aws.disableRequestCompression** - JVM system property: Java/Kotlin only

Turns on or off whether the SDK or tool will compress a payload prior to sending a request.

**Default value:** false

**Valid values:**

- **true** – Turn off request compression.
- **false** – Use request compression when possible.

**request\_min\_compression\_size\_bytes** - shared Amazon config file setting, **AWS\_REQUEST\_MIN\_COMPRESSION\_SIZE\_BYTES** - environment variable, **aws.requestMinCompressionSizeBytes** - JVM system property: Java/Kotlin only

Sets the minimum size in bytes of the request body that the SDK or tool should compress. Small payloads may become longer when compressed, thus, there is a lower limit where it makes sense to perform compression. This value is inclusive, a request size greater than or equal to the value is compressed.

**Default value:** 10240 bytes

**Valid values:** Integer value between 0 and 10485760 bytes inclusive.

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Support	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	Yes	
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	No	
<a href="#">SDK for Java 2.x</a>	Yes	
<a href="#">SDK for Java 1.x</a>	No	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	No	
<a href="#">SDK for Kotlin</a>	Yes	
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	No	
<a href="#">Tools for PowerShell</a>	Yes	

## Service-specific endpoints

Service-specific endpoint configuration provides the option to use an endpoint of your choosing for API requests and to have that choice persist. These settings provide flexibility to support local endpoints, VPC endpoints, and third-party local Amazon development environments. Different endpoints can be used for testing and production environments. You can specify an endpoint URL for individual Amazon Web Services.

Configure this functionality by using the following:

**endpoint\_url** - shared Amazon config file setting, **AWS\_ENDPOINT\_URL** - environment variable, **aws.endpointUrl** - JVM system property: Java/Kotlin only

When specified directly within a profile or as an environment variable, this setting specifies the endpoint that is used for all service requests. This endpoint is overridden by any configured service-specific endpoint.

You can also use this setting within a `services` section of a shared Amazon config file to set a custom endpoint for a specific service. For a list of all service identifier keys to use for subsections within the `services` section, see [Identifiers for service-specific endpoints](#).

**Default value:** none

**Valid values:** A URL including the scheme and host for the endpoint. The URL can optionally contain a path component that contains one or more path segments.

**AWS\_ENDPOINT\_URL\_<SERVICE>** - environment variable, **aws.endpointUrl<ServiceName>** - JVM system property: Java/Kotlin only

**AWS\_ENDPOINT\_URL\_<SERVICE>**, where `<SERVICE>` is the Amazon Web Service identifier, sets a custom endpoint for a specific service. For a list of all service-specific environment variables, see [Identifiers for service-specific endpoints](#).

This service-specific endpoint overrides any global endpoint set in `AWS_ENDPOINT_URL`.

**Default value:** none

**Valid values:** A URL including the scheme and host for the endpoint. The URL can optionally contain a path component that contains one or more path segments.



**ignore\_configured\_endpoint\_urls** - shared Amazon config file setting, **AWS\_IGNORE\_CONFIGURED\_ENDPOINT\_URLS** - environment variable, **aws.ignoreConfiguredEndpointUrls** - JVM system property: Java/Kotlin only

This setting is used to ignore all custom endpoints configurations.

Note that any explicit endpoint set in the code or on a service client itself is used regardless of this setting. For example, including the `--endpoint-url` command line parameter with an Amazon CLI command or passing an endpoint URL into a client constructor will always take effect.

**Default value:** `false`

**Valid values:**

- **true** – The SDK or tool does not read any custom configuration options from the shared config file or from environment variables for setting an endpoint URL.
- **false** – The SDK or tool uses any available user-provided endpoints from the shared config file or from environment variables.

## Configure endpoints using environment variables

To route requests for all services to a custom endpoint URL, set the `AWS_ENDPOINT_URL` global environment variable.

```
export AWS_ENDPOINT_URL=http://localhost:4567
```

To route requests for a specific Amazon Web Service to a custom endpoint URL, use the `AWS_ENDPOINT_URL_<SERVICE>` environment variable. Amazon DynamoDB has a `serviceId` of [DynamoDB](#). For this service, the endpoint URL environment variable is `AWS_ENDPOINT_URL_DYNAMODB`. This endpoint takes precedence over the global endpoint set in `AWS_ENDPOINT_URL` for this service.

```
export AWS_ENDPOINT_URL_DYNAMODB=http://localhost:5678
```

As another example, Amazon Elastic Beanstalk has a `serviceId` of [Elastic Beanstalk](#). The Amazon Web Service identifier is based on the API model's `serviceId` by replacing all spaces with underscores and uppercasing all letters. To set the endpoint for this service, the corresponding

environment variable is `AWS_ENDPOINT_URL_ELASTIC_BEANSTALK`. For a list of all service-specific environment variables, see [Identifiers for service-specific endpoints](#).

```
export AWS_ENDPOINT_URL_ELASTIC_BEANSTALK=http://localhost:5567
```

## Configure endpoints using the shared config file

In the shared config file, `endpoint_url` is used in different places for different functionality.

- `endpoint_url` specified directly within a profile makes that endpoint the global endpoint.
- `endpoint_url` nested under a service identifier key within a `services` section makes that endpoint apply to requests made only to that service. For details on defining a `services` section in your shared config file, see [Format of the config file](#).

The following example uses a `services` definition to configure a service-specific endpoint URL to be used for Amazon S3 and a custom global endpoint to be used for all other services:

```
[profile dev-s3-specific-and-global]  
endpoint_url = http://localhost:1234  
services = s3-specific  
  
[services s3-specific]  
s3 =  
  endpoint_url = https://play.min.io:9000
```

A single profile can configure endpoints for multiple services. This example shows how to set the service-specific endpoint URLs for Amazon S3 and Amazon Elastic Beanstalk in the same profile. Amazon Elastic Beanstalk has a `serviceId` of [Elastic Beanstalk](#). The Amazon Web Service identifier is based on the API model's `serviceId` by replacing all spaces with underscores and lowercasing all letters. Thus, the service identifier key becomes `elastic_beanstalk` and settings for this service begin on the line `elastic_beanstalk =`. For a list of all service identifier keys to use in the `services` section, see [Identifiers for service-specific endpoints](#).

```
[services testing-s3-and-eb]  
s3 =  
  endpoint_url = http://localhost:4567  
elastic_beanstalk =  
  endpoint_url = http://localhost:8000
```

```
[profile dev]
services = testing-s3-and-eb
```

The service configuration section can be used from multiple profiles. For example, two profiles can use the same services definition while altering other profile properties:

```
[services testing-s3]
s3 =
  endpoint_url = https://localhost:4567

[profile testing-json]
output = json
services = testing-s3

[profile testing-text]
output = text
services = testing-s3
```

## Configure endpoints in profiles using role-based credentials

If your profile has role-based credentials configured through a `source_profile` parameter for IAM assume role functionality, the SDK only uses service configurations for the specified profile. It does not use profiles that are role chained to it. For example, using the following shared config file:

```
[profile A]
credential_source = Ec2InstanceMetadata
endpoint_url = https://profile-a-endpoint.aws/

[profile B]
source_profile = A
role_arn = arn:aws:iam::123456789012:role/roleB
services = profileB

[services profileB]
ec2 =
  endpoint_url = https://profile-b-ec2-endpoint.aws
```

If you use profile B and make a call in your code to Amazon EC2, the endpoint resolves as `https://profile-b-ec2-endpoint.aws`. If your code makes a request to any other service,

the endpoint resolution will not follow any custom logic. The endpoint does not resolve to the global endpoint defined in profile A. For a global endpoint to take effect for profile B, you would need to set `endpoint_url` directly within profile B. For more information on the `source_profile` setting, see [Assume role credential provider](#).

## Precedence of settings

The settings for this feature can be used at the same time but only one value will take priority per service. For API calls made to a given Amazon Web Service, the following order is used to select a value:

1. Any explicit setting set in the code or on a service client itself takes precedence over anything else.
  - For the Amazon CLI, this is the value provided by the `--endpoint-url` command line parameter. For an SDK, explicit assignments can take the form of a parameter that you set when you instantiate an Amazon Web Service client or configuration object.
2. The value provided by a service-specific environment variable such as `AWS_ENDPOINT_URL_DYNAMODB`.
3. The value provided by the `AWS_ENDPOINT_URL` global endpoint environment variable.
4. The value provided by the `endpoint_url` setting nested under a service identifier key within a `services` section of the shared config file.
5. The value provided by the `endpoint_url` setting specified directly within a profile of the shared config file.
6. Any default endpoint URL for the respective Amazon Web Service is used last.

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Support	Notes or more information
<a href="#">Amazon CLI v2</a>	Yes	
<a href="#">SDK for C++</a>	No	

SDK	Support	Notes or more information
<a href="#">SDK for Go V2 (1.x)</a>	Yes	
<a href="#">SDK for Go 1.x (V1)</a>	No	
<a href="#">SDK for Java 2.x</a>	No	
<a href="#">SDK for Java 1.x</a>	No	
<a href="#">SDK for JavaScript 3.x</a>	Yes	
<a href="#">SDK for JavaScript 2.x</a>	No	
<a href="#">SDK for Kotlin</a>	Yes	
<a href="#">SDK for .NET 3.x</a>	Yes	
<a href="#">SDK for PHP 3.x</a>	Yes	
<a href="#">SDK for Python (Boto3)</a>	Yes	
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	No	
<a href="#">Tools for PowerShell</a>	Yes	

## Identifiers for service-specific endpoints

For information on how and where to use the identifiers in the following table, see [Service-specific endpoints](#).

<b>serviceId</b>	<b>Service ID</b>	<b>Environment Variable</b>
AccessAnalyzer	access-analyzer	AWS_ENDPOINT_URL_ACCESSANALYZER
Account	account	AWS_ENDPOINT_URL_ACCOUNT
ACM	acm	AWS_ENDPOINT_URL_ACM
ACM PCA	acm-pca	AWS_ENDPOINT_URL_ACM_PCA
Alexa For Business	alexa-for-business	AWS_ENDPOINT_URL_ALEXA_FOR_BUSINESS
amp	amp	AWS_ENDPOINT_URL_AMP
Amplify	amplify	AWS_ENDPOINT_URL_AMPLIFY
AmplifyBackend	amplify-backend	AWS_ENDPOINT_URL_AMPLIFYBACKEND
AmplifyUIBuilder	amplify-ui-builder	AWS_ENDPOINT_URL_AMPLIFYUIBUILDER
API Gateway	api-gateway	AWS_ENDPOINT_URL_API_GATEWAY
ApiGatewayManagementApi	api-gateway-management-api	AWS_ENDPOINT_URL_APIGATEWAYMANAGEMENTAPI

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
ApiGatewayV2	apigatewayv2	AWS_ENDPOINT_URL_APIGATEWAYV2
AppConfig	appprofile	AWS_ENDPOINT_URL_APPCONFIG
AppConfigData	appprofiledata	AWS_ENDPOINT_URL_APPCONFIGDATA
AppFabric	appprofilefabric	AWS_ENDPOINT_URL_APPFABRIC
Appflow	appprofileflow	AWS_ENDPOINT_URL_APPFLOW
AppIntegrations	appprofileintegrations	AWS_ENDPOINT_URL_APPINTEGRATIONS
Application Auto Scaling	applicationautoscaling	AWS_ENDPOINT_URL_APPLICATION_AUTO_SCALING
Application Insights	applicationinsights	AWS_ENDPOINT_URL_APPLICATION_INSIGHTS
ApplicationCostProfiler	applicationcostprofiler	AWS_ENDPOINT_URL_APPLICATIONCOSTPROFILER

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
App Mesh	aws- <code>aws-endpoint-url-app-mesh</code>	
AppRunner	aws- <code>aws-endpoint-url-apprunner</code>	
AppStream	aws- <code>aws-endpoint-url-appstream</code>	
AppSync	aws- <code>aws-endpoint-url-appsync</code>	
ARC Zonal Shift	aws- <code>aws-endpoint-url-arc-zonal-shift</code>	
Artifact	aws- <code>aws-endpoint-url-artifact</code>	
Athena	aws- <code>aws-endpoint-url-athena</code>	
AuditManager	aws- <code>aws-endpoint-url-auditmanager</code>	
Auto Scaling	aws- <code>aws-endpoint-url-auto-scaling</code>	
Auto Scaling Plans	aws- <code>aws-endpoint-url-auto-scaling-plans</code>	
b2bi	aws- <code>aws-endpoint-url-b2bi</code>	



<b>serviceId</b>	<b>Service-specific endpoint</b>	<b>environment variable</b>
Backup	b: AWS_ENDPOINT_URL_BACKUP	
Backup Gateway	b: AWS_ENDPOINT_URL_BACKUP_GATEWAY	
BackupStorage	b: AWS_ENDPOINT_URL_BACKUPSTORAGE	
Batch	b: AWS_ENDPOINT_URL_BATCH	
BCM Data Exports	b: AWS_ENDPOINT_URL_BCM_DATA_EXPORTS	
Bedrock	b: AWS_ENDPOINT_URL_BEDROCK	
Bedrock Agent	b: AWS_ENDPOINT_URL_BEDROCK_AGENT	
Bedrock Agent Runtime	b: AWS_ENDPOINT_URL_BEDROCK_AGENT_RUNTIME	
Bedrock Runtime	b: AWS_ENDPOINT_URL_BEDROCK_RUNTIME	
billingconductor	b: AWS_ENDPOINT_URL_BILLINGCONDUCTOR	
Braket	b: AWS_ENDPOINT_URL_BRAKET	

<b>serviceId</b>	<b>Service ID</b>	<b>environment variable</b>
Budgets	budgets	AWS_ENDPOINT_URL_BUDGETS
Cost Explorer	cost-explorer	AWS_ENDPOINT_URL_COST_EXPLORER
chatbot	chatbot	AWS_ENDPOINT_URL_CHATBOT
Chime	chime	AWS_ENDPOINT_URL_CHIME
Chime SDK Identity	chime-sdk-identity	AWS_ENDPOINT_URL_CHIME_SDK_IDENTITY
Chime SDK Media Pipelines	chime-sdk-media-pipelines	AWS_ENDPOINT_URL_CHIME_SDK_MEDIA_PIPELINES
Chime SDK Meetings	chime-sdk-meetings	AWS_ENDPOINT_URL_CHIME_SDK_MEETINGS
Chime SDK Messaging	chime-sdk-messaging	AWS_ENDPOINT_URL_CHIME_SDK_MESSAGING
Chime SDK Voice	chime-sdk-voice	AWS_ENDPOINT_URL_CHIME_SDK_VOICE

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
CleanRooms	c: AWS_ENDPOINT_URL_CLEANROOMS	
CleanRoomsML	c: AWS_ENDPOINT_URL_CLEANROOMSML	
Cloud9	c: AWS_ENDPOINT_URL_CLOUD9	
CloudControl	c: AWS_ENDPOINT_URL_CLOUDCONTROL	
CloudDirectory	c: AWS_ENDPOINT_URL_CLOUDDIRECTORY	
CloudFormation	c: AWS_ENDPOINT_URL_CLOUDFORMATION	
CloudFront	c: AWS_ENDPOINT_URL_CLOUDFRONT	
CloudFront KeyValueCollection	c: AWS_ENDPOINT_URL_CLOUDFRONT_KEYVALUESTORE	
CloudHSM	c: AWS_ENDPOINT_URL_CLOUDHSM	
CloudHSM V2	c: AWS_ENDPOINT_URL_CLOUDHSM_V2	

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
CloudSearch	c: AWS_ENDPOINT_URL_CLOUDSEARCH	
CloudSearch Domain	c: AWS_ENDPOINT_URL_CLOUDSEARCH_DOMAIN	
CloudTrail	c: AWS_ENDPOINT_URL_CLOUDTRAIL	
CloudTrail Data	c: AWS_ENDPOINT_URL_CLOUDTRAIL_DATA	
CloudWatch	c: AWS_ENDPOINT_URL_CLOUDWATCH	
codeartifact	c: AWS_ENDPOINT_URL_CODEARTIFACT	
CodeBuild	c: AWS_ENDPOINT_URL_CODEBUILD	
CodeCatalyst	c: AWS_ENDPOINT_URL_CODECATALYST	
CodeCommit	c: AWS_ENDPOINT_URL_CODECOMMIT	

<b>serviceId</b>	<b>Service endpoint URL</b>	<b>environment variable</b>
CodeDeploy	<code>AWS_ENDPOINT_URL_CODEDEPLOY</code>	
CodeGuru Reviewer	<code>AWS_ENDPOINT_URL_CODEGURU_REVIEWER</code>	
CodeGuru Security	<code>AWS_ENDPOINT_URL_CODEGURU_SECURITY</code>	
CodeGuruProfiler	<code>AWS_ENDPOINT_URL_CODEGURUPROFILER</code>	
CodePipeline	<code>AWS_ENDPOINT_URL_CODEPIPELINE</code>	
CodeStar	<code>AWS_ENDPOINT_URL_CODESTAR</code>	
CodeStar connections	<code>AWS_ENDPOINT_URL_CODESTAR_CONNECTIONS</code>	
codestar notifications	<code>AWS_ENDPOINT_URL_CODESTAR_NOTIFICATIONS</code>	
Cognito Identity	<code>AWS_ENDPOINT_URL_COGNITO_IDENTITY</code>	

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
Cognito Identity Provider	<code>AWS_ENDPOINT_URL_COGNITO_IDENTITY_PROVIDER</code>	
Cognito Sync	<code>AWS_ENDPOINT_URL_COGNITO_SYNC</code>	
Comprehend	<code>AWS_ENDPOINT_URL_COMPREHEND</code>	
ComprehendMedical	<code>AWS_ENDPOINT_URL_COMPREHENDMEDICAL</code>	
Compute Optimizer	<code>AWS_ENDPOINT_URL_COMPUTE_OPTIMIZER</code>	
Config Service	<code>AWS_ENDPOINT_URL_CONFIG_SERVICE</code>	
Connect	<code>AWS_ENDPOINT_URL_CONNECT</code>	
Connect Contact Lens	<code>AWS_ENDPOINT_URL_CONNECT_CONTACT_LENS</code>	
ConnectCampaigns	<code>AWS_ENDPOINT_URL_CONNECTCAMPAIGNS</code>	

<b>serviceId</b>	<b>Service ID</b>	<b>Environment Variable</b>
ConnectCases	connectcases	AWS_ENDPOINT_URL_CONNECTCASES
ConnectParticipant	connectparticipant	AWS_ENDPOINT_URL_CONNECTPARTICIPANT
ControlTower	controltower	AWS_ENDPOINT_URL_CONTROLTOWER
Cost Optimization Hub	costoptimizationhub	AWS_ENDPOINT_URL_COST_OPTIMIZATION_HUB
Cost and Usage Report Service	costandusagereport	AWS_ENDPOINT_URL_COST_AND_USAGE_REPO RT_SERVICE
Customer Profiles	customerprofiles	AWS_ENDPOINT_URL_CUSTOMER_PROFILES
DataBrew	databrew	AWS_ENDPOINT_URL_DATABREW
DataExchange	dataexchange	AWS_ENDPOINT_URL_DATAEXCHANGE

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
Data Pipeline	<code>data-pipeline</code>	<code>AWS_ENDPOINT_URL_DATA_PIPELINE</code>
DataSync	<code>data-sync</code>	<code>AWS_ENDPOINT_URL_DATASYNC</code>
DataZone	<code>datazone</code>	<code>AWS_ENDPOINT_URL_DATAZONE</code>
DAX	<code>dax</code>	<code>AWS_ENDPOINT_URL_DAX</code>
Detective	<code>detective</code>	<code>AWS_ENDPOINT_URL_DETECTIVE</code>
Device Farm	<code>device-farm</code>	<code>AWS_ENDPOINT_URL_DEVICE_FARM</code>
DevOps Guru	<code>devops-guru</code>	<code>AWS_ENDPOINT_URL_DEVOPS_GURU</code>
Direct Connect	<code>direct-connect</code>	<code>AWS_ENDPOINT_URL_DIRECT_CONNECT</code>
Application Discovery Service	<code>application-discovery-service</code>	<code>AWS_ENDPOINT_URL_APPLICATION_DISCOVERY_SERVICE</code>
DLM	<code>dlm</code>	<code>AWS_ENDPOINT_URL_DLM</code>



<b>serviceId</b>	<b>Service-specific endpoint key for AWS CLI configuration file</b>	<b>environment variable</b>
Database Migration Service	dm	AWS_ENDPOINT_URL_DATABASE_MIGRATION_SERVICE
DocDB	dc	AWS_ENDPOINT_URL_DOCDB
DocDB Elastic	de	AWS_ENDPOINT_URL_DOCDB_ELASTIC
drs	dr	AWS_ENDPOINT_URL_DRS
Directory Service	ds	AWS_ENDPOINT_URL_DIRECTORY_SERVICE
DynamoDB	dy	AWS_ENDPOINT_URL_DYNAMODB
DynamoDB Streams	ds	AWS_ENDPOINT_URL_DYNAMODB_STREAMS
EBS	el	AWS_ENDPOINT_URL_EBS
EC2	ec	AWS_ENDPOINT_URL_EC2
EC2 Instance Connect	ec	AWS_ENDPOINT_URL_EC2_INSTANCE_CONNECT
ECR	ec	AWS_ENDPOINT_URL_ECR

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
ECR PUBLIC	ecr-public	AWS_ENDPOINT_URL_ECR_PUBLIC
ECS	ecs	AWS_ENDPOINT_URL_ECS
EFS	efs	AWS_ENDPOINT_URL_EFS
EKS	eks	AWS_ENDPOINT_URL_EKS
EKS Auth	eks-auth	AWS_ENDPOINT_URL_EKS_AUTH
Elastic Inference	elastic-inference	AWS_ENDPOINT_URL_ELASTIC_INFERENCE
ElastiCache	elasticache	AWS_ENDPOINT_URL_ELASTICACHE
Elastic Beanstalk	elasticbeanstalk	AWS_ENDPOINT_URL_ELASTIC_BEANSTALK
Elastic Transcoder	elastictranscoder	AWS_ENDPOINT_URL_ELASTIC_TRANSCODER
Elastic Load Balancing	elasticloadbalancing	AWS_ENDPOINT_URL_ELASTIC_LOAD_BALANCING

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
Elastic Load Balancing v2	elasticloadbalancingv2	AWS_ENDPOINT_URL_ELASTIC_LOAD_BALANCING_V2
EMR	emr	AWS_ENDPOINT_URL_EMR
EMR containers	emrcontainers	AWS_ENDPOINT_URL_EMR_CONTAINERS
EMR Serverless	emrserverless	AWS_ENDPOINT_URL_EMR_SERVERLESS
EntityResolution	entityresolution	AWS_ENDPOINT_URL_ENTITYRESOLUTION
Elasticsearch Service	elasticsearch	AWS_ENDPOINT_URL_ELASTICSEARCH_SERVICE
EventBridge	eventbridge	AWS_ENDPOINT_URL_EVENTBRIDGE
Evidently	evidently	AWS_ENDPOINT_URL_EVIDENTLY
finspace	finspace	AWS_ENDPOINT_URL_FINSPEC
finspace data	finspace-data	AWS_ENDPOINT_URL_FINSPEC_DATA

<b>serviceId</b>	<b>Service ID</b>	<b>environment variable</b>
Firehose	f: AWS_ENDPOINT_URL_FIREHOSE	
fis	f: AWS_ENDPOINT_URL_FIS	
FMS	fr: AWS_ENDPOINT_URL_FMS	
forecast	fo: AWS_ENDPOINT_URL_FORECAST	
forecastquery	fo: AWS_ENDPOINT_URL_FORECASTQUERY ur	
FraudDetector	f: AWS_ENDPOINT_URL_FRAUDETECTOR ct	
FreeTier	f: AWS_ENDPOINT_URL_FREETIER	
FSx	f: AWS_ENDPOINT_URL_FSX	
GameLift	g: AWS_ENDPOINT_URL_GAMELIFT	
Glacier	g: AWS_ENDPOINT_URL_GLACIER	
Global Accelerator	g: AWS_ENDPOINT_URL_GLOBAL_ACCELERATOR ce	
Glue	g: AWS_ENDPOINT_URL_GLUE	
grafana	g: AWS_ENDPOINT_URL_GRAFANA	

<b>serviceId</b>	<b>Service endpoint</b>	<b>environment variable</b>
Greengrass	g: AWS_ENDPOINT_URL_GREENGRASS	
GreengrassV2	g: AWS_ENDPOINT_URL_GREENGRASSV2	
GroundStation	g: AWS_ENDPOINT_URL_GROUNDSTATION	
GuardDuty	g: AWS_ENDPOINT_URL_GUARDDUTY	
Health	h: AWS_ENDPOINT_URL_HEALTH	
HealthLake	h: AWS_ENDPOINT_URL_HEALTHLAKE	
Honeycode	h: AWS_ENDPOINT_URL_HONEYCODE	
IAM	i: AWS_ENDPOINT_URL_IAM	
identitystore	i: AWS_ENDPOINT_URL_IDENTITYSTORE	
imagebuilder	i: AWS_ENDPOINT_URL_IMAGEBUILDER	

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
ImportExport	<code>AWS_ENDPOINT_URL_IMPORTEXPORT</code>	
Inspector	<code>AWS_ENDPOINT_URL_INSPECTOR</code>	
Inspector Scan	<code>AWS_ENDPOINT_URL_INSPECTOR_SCAN</code>	
Inspector2	<code>AWS_ENDPOINT_URL_INSPECTOR2</code>	
InternetMonitor	<code>AWS_ENDPOINT_URL_INTERNETMONITOR</code>	
IoT	<code>AWS_ENDPOINT_URL_IOT</code>	
IoT Data Plane	<code>AWS_ENDPOINT_URL_IOT_DATA_PLANE</code>	
IoT Jobs Data Plane	<code>AWS_ENDPOINT_URL_IOT_JOBS_DATA_PLANE</code>	
IoT 1Click Devices Service	<code>AWS_ENDPOINT_URL_IOT_1CLICK_DEVICES_SERVICE</code>	

<b>serviceId</b>	<b>Service endpoint</b>	<b>environment variable</b>
IoT 1Click Projects	iot-1click-projects.amazonaws.com	AWS_ENDPOINT_URL_IOT_1CLICK_PROJECTS
IoTAnalytics	iotanalytics.amazonaws.com	AWS_ENDPOINT_URL_IOTANALYTICS
IotDeviceAdvisor	iot-deviceadvisor.amazonaws.com	AWS_ENDPOINT_URL_IOTDEVICEADVISOR
IoT Events	iot-events.amazonaws.com	AWS_ENDPOINT_URL_IOT_EVENTS
IoT Events Data	iot-events-data.amazonaws.com	AWS_ENDPOINT_URL_IOT_EVENTS_DATA
IoTFleetHub	iotfleethub.amazonaws.com	AWS_ENDPOINT_URL_IOTFLEETHUB
IoTFleetWise	iotfleetwise.amazonaws.com	AWS_ENDPOINT_URL_IOTFLEETWISE
IoTSecureTunneling	iot-securetunneling.amazonaws.com	AWS_ENDPOINT_URL_IOTSECURETUNNELING
IoTSiteWise	iot-sitewise.amazonaws.com	AWS_ENDPOINT_URL_IOTSITWISE

<b>serviceId</b>	<b>Service-specific endpoint</b>	<b>environment variable</b>
IoTThingsGraph	iot: AWS_ENDPOINT_URL_IOTTHINGSGRAPH	
IoTTwinMaker	iot: AWS_ENDPOINT_URL_IOTTWINMAKER	
IoT Wireless	iot: AWS_ENDPOINT_URL_IOT_WIRELESS	
ivs	ivs: AWS_ENDPOINT_URL_IVS	
IVS RealTime	ivs: AWS_ENDPOINT_URL_IVS_REALTIME	
ivschat	ivs: AWS_ENDPOINT_URL_IVSCHAT	
Kafka	kafka: AWS_ENDPOINT_URL_KAFKA	
KafkaConnect	kafka: AWS_ENDPOINT_URL_KAFKACONNECT	
kendra	kendra: AWS_ENDPOINT_URL_KENDRA	
Kendra Ranking	kendra: AWS_ENDPOINT_URL_KENDRA_RANKING	
Keyspaces	kafka: AWS_ENDPOINT_URL_KEYSPACES	



<b>serviceId</b>	<b>Service-specific endpoint</b>	<b>environment variable</b>
Kinesis	k: AWS_ENDPOINT_URL_KINESIS	
Kinesis Video Archived Media	k: AWS_ENDPOINT_URL_KINESIS_VIDEO_ARCHIVED_MEDIA	
Kinesis Video Media	k: AWS_ENDPOINT_URL_KINESIS_VIDEO_MEDIA	
Kinesis Video Signaling	k: AWS_ENDPOINT_URL_KINESIS_VIDEO_SIGNALING	
Kinesis Video WebRTC Storage	k: AWS_ENDPOINT_URL_KINESIS_VIDEO_WEBRTC_STORAGE	
Kinesis Analytics	k: AWS_ENDPOINT_URL_KINESIS_ANALYTICS	
Kinesis Analytics V2	k: AWS_ENDPOINT_URL_KINESIS_ANALYTICS_V2	

<b>serviceId</b>	<b>Service-specific endpoint</b>	<b>environment variable</b>
Kinesis Video	kinesis-video	AWS_ENDPOINT_URL_KINESIS_VIDEO
KMS	kms	AWS_ENDPOINT_URL_KMS
LakeFormation	lakeformation	AWS_ENDPOINT_URL_LAKEFORMATION
Lambda	lambda	AWS_ENDPOINT_URL_LAMBDA
Launch Wizard	launch-wizard	AWS_ENDPOINT_URL_LAUNCH_WIZARD
Lex Model Building Service	lex-model-building-service	AWS_ENDPOINT_URL_LEX_MODEL_BUILDING_SERVICE
Lex Runtime Service	lex-runtime-service	AWS_ENDPOINT_URL_LEX_RUNTIME_SERVICE
Lex Models V2	lex-models-v2	AWS_ENDPOINT_URL_LEX_MODELS_V2
Lex Runtime V2	lex-runtime-v2	AWS_ENDPOINT_URL_LEX_RUNTIME_V2
License Manager	license-manager	AWS_ENDPOINT_URL_LICENSE_MANAGER

<b>serviceId</b>	<b>Service-specific endpoint</b>	<b>environment variable</b>
License Manager Linux Subscriptions	LinuxSubscriptionsEndpoint	AWS_ENDPOINT_URL_LICENSE_MANAGER_LINUX_SUBSCRIPTIONS
License Manager User Subscriptions	UserSubscriptionsEndpoint	AWS_ENDPOINT_URL_LICENSE_MANAGER_USER_SUBSCRIPTIONS
Lightsail	LightsailEndpoint	AWS_ENDPOINT_URL_LIGHTSAIL
Location	LocationEndpoint	AWS_ENDPOINT_URL_LOCATION
CloudWatch Logs	CloudWatchLogsEndpoint	AWS_ENDPOINT_URL_CLOUDWATCH_LOGS
CloudWatch Logs	CloudWatchLogsEndpoint	AWS_ENDPOINT_URL_CLOUDWATCH_LOGS
LookoutEquipment	LookoutEquipmentEndpoint	AWS_ENDPOINT_URL_LOOKOUTEQUIPMENT
LookoutMetrics	LookoutMetricsEndpoint	AWS_ENDPOINT_URL_LOOKOUTMETRICS

<b>serviceId</b>	<b>Service-specific endpoint</b>	<b>environment variable</b>
LookoutVision	LookoutVisionEndpoint	AWS_ENDPOINT_URL_LOOKOUTVISION
Macie2	Macie2Endpoint	AWS_ENDPOINT_URL_MACIE2
Machine Learning	MachineLearningEndpoint	AWS_ENDPOINT_URL_MACHINE_LEARNING
ManagedBlockchain	ManagedBlockchainEndpoint	AWS_ENDPOINT_URL_MANAGEDBLOCKCHAIN
ManagedBlockchain Query	ManagedBlockchainQueryEndpoint	AWS_ENDPOINT_URL_MANAGEDBLOCKCHAIN_QUERY
Marketplace Agreement	MarketplaceAgreementEndpoint	AWS_ENDPOINT_URL_MARKETPLACE_AGREEMENT
Marketplace Catalog	MarketplaceCatalogEndpoint	AWS_ENDPOINT_URL_MARKETPLACE_CATALOG
Marketplace Deployment	MarketplaceDeploymentEndpoint	AWS_ENDPOINT_URL_MARKETPLACE_DEPLOYMENT

<b>serviceId</b>	<b>Service ID</b>	<b>AWS_ENDPOINT_URL_&lt;SERVICE&gt;</b>	<b>environment variable</b>
Marketplace Entitlement Service	marketplace-entitlement	AWS_ENDPOINT_URL_MARKETPLACE_ENTITLEMENT_SERVICE	
Marketplace Commerce Analytics	marketplace-commerce-analytics	AWS_ENDPOINT_URL_MARKETPLACE_COMMERCIAL_ANALYTICS	
MediaConnect	mediaconnect	AWS_ENDPOINT_URL_MEDIACONNECT	
MediaConvert	mediaconvert	AWS_ENDPOINT_URL_MEDIACONVERT	
MediaLive	medialive	AWS_ENDPOINT_URL_MEDIALIVE	
MediaPackage	mediapackage	AWS_ENDPOINT_URL_MEDIAPACKAGE	
MediaPackage Vod	mediapackage-vod	AWS_ENDPOINT_URL_MEDIAPACKAGE_VOD	
MediaPackageV2	mediapackagev2	AWS_ENDPOINT_URL_MEDIAPACKAGEV2	

<b>serviceId</b>	<b>Service ID</b>	<b>AWS_ENDPOINT_URL_&lt;SERVICE&gt;</b>	<b>environment variable</b>
MediaStore	media-store	AWS_ENDPOINT_URL_MEDIASTORE	
MediaStore Data	media-store-data	AWS_ENDPOINT_URL_MEDIASTORE_DATA	
MediaTailor	media-tailor	AWS_ENDPOINT_URL_MEDIATAILOR	
Medical Imaging	medical-imaging	AWS_ENDPOINT_URL_MEDICAL_IMAGING	
MemoryDB	memorydb	AWS_ENDPOINT_URL_MEMORYDB	
Marketplace Metering	marketplace-metering	AWS_ENDPOINT_URL_MARKETPLACE_METERING	
Migration Hub	migration-hub	AWS_ENDPOINT_URL_MIGRATION_HUB	
mgn	mgn	AWS_ENDPOINT_URL_MGN	
Migration Hub Refactor Spaces	migration-hub-refactor-spaces	AWS_ENDPOINT_URL_MIGRATION_HUB_REFACTOR_SPACES	

<b>serviceId</b>	<b>Service ID</b>	<b>AWS_ENDPOINT_URL_&lt;SERVICE&gt;</b>	<b>environment variable</b>
MigrationHub Config	m: h g	AWS_ENDPOINT_URL_MIGRATIONHUB_CONFIG	
MigrationHubOrches trator	m: h t:	AWS_ENDPOINT_URL_MIGRATIONHUBORCHESTRATOR	
MigrationHubStrategy	m: h g)	AWS_ENDPOINT_URL_MIGRATIONHUBSTRATEGY	
Mobile	m:	AWS_ENDPOINT_URL_MOBILE	
mq	m:	AWS_ENDPOINT_URL_MQ	
MTurk	m:	AWS_ENDPOINT_URL_MTURK	
MWAA	m:	AWS_ENDPOINT_URL_MWAA	
Neptune	n:	AWS_ENDPOINT_URL_NEPTUNE	
Neptune Graph	n: r:	AWS_ENDPOINT_URL_NEPTUNE_GRAPH	
neptunedata	n: t:	AWS_ENDPOINT_URL_NEPTUNEDATA	

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
Network Firewall	<code>aws://networkfirewall</code>	<code>AWS_ENDPOINT_URL_NETWORK_FIREWALL</code>
NetworkManager	<code>aws://networkmanager</code>	<code>AWS_ENDPOINT_URL_NETWORKMANAGER</code>
NetworkMonitor	<code>aws://networkmonitor</code>	<code>AWS_ENDPOINT_URL_NETWORKMONITOR</code>
nimble	<code>aws://nimble</code>	<code>AWS_ENDPOINT_URL_NIMBLE</code>
OAM	<code>aws://oam</code>	<code>AWS_ENDPOINT_URL_OAM</code>
Omics	<code>aws://omics</code>	<code>AWS_ENDPOINT_URL_OMICS</code>
OpenSearch	<code>aws://opensearch</code>	<code>AWS_ENDPOINT_URL_OPENSEARCH</code>
OpenSearchServerless	<code>aws://opensearchserverless</code>	<code>AWS_ENDPOINT_URL_OPENSEARCHSERVERLESS</code>
OpsWorks	<code>aws://opsworks</code>	<code>AWS_ENDPOINT_URL_OPSWORKS</code>
OpsWorksCM	<code>aws://opsworkscm</code>	<code>AWS_ENDPOINT_URL_OPSWORKSCM</code>
Organizations	<code>aws://organizations</code>	<code>AWS_ENDPOINT_URL_ORGANIZATIONS</code>



<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
OSIS	o: AWS_ENDPOINT_URL_OSIS	
Outposts	o: AWS_ENDPOINT_URL_OUTPOSTS	
p8data	p: AWS_ENDPOINT_URL_P8DATA	
p8data	p: AWS_ENDPOINT_URL_P8DATA	
Panorama	p: AWS_ENDPOINT_URL_PANORAMA	
Payment Cryptography	p: AWS_ENDPOINT_URL_PAYMENT_CRYPTOGRAPHY	
Payment Cryptography Data	p: AWS_ENDPOINT_URL_PAYMENT_CRYPTOGRAPHY_DATA	
Pca Connector Ad	p: AWS_ENDPOINT_URL_PCA_CONNECTOR_AD	
Personalize	p: AWS_ENDPOINT_URL_PERSONALIZE	
Personalize Events	p: AWS_ENDPOINT_URL_PERSONALIZE_EVENTS	

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
Personalize Runtime	personalize-runtime	AWS_ENDPOINT_URL_PERSONALIZE_RUNTIME
PI	pinpoint	AWS_ENDPOINT_URL_PI
Pinpoint	pinpoint	AWS_ENDPOINT_URL_PINPOINT
Pinpoint Email	pinpoint-email	AWS_ENDPOINT_URL_PINPOINT_EMAIL
Pinpoint SMS Voice	pinpoint-sms-voice	AWS_ENDPOINT_URL_PINPOINT_SMS_VOICE
Pinpoint SMS Voice V2	pinpoint-sms-voice-v2	AWS_ENDPOINT_URL_PINPOINT_SMS_VOICE_V2
Pipes	pinpoint-pipes	AWS_ENDPOINT_URL_PIPES
Polly	polly	AWS_ENDPOINT_URL_POLLY
Pricing	pricing	AWS_ENDPOINT_URL_PRICING
PrivateNetworks	private-networks	AWS_ENDPOINT_URL_PRIVATENETWORKS
Proton	proton	AWS_ENDPOINT_URL_PROTON

<b>serviceId</b>	<b>Service-specific endpoint</b>	<b>environment variable</b>
QBusiness	qbusiness.amazonaws.com	AWS_ENDPOINT_URL_QBUSINESS
QConnect	qconnect.amazonaws.com	AWS_ENDPOINT_URL_QCONNECT
QLDB	qldb.amazonaws.com	AWS_ENDPOINT_URL_QLDB
QLDB Session	qldb-session.amazonaws.com	AWS_ENDPOINT_URL_QLDB_SESSION
QuickSight	quicksight.amazonaws.com	AWS_ENDPOINT_URL_QUICKSIGHT
RAM	ram.amazonaws.com	AWS_ENDPOINT_URL_RAM
rbin	rbin.amazonaws.com	AWS_ENDPOINT_URL_RBIN
RDS	rds.amazonaws.com	AWS_ENDPOINT_URL_RDS
RDS Data	rds-data.amazonaws.com	AWS_ENDPOINT_URL_RDS_DATA
Redshift	redshift.amazonaws.com	AWS_ENDPOINT_URL_REDSHIFT
Redshift Data	redshift-data.amazonaws.com	AWS_ENDPOINT_URL_REDSHIFT_DATA
Redshift Serverless	redshift-serverless.amazonaws.com	AWS_ENDPOINT_URL_REDSHIFT_SERVERLESS

<b>serviceId</b>	<b>Service endpoint URL</b>	<b>environment variable</b>
Rekognition	<code>AWS_ENDPOINT_URL_REKOGNITION</code>	
repostspace	<code>AWS_ENDPOINT_URL_REPOSTSPACE</code>	
resiliencehub	<code>AWS_ENDPOINT_URL_RESILIENCEHUB</code>	
Resource Explorer 2	<code>AWS_ENDPOINT_URL_RESOURCE_EXPLORER_2</code>	
Resource Groups	<code>AWS_ENDPOINT_URL_RESOURCE_GROUPS</code>	
Resource Groups Tagging API	<code>AWS_ENDPOINT_URL_RESOURCE_GROUPS_TAGGING_API</code>	
RoboMaker	<code>AWS_ENDPOINT_URL_ROBOMAKER</code>	
RolesAnywhere	<code>AWS_ENDPOINT_URL_ROLESEANYWHERE</code>	
Route 53	<code>AWS_ENDPOINT_URL_ROUTE_53</code>	

<b>serviceId</b>	<b>Service ID</b>	<b>Environment Variable</b>
	<b>Service ID</b>	<b>environment variable</b>
Route53 Recovery Cluster	Route53RecoveryCluster	AWS_ENDPOINT_URL_ROUTE53_RECOVERY_CLUSTER
Route53 Recovery Control Config	Route53RecoveryControlConfig	AWS_ENDPOINT_URL_ROUTE53_RECOVERY_CONTROL_CONFIG
Route53 Recovery Readiness	Route53RecoveryReadiness	AWS_ENDPOINT_URL_ROUTE53_RECOVERY_READINESS
Route 53 Domains	Route53Domains	AWS_ENDPOINT_URL_ROUTE_53_DOMAINS
Route53Resolver	Route53Resolver	AWS_ENDPOINT_URL_ROUTE53RESOLVER
RUM	AmazonRecoveryMemoryUsage	AWS_ENDPOINT_URL_RUM
S3	AmazonS3	AWS_ENDPOINT_URL_S3
S3 Control	AmazonS3Control	AWS_ENDPOINT_URL_S3_CONTROL
S3Outposts	AmazonS3Outposts	AWS_ENDPOINT_URL_S3OUTPOSTS

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
SageMaker	sagemaker	AWS_ENDPOINT_URL_SAGEMAKER
SageMaker A2I Runtime	sagemaker-a2i-runtime	AWS_ENDPOINT_URL_SAGEMAKER_A2I_RUNTIME
Sagemaker Edge	sagemaker-edge	AWS_ENDPOINT_URL_SAGEMAKER_EDGE
SageMaker FeatureStore Runtime	sagemaker-featurestore-runtime	AWS_ENDPOINT_URL_SAGEMAKER_FEATURESTORE_RUNTIME
SageMaker Geospatial	sagemaker-geospatial	AWS_ENDPOINT_URL_SAGEMAKER_GEOSPATIAL
SageMaker Metrics	sagemaker-metrics	AWS_ENDPOINT_URL_SAGEMAKER_METRICS
SageMaker Runtime	sagemaker-runtime	AWS_ENDPOINT_URL_SAGEMAKER_RUNTIME
savingsplans	savingsplans	AWS_ENDPOINT_URL_SAVINGSPLANS

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
Scheduler	Service-specific endpoint URL for Scheduler	AWS_ENDPOINT_URL_SCHEDULER
schemas	Service-specific endpoint URL for schemas	AWS_ENDPOINT_URL_SCHEMAS
SimpleDB	Service-specific endpoint URL for SimpleDB	AWS_ENDPOINT_URL_SIMPLEDB
Secrets Manager	Service-specific endpoint URL for Secrets Manager	AWS_ENDPOINT_URL_SECRETS_MANAGER
SecurityHub	Service-specific endpoint URL for SecurityHub	AWS_ENDPOINT_URL_SECURITYHUB
SecurityLake	Service-specific endpoint URL for SecurityLake	AWS_ENDPOINT_URL_SECURITYLAKE
ServerlessApplicationRepository	Service-specific endpoint URL for ServerlessApplicationRepository	AWS_ENDPOINT_URL_SERVERLESSAPPLICATIONREPOSITORY
Service Quotas	Service-specific endpoint URL for Service Quotas	AWS_ENDPOINT_URL_SERVICE_QUOTAS
Service Catalog	Service-specific endpoint URL for Service Catalog	AWS_ENDPOINT_URL_SERVICE_CATALOG

<b>serviceId</b>	<b>Service-specific endpoint key for AWS CLI configuration file</b>	<b>environment variable</b>
Service Catalog AppRegistry	servicecatalogappregistry	AWS_ENDPOINT_URL_SERVICE_CATALOG_APP_REGISTRY
ServiceDiscovery	servicediscovery	AWS_ENDPOINT_URL_SERVICEDISCOVERY
SES	ses	AWS_ENDPOINT_URL_SES
SESV2	sesv2	AWS_ENDPOINT_URL_SESV2
Shield	shield	AWS_ENDPOINT_URL_SHIELD
signer	signer	AWS_ENDPOINT_URL_SIGNER
SimSpaceWeaver	simspacesweaver	AWS_ENDPOINT_URL_SIMSPACEWEAVER
SMS	sms	AWS_ENDPOINT_URL_SMS
Snow Device Management	snowdevicemanagement	AWS_ENDPOINT_URL_SNOW_DEVICE_MANAGEMENT
Snowball	snowball	AWS_ENDPOINT_URL_SNOWBALL
SNS	sns	AWS_ENDPOINT_URL_SNS



<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
SQS	service-specific endpoint URL for Amazon SQS	AWS_ENDPOINT_URL_SQS
SSM	service-specific endpoint URL for Amazon SSM	AWS_ENDPOINT_URL_SSM
SSM Contacts	service-specific endpoint URL for Amazon SSM Contacts	AWS_ENDPOINT_URL_SSM_CONTACTS
SSM Incidents	service-specific endpoint URL for Amazon SSM Incidents	AWS_ENDPOINT_URL_SSM_INCIDENTS
Ssm Sap	service-specific endpoint URL for Amazon SSM SAP	AWS_ENDPOINT_URL_SSM_SAP
SSO	service-specific endpoint URL for Amazon SSO	AWS_ENDPOINT_URL_SSO
SSO Admin	service-specific endpoint URL for Amazon SSO Admin	AWS_ENDPOINT_URL_SSO_ADMIN
SSO OIDC	service-specific endpoint URL for Amazon SSO OIDC	AWS_ENDPOINT_URL_SSO_OIDC
SFN	service-specific endpoint URL for Amazon SFN	AWS_ENDPOINT_URL_SFN
Storage Gateway	service-specific endpoint URL for Amazon Storage Gateway	AWS_ENDPOINT_URL_STORAGE_GATEWAY
STS	service-specific endpoint URL for Amazon STS	AWS_ENDPOINT_URL_STS
SupplyChain	service-specific endpoint URL for Amazon SupplyChain	AWS_ENDPOINT_URL_SUPPLYCHAIN
Support	service-specific endpoint URL for Amazon Support	AWS_ENDPOINT_URL_SUPPORT

<b>serviceId</b>	<b>Service ID</b>	<b>environment variable</b>
Support App	support-app	AWS_ENDPOINT_URL_SUPPORT_APP
SWF	swf	AWS_ENDPOINT_URL_SWF
synthetics	synthetics	AWS_ENDPOINT_URL_SYNTHETICS
Textract	textract	AWS_ENDPOINT_URL_TEXTRACT
Timestream InfluxDB	timestream-influxdb	AWS_ENDPOINT_URL_TIMESTREAM_INFLUXDB
Timestream Query	timestream-query	AWS_ENDPOINT_URL_TIMESTREAM_QUERY
Timestream Write	timestream-write	AWS_ENDPOINT_URL_TIMESTREAM_WRITE
tnb	tnb	AWS_ENDPOINT_URL_TNB
Transcribe	transcribe	AWS_ENDPOINT_URL_TRANSCRIBE
Transfer	transfer	AWS_ENDPOINT_URL_TRANSFER
Translate	translate	AWS_ENDPOINT_URL_TRANSLATE

<b>serviceId</b>	<b>Service-specific endpoint URL</b>	<b>environment variable</b>
TrustedAdvisor	tr: AWS_ENDPOINT_URL_TRUSTEDADVISOR	
VerifiedPermissions	vp: AWS_ENDPOINT_URL_VERIFIEDPERMISSIONS	
Voice ID	vi: AWS_ENDPOINT_URL_VOICE_ID	
VPC Lattice	vl: AWS_ENDPOINT_URL_VPC_LATTICE	
WAF	wi: AWS_ENDPOINT_URL_WAF	
WAF Regional	wi: AWS_ENDPOINT_URL_WAF_REGIONAL	
WAFV2	wi: AWS_ENDPOINT_URL_WAFV2	
WellArchitected	wi: AWS_ENDPOINT_URL_WELLARCHITECTED	
Wisdom	wi: AWS_ENDPOINT_URL_WISDOM	
WorkDocs	wi: AWS_ENDPOINT_URL_WORKDOCS	
WorkLink	wi: AWS_ENDPOINT_URL_WORKLINK	
WorkMail	wi: AWS_ENDPOINT_URL_WORKMAIL	

<b>serviceId</b>	<b>Service ID</b>	<b>environment variable</b>
WorkMailMessageFlow	w: AWS_ENDPOINT_URL_WORKMAILMESSAGEFLOW	
WorkSpaces	w: AWS_ENDPOINT_URL_WORKSPACES	
WorkSpaces Thin Client	w: AWS_ENDPOINT_URL_WORKSPACES_THIN_CLIENT	
WorkSpaces Web	w: AWS_ENDPOINT_URL_WORKSPACES_WEB	
XRay	x: AWS_ENDPOINT_URL_XRAY	

## Smart configuration defaults

With the smart configuration defaults feature, Amazon SDKs can provide predefined, optimized default values for other configuration settings.

Configure this functionality by using the following:

**defaults\_mode** - shared Amazon config file setting, **AWS\_DEFAULTS\_MODE** - environment variable, **aws.defaultsMode** - JVM system property: Java/Kotlin only

With this setting, you can choose a mode that aligns with your application architecture, which then provides optimized default values for your application. If an Amazon SDK setting has

a value explicitly set, then that value always takes precedence. If an Amazon SDK setting does not have a value explicitly set, and `defaults_mode` is not equal to `legacy`, then this feature can provide different default values for various settings optimized for your application. Settings may include the following: HTTP communication settings, retry behavior, service Regional endpoint settings, and, potentially, any SDK-related configuration. Customers who use this feature can get new configuration defaults tailored to common usage scenarios. If your `defaults_mode` is not equal to `legacy`, we recommend performing tests of your application when you upgrade the SDK, because the default values provided might change as best practices evolve.

**Default value:** `legacy`

Note: New major versions of SDKs will default to `standard`.

**Valid values:**

- `legacy` – Provides default settings that vary by SDK and existed before establishment of `defaults_mode`.
- `standard` – Provides the latest recommended default values that should be safe to run in most scenarios.
- `in-region` – Builds on the `standard` mode and includes optimization tailored for applications that call Amazon Web Services from within the same Amazon Web Services Region.
- `cross-region` – Builds on the `standard` mode and includes optimization tailored for applications that call Amazon Web Services in a different Region.
- `mobile` – Builds on the `standard` mode and includes optimization tailored for mobile applications.
- `auto` – Builds on the `standard` mode and includes experimental features. The SDK attempts to discover the runtime environment to determine the appropriate settings automatically. The auto detection is heuristics-based and does not provide 100% accuracy. If the runtime environment can't be determined, `standard` mode is used. The auto detection might query [Instance metadata and user data](#), which might introduce latency. If startup latency is critical to your application, we recommend choosing an explicit `defaults_mode` instead.

Example of setting this value in the config file:

```
[default]
defaults_mode = standard
```

The following parameters might be optimized based on the selection of `defaults_mode`:

- `retryMode` – Specifies how the SDK attempts retries. See [Retry behavior](#).
- `stsRegionalEndpoints` – Specifies how the SDK determines the Amazon Web Service endpoint that it uses to talk to the Amazon Security Token Service (Amazon STS). See [Amazon STS Regionalized endpoints](#).
- `s3UsEast1RegionalEndpoints` – Specifies how the SDK determines the Amazon service endpoint that it uses to talk to the Amazon S3 for the `us-east-1` Region.
- `connectTimeoutInMillis` – After making an initial connection attempt on a socket, the amount of time before timing out. If the client does not receive a completion of the connect handshake, the client gives up and fails the operation.
- `tlsNegotiationTimeoutInMillis` – The maximum amount of time that a TLS handshake can take from the time the CLIENT HELLO message is sent to the time the client and server have fully negotiated ciphers and exchanged keys.

The default value for each setting changes depending on the `defaults_mode` selected for your application. These values are currently set as follows (subject to change):

Parameter	standard mode	in-region mode	cross-region mode	mobile mode
<code>retryMode</code>	standard	standard	standard	standard
<code>stsRegionalEndpoints</code>	regional	regional	regional	regional
<code>s3UsEast1RegionalEndpoints</code>	regional	regional	regional	regional
<code>connectTimeoutInMillis</code>	3100	1100	3100	30000
<code>tlsNegotiationTime</code>	3100	1100	3100	30000

Parameter	standard mode	in-region mode	cross-region mode	mobile mode
outInMillis				

For example, if the `defaults_mode` you selected was `standard`, then the value of `standard` would be assigned for `retry_mode` (from the valid `retry_mode` options) and the value of `regional` would be assigned for `stsRegionalEndpoints` (from the valid `stsRegionalEndpoints` options).

## Compatibility with Amazon SDKs

The following SDKs support the features and settings described in this topic. Any partial exceptions are noted. Any JVM system property settings are supported by the Amazon SDK for Java and the Amazon SDK for Kotlin only.

SDK	Supported	Notes or more information
<a href="#">Amazon CLI v2</a>	No	
<a href="#">SDK for C++</a>	Yes	Parameters not optimized: <code>stsRegionalEndpoints</code> , <code>s3UsEast1RegionalEndpoints</code> , <code>tlsNegotiationTimeoutInMillis</code> .
<a href="#">SDK for Go V2 (1.x)</a>	Yes	Parameters not optimized: <code>retryMode</code> , <code>stsRegionalEndpoints</code> , <code>s3UsEast1RegionalEndpoints</code> .
<a href="#">SDK for Go 1.x (V1)</a>	No	
<a href="#">SDK for Java 2.x</a>	Yes	Parameters not optimized: <code>stsRegionalEndpoints</code> .

SDK	Supported	Notes or more information
<a href="#">SDK for Java 1.x</a>	No	
<a href="#">SDK for JavaScript 3.x</a>	Yes	Parameters not optimized: <code>stsRegionalEndpoints</code> , <code>s3UsEast1RegionalEndpoints</code> , <code>tlsNegotiationTimeoutInMillis</code> is <code>connectTimeoutInMillis</code> is called <code>connectionTimeout</code> .
<a href="#">SDK for JavaScript 2.x</a>	No	
<a href="#">SDK for Kotlin</a>	No	
<a href="#">SDK for .NET 3.x</a>	Yes	Parameters not optimized : <code>connectTimeoutInMillis</code> , <code>tlsNegotiationTimeoutInMillis</code> is .
<a href="#">SDK for PHP 3.x</a>	Yes	Parameters not optimized : <code>tlsNegotiationTimeoutInMillis</code> .
<a href="#">SDK for Python (Boto3)</a>	Yes	Parameters not optimized : <code>tlsNegotiationTimeoutInMillis</code> .
<a href="#">SDK for Ruby 3.x</a>	Yes	
<a href="#">SDK for Rust</a>	No	



SDK	Supported	Notes or more information
<a href="#">Tools for PowerShell</a>	Yes	Parameters not optimized :connectTimeoutInMi llis ,tlsNegoti ationTimeoutInMill is .

# Amazon Common Runtime (CRT) libraries

The Amazon Common Runtime (CRT) libraries are a base library of the SDKs. The CRT is a modular family of independent packages, written in C. Each package provides good performance and minimal footprint for different required functionalities. These functionalities are common and shared across all SDKs providing better code reuse, optimization, and accuracy. The packages are:

- [awslabs/aws-c-auth](#): Amazon client-side authentication (standard credential providers and signing (sigv4))
- [awslabs/aws-c-cal](#): Cryptographic primitive types, hashes (MD5, SHA256, SHA256 HMAC), signers, AES
- [awslabs/aws-c-common](#): Basic data structures, threading/synchronization primitive types, buffer management, stdlib-related functions
- [awslabs/aws-c-compression](#): Compression algorithms (Huffman encoding/decoding)
- [awslabs/aws-c-event-stream](#): Event stream message processing (headers, prelude, payload, crc/trailer), remote procedure call (RPC) implementation over event streams
- [awslabs/aws-c-http](#): C99 implementation of the HTTP/1.1 and HTTP/2 specifications
- [awslabs/aws-c-io](#): Sockets (TCP, UDP), DNS, pipes, event loops, channels, SSL/TLS
- [awslabs/aws-c-iot](#): C99 implementation of Amazon IoT cloud services integration with devices
- [awslabs/aws-c-mqtt](#): Standard, lightweight messaging protocol for the Internet of Things (IoT)
- [awslabs/aws-c-s3](#): C99 library implementation for communicating with the Amazon S3 service, designed for maximizing throughput on high bandwidth Amazon EC2 instances
- [awslabs/aws-c-sdkutils](#): A utilities library for parsing and managing Amazon profiles
- [awslabs/aws-checksums](#): Cross-platform hardware-accelerated CRC32c and CRC32 with fallback to efficient software implementations
- [awslabs/aws-lc](#): General-purpose cryptographic library maintained by the Amazon Cryptography team for Amazon and their customers, based on code from the Google BoringSSL project and the OpenSSL project
- [awslabs/s2n](#): C99 implementation of the TLS/SSL protocols, designed to be small and fast with security as a priority

The CRT is available through all SDKs except Go.

## CRT dependencies

The CRT libraries form a complex net of relationships and dependencies. Knowing these relationships is helpful if you need to build the CRT directly from source. However, most users access CRT functionality through their language SDK (such as Amazon SDK for C++ or Amazon SDK for Java) or their language IoT device SDK (such as Amazon IoT SDK for C++ or Amazon IoT SDK for Java). In the following diagram, the Language CRT Bindings box refers to the package wrapping the CRT libraries for a specific language SDK. This is a collection of packages of the form `aws-crt-*`, where '\*' is an SDK language (such as [aws-crt-cpp](#) or [aws-crt-java](#)).

*The following is an illustration of the hierarchical dependencies of the CRT libraries.*

# Maintenance and support

For an overview of tools that can help you develop applications on Amazon, see [Tools to Build on Amazon](#). For information on support, see the [Amazon Knowledge Center](#).

The following topics cover the maintenance and version support policies for Amazon SDKs.

## Topics

- [Amazon SDKs and Tools maintenance policy](#)
- [Amazon SDKs and Tools version support matrix](#)
- [IDE Toolkits](#)

## Amazon SDKs and Tools maintenance policy

### Overview

This document outlines the maintenance policy for Amazon Software Development Kits (SDKs) and Tools, including Mobile and IoT SDKs, and their underlying dependencies. Amazon regularly provides the Amazon SDKs and Tools with updates that may contain support for new or updated Amazon APIs, new features, enhancements, bug fixes, security patches, or documentation updates. Updates may also address changes with dependencies, language runtimes, and operating systems. Amazon SDK releases are published to package managers (e.g. Maven, NuGet, PyPI), and are available as source code on GitHub.

We recommend users to stay up-to-date with SDK releases to keep up with the latest features, security updates, and underlying dependencies. Continued use of an unsupported SDK version is not recommended and is done at the user's discretion.

### Versioning

The Amazon SDK release versions are in the form of X.Y.Z where X represents the major version. Increasing the major version of an SDK indicates that this SDK underwent significant and substantial changes to support new idioms and patterns in the language. Major versions are introduced when public interfaces (e.g. classes, methods, types, etc.), behaviors, or semantics have changed. Applications need to be updated in order for them to work with the newest SDK version. It is important to update major versions carefully and in accordance with the upgrade guidelines provided by Amazon.

## SDK major version lifecycle

The lifecycle for major SDKs and Tools versions consists of 5 phases, which are outlined below.

- *Developer Preview (Phase 0)* - During this phase, SDKs are not supported, should not be used in production environments, and are meant for early access and feedback purposes only. It is possible for future releases to introduce breaking changes. Once Amazon identifies a release to be a stable product, it may mark it as a Release Candidate. Release Candidates are ready for GA release unless significant bugs emerge, and will receive full Amazon support.
- *General Availability (GA) (Phase 1)* - During this phase, SDKs are fully supported. Amazon will provide regular SDK releases that include support for new services, API updates for existing services, as well as bug and security fixes. For Tools, Amazon will provide regular releases that include new feature updates and bug fixes. Amazon will support the GA version of an SDK for *at least 24 months*.
- *Maintenance Announcement (Phase 2)* - Amazon will make a public announcement at least 6 months before an SDK enters maintenance mode. During this period, the SDK will continue to be fully supported. Typically, maintenance mode is announced at the same time as the next major version is transitioned to GA.
- *Maintenance (Phase 3)* - During the maintenance mode, Amazon limits SDK releases to address critical bug fixes and security issues only. An SDK will not receive API updates for new or existing services, or be updated to support new regions. Maintenance mode has a *default duration of 12 months*, unless otherwise specified.
- *End-of-Support (Phase 4)* - When an SDK reaches end-of support, it will no longer receive updates or releases. Previously published releases will continue to be available via public package managers and the code will remain on GitHub. The GitHub repository may be archived. Use of an SDK which has reached end-of-support is done at the user's discretion. We recommend users upgrade to the new major version.

*The following is a visual illustration of the SDK major version lifecycle. Please note that the timelines shown below are illustrative and not binding.*

## Dependency lifecycle

Most Amazon SDKs have underlying dependencies, such as language runtimes, operating systems, or third party libraries and frameworks. These dependencies are typically tied to the language

community or the vendor who owns that particular component. Each community or vendor publishes their own end-of-support schedule for their product.

The following terms are used to classify underlying third party dependencies:

- *Operating System (OS)*: Examples include Amazon Linux AMI, Amazon Linux 2, Windows 2008, Windows 2012, Windows 2016, etc.
- *Language Runtime*: Examples include Java 7, Java 8, Java 11, .NET Core, .NET Standard, .NET PCL, etc.
- *Third party Library / Framework*: Examples include OpenSSL, .NET Framework 4.5, Java EE, etc.

Our policy is to continue supporting SDK dependencies for at least 6 months after the community or vendor ends support for the dependency. This policy, however, could vary depending on the specific dependency.

#### Note

Amazon reserves the right to stop support for an underlying dependency without increasing the major SDK version

## Communication methods

Maintenance announcements are communicated in several ways:

- An email announcement is sent to affected accounts, announcing our plans to end support for the specific SDK version. The email will outline the path to end-of-support, specify the campaign timelines, and provide upgrade guidance.
- Amazon SDK documentation, such as API reference documentation, user guides, SDK product marketing pages, and GitHub readme(s) are updated to indicate the campaign timeline and provide guidance on upgrading affected applications.
- An Amazon blog post is published that outlines the path to end-of-support, as well as reiterates the campaign timelines.
- Deprecation warnings are added to the SDKs, outlining the path to end-of-support and linking to the SDK documentation.

To see the list of available major versions of Amazon SDKs and Tools and where they are in their maintenance lifecycle, see [the section called "Version support matrix"](#).

## Amazon SDKs and Tools version support matrix

The matrix below shows the list of available Amazon Software Development Kit (SDK) major versions and where they are in the maintenance lifecycle with associated timelines. For detailed information on the lifecycle for the major versions of Amazon SDKs and Tools and their underlying dependencies, see [the section called "Maintenance policy"](#).

SDK	Major version	Current Phase	General Availability Date	Notes
<a href="#">Amazon CLI</a>	1.x	General Availability	9/2/2013	
<a href="#">Amazon CLI</a>	2.x	General Availability	2/10/2020	
<a href="#">SDK for C++</a>	1.x	General Availability	9/2/2015	
<a href="#">SDK for Go V2</a>	V2 1.x	General Availability	1/19/2021	
<a href="#">SDK for Go</a>	1.x	Maintenance Announcement	11/19/2015	See <a href="#">announcement</a> for details and dates
<a href="#">SDK for Java</a>	1.x	Maintenance Announcement	3/25/2010	See <a href="#">announcement</a> for details and dates
<a href="#">SDK for Java</a>	2.x	General Availability	11/20/2018	
<a href="#">SDK for JavaScript</a>	1.x	End-of-Support	5/6/2013	

SDK	Major version	Current Phase	General Availability Date	Notes
<a href="#">SDK for JavaScript</a>	2.x	Maintenance Announcement	6/19/2014	See <a href="#">announcement</a> for details and dates
<a href="#">SDK for JavaScript</a>	3.x	General Availability	12/15/2020	
<a href="#">SDK for Kotlin</a>	1.x	General Availability	11/27/2023	
<a href="#">SDK for .NET</a>	1.x	End-of-Support	11/2009	
<a href="#">SDK for .NET</a>	2.x	End-of-Support	11/8/2013	
<a href="#">SDK for .NET</a>	3.x	General Availability	7/28/2015	
<a href="#">SDK for PHP</a>	2.x	End-of-Support	11/2/2012	
<a href="#">SDK for PHP</a>	3.x	General Availability	5/27/2015	
<a href="#">SDK for Python (Boto2)</a>	1.x	End-of-Support	7/13/2011	
<a href="#">SDK for Python (Boto3)</a>	1.x	General Availability	6/22/2015	
<a href="#">SDK for Python (Botocore)</a>	1.x	General Availability	6/22/2015	
<a href="#">SDK for Ruby</a>	1.x	End-of-Support	7/14/2011	
<a href="#">SDK for Ruby</a>	2.x	End-of-Support	2/15/2015	



SDK	Major version	Current Phase	General Availability Date	Notes
<a href="#">SDK for Ruby</a>	3.x	General Availability	8/29/2017	
<a href="#">SDK for Rust</a>	1.x	General Availability	11/27/2023	
<a href="#">SDK for Swift</a>	1.x	Developer Preview		
<a href="#">Tools for PowerShell</a>	2.x	End-of-Support	11/8/2013	
<a href="#">Tools for PowerShell</a>	3.x	End-of-Support	7/29/2015	
<a href="#">Tools for PowerShell</a>	4.x	General Availability	11/21/2019	

## IDE Toolkits

Amazon Integrated Development Environment (IDE) Toolkits are plugins and extensions that enable access to Amazon services from your IDE.

For detailed information about each of the IDE Toolkits, see these Toolkit User Guides:

- [Amazon Toolkit for Visual Studio](#)
- [Amazon Toolkit for Visual Studio Code](#)
- [Amazon Toolkit for JetBrains](#)

The following sections contain support information, maintenance reports, and notifications for Amazon IDE Toolkits.

## Telemetry Notification

Amazon IDE Toolkits may collect and store client-side telemetry data to inform decisions regarding future Amazon Toolkit releases. The data collected quantifies your usage of the Amazon Toolkit.

To learn more about the telemetry data collected across all of the Amazon IDE Toolkits, see the [commonDefinitions.json](#) document in the `aws-toolkit-common` Github repository.

For detailed information about the telemetry data collected by each of the Amazon IDE Toolkits, reference the resource documents in the following Amazon Toolkits' Github repositories:

- [Amazon Toolkit for Visual Studio](#)
- [Amazon Toolkit for Visual Studio Code](#)
- [Amazon Toolkit for JetBrains](#)

Certain Amazon services that are accessible from the Toolkits may collect additional client-side telemetry data, such as Amazon CodeWhisperer. For detailed information about the type of data collected by CodeWhisperer or how to opt out of client-side telemetry for CodeWhisperer, see the [Sharing your data with Amazon](#) topic in the *Amazon CodeWhisperer User Guide*.

# Document history for Amazon SDKs and Tools Reference Guide

The following table describes important additions and updates to the *Amazon SDKs and Tools Reference Guide*. For notification about updates to this documentation, you can subscribe to the RSS feed.

Change	Description	Date
<a href="#">Settings updates</a>	Add JVM system configuration settings.	March 27, 2024
<a href="#">Compatibility table updates</a>	Updates to compatibility for SDK support, updates to IAM Identity Center procedures.	February 20, 2024
<a href="#">Container credential update. IMDS update.</a>	Adding support for Amazon EKS. Adding setting to disable IMDSv1 fallback.	December 29, 2023
<a href="#">Request compression</a>	Adding settings for request compression feature.	December 27, 2023
<a href="#">Compatibility tables</a>	Compatibility tables for SDK and tool features updated to include SDK for Kotlin, SDK for Rust, and Amazon Tools for PowerShell.	December 10, 2023
<a href="#">Authentication updates</a>	Updates to supported methods of authentication for SDKs and tools.	July 1, 2023
<a href="#">IAM best practices updates</a>	Updated guide to align with the IAM best practices . For more information, see <a href="#">Security best practices in IAM</a> .	February 27, 2023

---

<a href="#">SSO updates</a>	Updates to SSO credentials for the new SSO token configuration.	November 19, 2022
<a href="#">Settings updates</a>	Updates to support table for General configuration and for Amazon S3 Multi-Region Access Points.	November 17, 2022
<a href="#">Settings updates</a>	Updates to clarity of IMDS client and IMDS credentials. Updates to Environment variables.	November 4, 2022
<a href="#">Updating welcome page</a>	Announcing Amazon CodeWhisperer.	September 22, 2022
<a href="#">Service name change for single sign-on</a>	Updates to reflect that Amazon SSO is now referred to as Amazon IAM Identity Center.	July 26, 2022
<a href="#">Settings update</a>	Minor updates to config file details and to supported settings.	June 15, 2022
<a href="#">Update</a>	Massive update of almost all parts of this guide.	February 1, 2022
<a href="#">Initial release</a>	The first release of this guide is released to the public.	March 13, 2020

# Amazon Glossary

For the latest Amazon terminology, see the [Amazon glossary](#) in the *Amazon Web Services Glossary Reference*.