Developer Guide

AWS SDK for Swift

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

AWS SDK for Swift: Developer Guide

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

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

Services or capabilities described in AWS documentation might vary by Region. To see the differences applicable to the AWS European Sovereign Cloud Region, see the AWS European Sovereign Cloud User Guide.

Table of Contents

AWS SDK for Swift 1
What is the AWS SDK for Swift? 1
Supported targets 1
Get started with the SDK 1
About this guide 2
Maintenance and support for SDK major versions 2
Additional resources 3
Contributing to the SDK 3
Getting started 4
Setting up 4
Overview 4
Setting up your Swift development environment 8
Adding the AWS SDK for Swift as a dependency 9
Security and authentication when testing on macOS 9
Building and running a SwiftPM project 13
Importing AWS SDK for Swift libraries into source files 13
Next steps 14
Authenticating with AWS 14
Using console credentials 14
Using IAM Identity Center 15
Creating an application 17
Creating a project using the AWS SDK for Swift 17
Creating the project 18
Configuring the package 19
Accessing AWS services using Swift 20
Adding the example entry point 23
Additional information 24
Configuring service clients 25
Client configuration externally 25
Client configuration in code 27
Configuration data types 27
Configuring a client 28
AWS Region 30
Specifying the Region programmatically 28

Credential providers ... 31
The credential provider chain ... 32
Credential retrieval order ... 32
Credential identity resolvers ... 34
Getting credentials from an identity ... 36
Login credentials ... 36
SSO credentials ... 37
Static credentials ... 39
Amazon Cognito Identity ... 40
Additional information ... 41
Retry ... 41
Default configuration ... 41
Configuring retry ... 41
HTTP ... 43
Configuring HTTP timeouts ... 43
Customizing HTTP headers ... 44
HTTP/1 vs. HTTP/2 ... 45
Using a custom HTTP client ... 45
Using the SDK ... 46
Making AWS service requests ... 46
Creating and using AWS client objects ... 47
Specifying service client function parameters ... 48
Calling SDK functions ... 48
Handling errors ... 50
Overview ... 50
Error protocols ... 51
Handling errors ... 53
Pagination ... 56
Testing and debugging ... 58
Logging ... 58
Mocking the AWS SDK for Swift ... 60
Example: Mocking an Amazon S3 function ... 60
Waiters ... 66
How to use waiters ... 66
Example: Waiting for an S3 bucket to exist ... 67
Lambda functions ... 68

Overview ... 68
Set up ... 69
Creating a Lambda function ... 70
Build and test locally ... 74
Packaging and uploading the app ... 75
Additional information ... 24
Event streaming ... 76
Overview ... 76
Event streaming example ... 77
Additional information ... 24
Presigning requests ... 83
Overview ... 83
Presigning basics ... 83
Advanced presigning configuration ... 84
Integrating with Apple ... 85
Adding the AWS SDK for Swift to an existing Xcode project ... 85
Sign In With Apple ... 90
Working with AWS services ... 102
Amazon S3 ... 102
Multipart uploads ... 103
Binary streaming ... 108
Data integrity protection with checksums ... 112
Code examples ... 119
Amazon Bedrock ... 120
Get started ... 120
Actions ... 124
Amazon Bedrock Runtime ... 125
Amazon Nova ... 125
Amazon Nova Canvas ... 129
Amazon Nova Reel ... 131
Anthropic Claude ... 134
Meta Llama ... 137
Amazon Cognito Identity ... 141
Actions ... 124
Amazon Cognito Identity Provider ... 146
Actions ... 124

Scenarios 159
DynamoDB 174
Basics 175
Actions 124
Amazon EC2 212
Get started 120
Basics 175
Actions 124
AWS Glue 266
Basics 175
Actions 124
IAM 307
Actions 124
Lambda 328
Basics 175
Actions 124
Amazon RDS 359
Basics 175
Actions 124
Amazon S3 395
Basics 175
Actions 124
Scenarios 159
Amazon SNS 421
Get started 120
Actions 124
Scenarios 159
Amazon SQS 447
Get started 120
Actions 124
Scenarios 159
AWS STS 476
Actions 124
Amazon Transcribe Streaming 477
Actions 124
Scenarios 159

Security **485**
Data protection 485
Encryption in transit 486
Identity and Access Management 487
Audience 487
Authenticating with identities 488
Managing access using policies 489
How AWS services work with IAM 491
Troubleshooting AWS identity and access 491
Compliance Validation 493
Resilience 493
Infrastructure Security 494
Document history **495**

AWS SDK for Swift Developer Guide

What is the AWS SDK for Swift?

Welcome to the AWS SDK for Swift, a pure Swift SDK that makes it easier to develop tools that take advantage of AWS services, including Amazon S3, Amazon EC2, DynamoDB, and more, all using the Swift programming language.

Supported targets

Platform	Operating system(s)	Processor	Notes
Apple (Darwin)	macOS, iOS/iPadOS, tvOS, watchOS, and visionOS.	Intel x86, Apple Silicon (arm64).	All Apple targets support 64-bit only. The only exception is watchOS, which also supports the hybrid 64/32-bit processors used in the Series 1 through Series 3 Apple watches.
Linux	Amazon Linux 2 and Ubuntu.	Intel x86, arm64.	
AWS Lambda	Amazon Linux 2	Intel x86, arm64.	

The AWS SDK for Swift is not currently built, tested, or supported on Microsoft Windows. Support for Windows targets could be added in the future if customer demand warrants it.

Get started with the SDK

To set up your development environment, see ???. Then you can test drive the AWS SDK for Swift by creating your first project in ???.

For information on making requests to Amazon S3, DynamoDB, Amazon EC2, and other AWS services, see *Using the SDK*.

About this guide

The AWS SDK for Swift Developer Guide covers how to install, configure, and use the SDK to create applications and tools in Swift that make use of AWS services.

This guide contains the following sections:

Getting started

Explains how to create a project that imports the SDK for Swift using the Swift Package Manager in a shell environment on Linux and macOS, as well as how to add the SDK package to an Xcode project. Also included is a guide to building a project that uses the SDK to output a list of a user's Amazon S3 buckets.

Configuring service clients

Content describing how to configure service clients to meet your requirements.

Using the SDK

Guides covering typical usage scenarios including creating service clients, performing common tasks, and more.

Code examples

Code examples demonstrating how to use various features of the SDK for Swift, as well as how to achieve specific tasks using the SDK.

Security

Guides covering security topics in general, as well as considerations surrounding using the SDK for Swift in various contexts and while performing specific tasks.

Document history

History of this document.

Maintenance and support for SDK major versions

In general, the AWS SDK for Swift supports Apple's operating systems for two years after their release. Xcode and Swift versions are supported for a year after their release. For more details, see the SDK for Swift's versioning policy in the AWS SDK for Swift API Reference.

For information about maintenance and support for SDK major versions and their underlying dependencies, see the following in the AWS SDKs and Tools Reference Guide.

- AWS SDKs and Tools Maintenance Policy
- AWS SDKs and Tools Version Support Matrix

Additional resources

In addition to this guide, the following are valuable online resources for AWS SDK for Swift developers:

- AWS SDK for Swift API Reference
- AWS SDK for Swift code examples
- AWS SDK for Swift on GitHub

Contributing to the SDK

Developers can also contribute feedback through the following channels:

- Report bugs in the AWS SDK for Swift

Getting started with the AWS SDK for Swift

Learn how to install, set up, and use the SDK to create an application to access an AWS resource programmatically.

Topics

- Setting up the AWS SDK for Swift
- Authenticating with AWS using AWS SDK for Swift
- Creating a simple application using the AWS SDK for Swift

Setting up the AWS SDK for Swift

The AWS SDK for Swift is a cross-platform, open source Swift package that lets applications and tools built using Swift make full use of the services offered by AWS. Using the SDK, you can build Swift applications that work with Amazon S3, Amazon EC2, DynamoDB, and more.

The SDK package is imported into your project using the Swift Package Manager (SPM), which is part of the standard Swift toolchain.

When you've finished following the steps in this article, or have confirmed that everything is configured as described, you're ready to begin developing using the AWS SDK for Swift.

Overview

To make requests to AWS using the AWS SDK for Swift, you need the following:

- An active AWS account.
- A user in IAM Identity Center with permission to use the AWS services and resources your application will access.
- A development environment with version 6.0 or later of the Swift toolchain. If you don't have this, you'll install it as part of the steps below.
- Xcode users need version 16 or later of the Xcode application. This includes Swift 6.0.

After finishing these steps, you're ready to use the SDK to develop Swift projects that access AWS services.

Setting up AWS access

Before installing the Swift tools, configure your environment to let your project access AWS services. This section covers creating and configuring your AWS account, preparing IAM Identity Center for use, and setting the environment variables used by the SDK for Swift to fetch your access credentials.

Sign up for an AWS account

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

Grant programmatic AWS account access

Users need programmatic access if they want to interact with AWS outside of the AWS Management Console. The way to grant programmatic access depends on the type of user that's accessing AWS.

To grant users programmatic access, choose one of the following options.

Which user needs programmatic access?	To	By
IAM	(Recommended) Use console credentials as temporary credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions for the interface that you want to use. • For the AWS CLI, see Login for AWS local development in the *AWS Command Line Interface User Guide*. • For AWS SDKs, see Login for AWS local development in the *AWS SDKs and Tools Reference Guide*.
Workforce identity	Use temporary credentials to sign programmatic requests	Following the instructions for the interface that you want to use.

Which user needs programmatic access?	To	By
(Users managed in IAM Identity Center)	to the AWS CLI, AWS SDKs, or AWS APIs.	• For the AWS CLI, see Configuring the AWS CLI to use AWS IAM Identity Center in the *AWS Command Line Interface User Guide*. • For AWS SDKs, tools, and AWS APIs, see IAM Identity Center authentication in the *AWS SDKs and Tools Reference Guide*.
IAM	Use temporary credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions in Using temporary credentials with AWS resources in the *IAM User Guide*.

Which user needs programmatic access?	To	By
IAM	(Not recommended) Use long-term credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions for the interface that you want to use. • For the AWS CLI, see Authenticating using IAM user credentials in the *AWS Command Line Interface User Guide*. • For AWS SDKs and tools, see Authenticate using long-term credentials in the *AWS SDKs and Tools Reference Guide*. • For AWS APIs, see Managing access keys for IAM users in the *IAM User Guide*.

For more advanced cases regarding configuring the credentials and Region, see The .aws/credentials and .aws/config files, AWS Region, and Using environment variables in the AWS SDKs and Tools Reference Guide.

> **Note**
>
> If you plan to develop a macOS desktop application, keep in mind that due to sandbox restrictions, the SDK is unable to access your `~/.aws/config` and `~/.aws/credentials` files, as there is no entitlement available to grant access to the `~/.aws` directory. See the section called "Security and authentication when testing on macOS" for details.

Setting up your Swift development environment

The SDK requires at least version 6.0 of Swift. This can be installed either standalone or as part of the Xcode development environment on macOS.

- Swift 6.0 toolchain or newer.
- If you're developing on macOS using Xcode, you need a minimum of Xcode 16.
- An AWS account. If you don't have one already, you can create one using the AWS portal.

Prepare to install Swift

This isn't necessary on macOS.

The Swift install process for Linux doesn't automatically install `libcrypto` version 1.1, even though it's required by the compiler. To install it, be sure to install OpenSSL 1.1 or later, as well as its development package. Using the `yum` package manager, for example:

```
$ sudo yum install openssl openssl-devel
```

Using the `apt` package manager:

```
$ sudo apt install openssl libssl-dev
```

Installing Swift

On macOS, the easiest way to install Swift is to simply install Apple's Xcode IDE, which includes Swift and all the standard libraries and tools that go with it. This can be found on the macOS App Store.

If you're using Linux or Windows, or don't want to install Xcode on macOS, the Swift organization's web site has detailed instructions to install and set up the Swift toolchain.

Checking the Swift tools version number

If Swift is already installed, you can verify the version number using the command `swift --version`. The output will look similar to one of the following examples.

Checking the Swift version on macOS

```
$ swift --version
swift-driver version: 1.115 Apple Swift version 6.0 (swiftlang-6.0.0.9.10
 clang-1600.0.26.2)
Target: arm64-apple-macosx15.0
```

Checking the Swift version on Linux

```
$ swift --version
Swift version 6.0 (swift-6.0-RELEASE)
Target: x86_64-unknown-linux-gnu
```

For more information about the configuration and credentials files shared among the AWS Command Line Interface and the various AWS SDKs, see the AWS SDKs and Tools Reference Guide.

Adding the AWS SDK for Swift as a dependency

To add AWS SDK for Swift as a dependency to your Swift package, follow the steps in ???.

To add AWS SDK for Swift as a dependency to your Xcode project, follow the steps in ???.

Security and authentication when testing on macOS

Configuring the App Sandbox

If your SDK for Swift project is a desktop application that you're building in Xcode, you will need to enable the App Sandbox capability and turn on the "Outgoing Connections (Client)" entitlement so that the SDK can communicate with AWS.

First, open the macOS target's **Signing & Capabilities** panel, shown below.

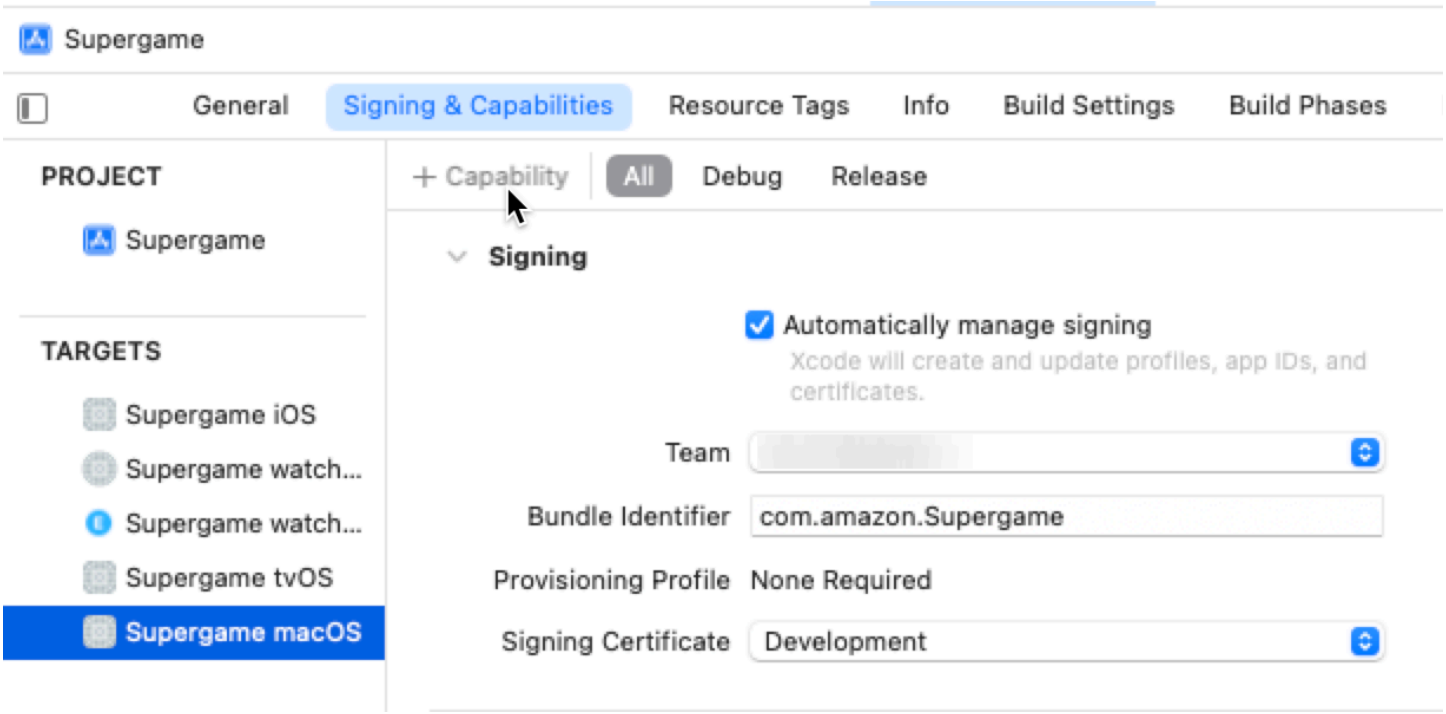

Click the **+ Capability** button near the top left of this panel to bring up the box listing the available capabilities. In this box, locate the "App Sandbox" capability and double-click on it to add it to your target.

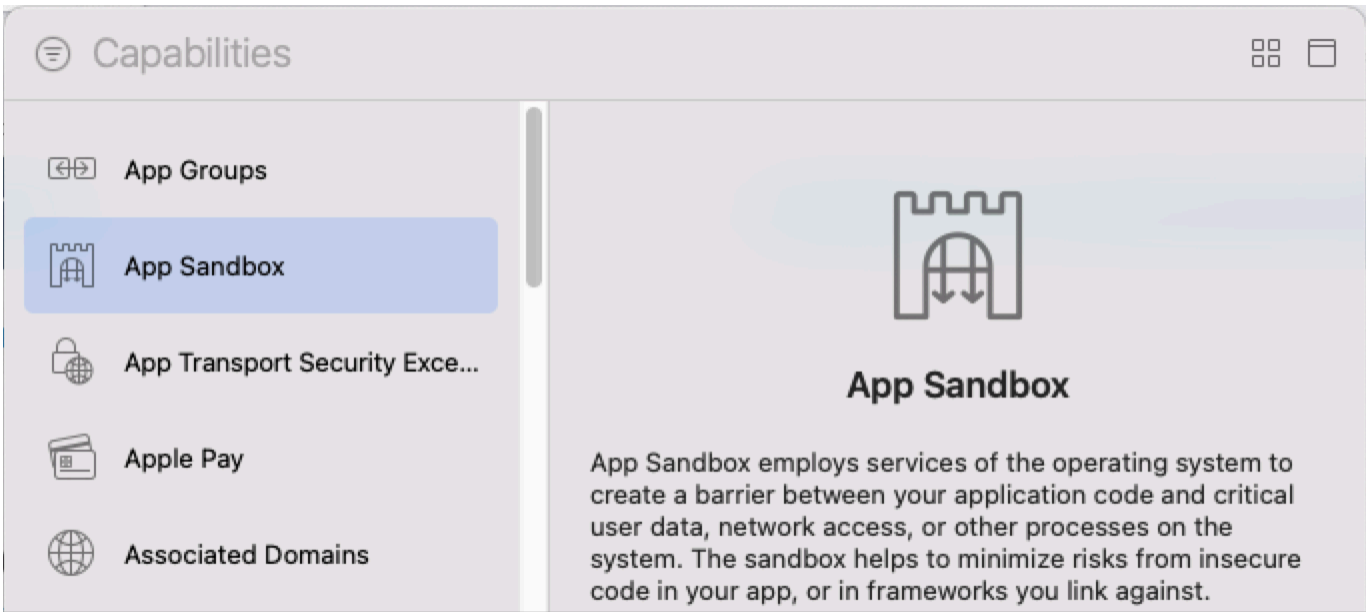

Next, back in your target's **Signing & Capabilities** panel, find the new **App Sandbox** section and make sure that next to **Network**, the **Outgoing Connections (Client)** checkbox is selected in the following image.

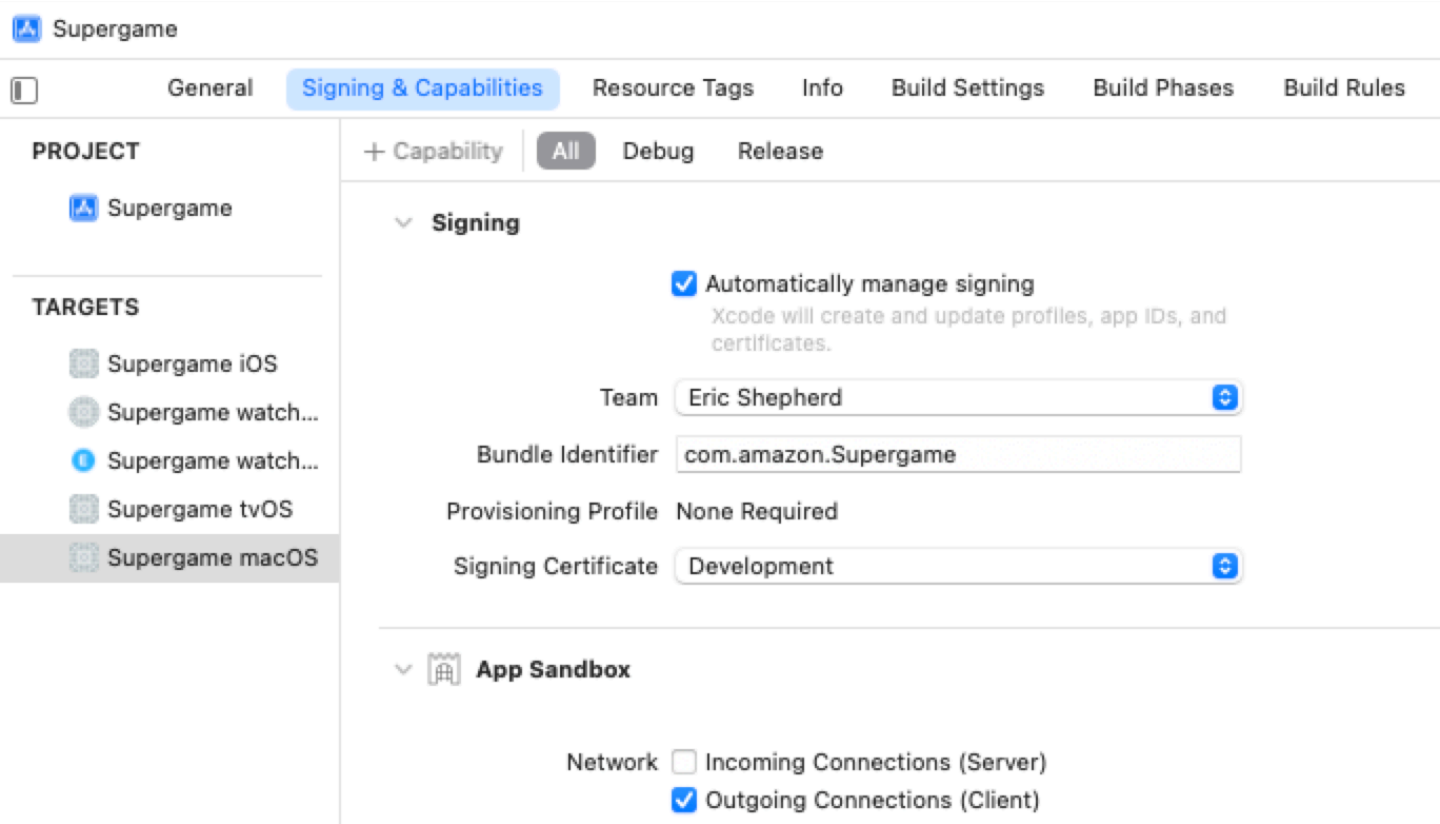

Using AWS access keys on macOS

Although shipping applications should use AWS IAM Identity Center, Amazon Cognito Identity, or similar technologies to handle authentication instead of requiring the direct use of AWS access keys, you may need to use these low-level credentials during development and testing.

macOS security features for desktop applications don't allow applications to access files without express user permission, so the SDK can't automatically configure clients using the contents of the CLI's `~/.aws/config` and `~/.aws/credentials` files. Instead, you need to use the `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY` environment variables to specify the authentication keys for your AWS account.

When running projects from within Xcode, the environment you have set up for your shell is not automatically inherited. Because of this, if you want to test your code using your AWS account's access key ID and secret access key, you need to set up the runtime environment for your application in the Xcode scheme for your development device. To configure the environment for a particular target in your Xcode project, switch to that target, then choose **Edit Scheme** in Xcode's **Product** menu.

This will open the scheme editor window for your project. Click on the **Run** phase in the left sidebar, then **Arguments** in the tab bar near the top of the window.

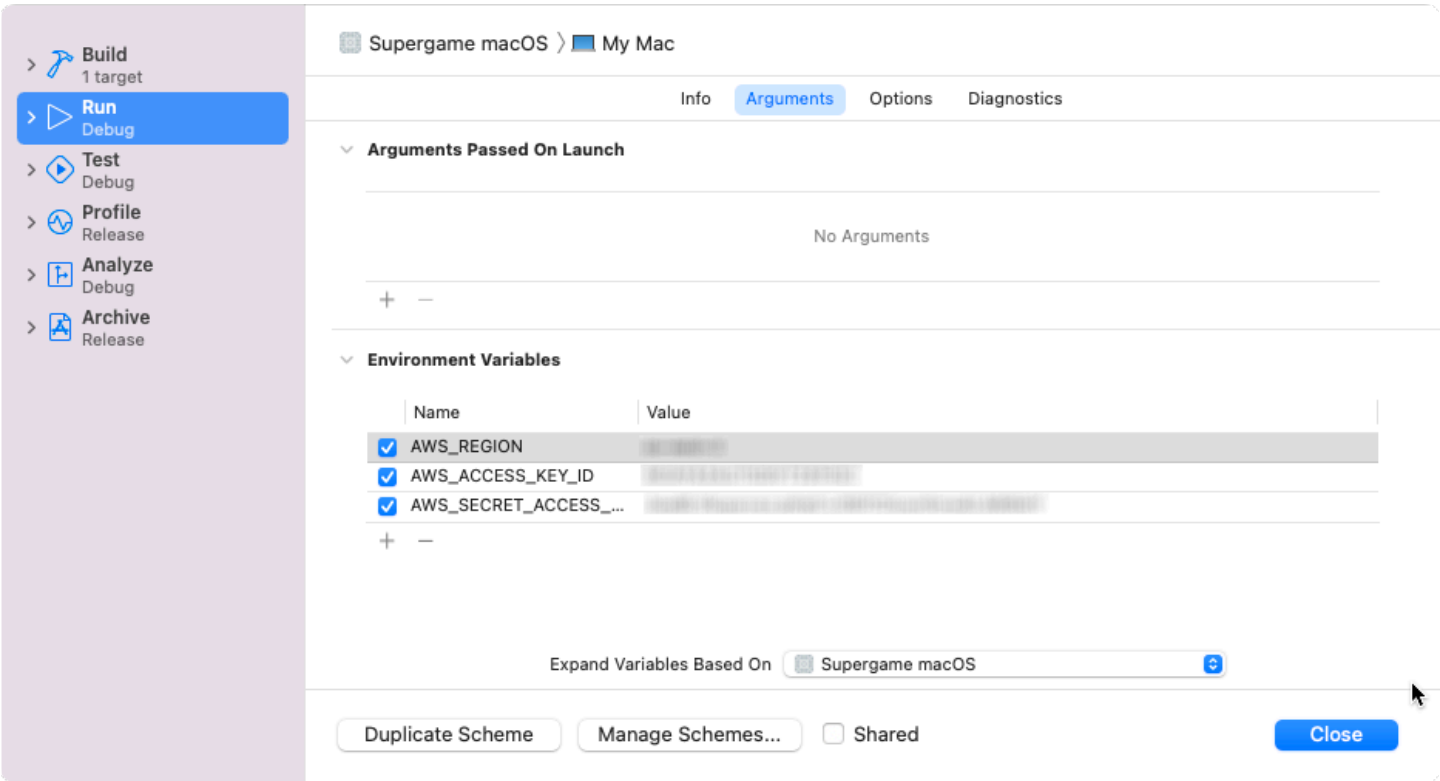

Under **Environment Variables**, click the **+** icon to add `AWS_REGION` and set its value to the desired region (in the screenshot above, it's set to "us-east-2"). Then add `AWS_ACCESS_KEY_ID` and its value, then `AWS_SECRET_ACCESS_KEY` and its value. If you're using temporary credentials (which is recommended), also add the `AWS_SESSION_TOKEN` variable and its value. Close the window once you have these configured and your scheme's Run configuration looks similar to the above.

> **Important**
>
> Be sure to uncheck the **Shared** checkbox before pushing your code to any public version control repository such as GitHub. Otherwise, your AWS access key and secret access key will be *included* in the publicly shared content. This is important enough to be worth double-checking regularly.

Your project should now be able to use the SDK to connect to AWS services.

Building and running a SwiftPM project

To build and run a Swift Package Manager project from a Linux or macOS terminal prompt, use the following commands.

Build a project

```
$ swift build
```

Run a project

```
$ swift run
$ swift run executable-name
$ swift run executable-name arg1, ...
```

If your project builds only one executable file, you can type **swift run** to build and run it. If your project outputs multiple executables, you can specify the file name of the executable you want to run. If you want to pass arguments to the program when you run it, you *must* specify the executable name before listing the arguments.

Get the built product output directory

```
$ swift build --show-bin-path
/home/janice/MyProject/.build/x86_64-unknown-linux-gnu/debug
```

Delete build artifacts

```
$ swift package clean
```

Delete build artifacts and all build caches

```
$ swift package reset
```

Importing AWS SDK for Swift libraries into source files

After the libraries are in place, you can use the Swift `import` directive to import the individual libraries into each source file that needs them. To use the functions for a given service, import

its library from the AWS SDK for Swift package into your source code file. Also import the `ClientRuntime` library, which contains utility functions and type definitions.

```
import Foundation
import ClientRuntime
import AWSS3
```

The standard Swift library `Foundation` is also imported because it's used by many features of the AWS SDK for Swift.

Next steps

Now that your tools and environment are ready for you to begin developing with AWS SDK for Swift, see *Getting started*, which demonstrates how to create and build a Swift project using AWS services.

Authenticating with AWS using AWS SDK for Swift

You must establish how your code authenticates with AWS when developing with AWS services. You can configure programmatic access to AWS resources in different ways depending on the environment and the AWS access available to you. For choices on all primary methods of authentication, and guidance on configuring it for the SDK, see Authentication and access in the *AWS SDKs and Tools Reference Guide*.

Using console credentials

For local development, we recommend that new users use their existing AWS Management Console sign-in credentials for programmatic access to AWS services. After a browser-based authentication flow, AWS generates temporary credentials that work across local development tools like the AWS CLI, AWS Tools for PowerShell and AWS SDKs. This feature simplifies the process of configuring and managing AWS CLI credentials, especially if you prefer interactive authentication over managing long-term access keys.

If you choose this method, follow the instructions to login with console credentials using the AWS CLI. See Login for AWS local development using console credentials for more details.

Once set up with AWS CLI, the default credential provider chain will automatically start using the login token cached by AWS CLI to make requests.

Using IAM Identity Center

This method includes installing the AWS CLI for ease of configuration and for regularly signing in to the AWS access portal.

If you choose this method, complete the procedure for IAM Identity Center authentication in the *AWS SDKs and Tools Reference Guide*. Afterwards, your environment should contain the following elements:

- The AWS CLI, which you use to start an AWS access portal session before you run your application.
- A shared AWSconfig file having a `[default]` profile with a set of configuration values that can be referenced from the SDK. To find the location of this file, see Location of the shared files in the *AWS SDKs and Tools Reference Guide*.
- The shared `config` file sets the `region` setting. This sets the default AWS Region that the SDK uses for AWS requests. This Region is used for SDK service requests that aren't specified with a Region to use.
- The SDK uses the profile's SSO token provider configuration to acquire credentials before sending requests to AWS. The `sso_role_name` value, which is an IAM role connected to an IAM Identity Center permission set, allows access to the AWS services used in your application.

  The following sample `config` file shows a default profile set up with SSO token provider configuration. The profile's `sso_session` setting refers to the named `sso-session` section. The `sso-session` section contains settings to initiate an AWS access portal session.

```
[default]
sso_session = my-sso
sso_account_id = 111122223333
sso_role_name = SampleRole
region = us-east-1
output = json

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

Starting an AWS access portal session

Before running an application that accesses AWS services, you need an active AWS access portal session for the SDK to use IAM Identity Center authentication to resolve credentials. Depending on your configured session lengths, your access will eventually expire and the SDK will encounter an authentication error. To sign in to the AWS access portal, run the following command in the AWS CLI.

```
$ aws sso login
```

If you followed the guidance and have a default profile setup, you do not need to call the command with a `--profile` option. If your SSO token provider configuration is using a named profile, the command is `aws sso login --profile named-profile`.

To optionally test if you already have an active session, you can run the following AWS CLI command.

```
$ aws sts get-caller-identity
```

If your session is active, the response to this command reports the IAM Identity Center account and permission set configured in the shared `config` file.

> **Note**
>
> If you already have an active AWS access portal session and run `aws sso login`, you will not be required to provide credentials.
> The sign-in process might prompt you to allow the AWS CLI access to your data. On the approval screen, the client name begins with `botocore-client-`, named for `botocore`, a Python library that the AWS CLI uses.

More authentication information

Human users, also known as *human identities*, are the people, administrators, developers, operators, and consumers of your applications. They must have an identity to access your AWS environments and applications. Human users that are members of your organization - that means you, the developer - are known as *workforce identities*.

Use temporary credentials when accessing AWS. You can use an identity provider for your human users to provide federated access to AWS accounts by assuming roles, which provide temporary credentials. For centralized access management, we recommend that you use AWS IAM Identity Center (IAM Identity Center) to manage access to your accounts and permissions within those accounts. For more alternatives, see the following:

- To learn more about best practices, see Security best practices in IAM in the *IAM User Guide*.
- To create short-term AWS credentials, see Temporary Security Credentials in the *IAM User Guide*.
- To learn about other credential providers supported by teh AWS SDKs, see Standardized credential providers in the *AWS SDKs and Tools Reference Guide*.

Creating a simple application using the AWS SDK for Swift

This chapter explores how to use the Swift Package Manager — part of the standard Swift toolchain — to create and build a small project. The project uses the AWS SDK for Swift to output a list of available Amazon Simple Storage Service (Amazon S3) buckets.

Creating a project using the AWS SDK for Swift

This chapter demonstrates how to create a small program that lists all the Amazon S3 buckets available on the default user account.

Goals for this project:

- Create a project using Swift Package Manager.
- Add the AWS SDK for Swift to the project.
- Configure the project's `Package.swift` file to describe the project and its dependencies.
- Write code that uses Amazon S3 to get a list of the buckets on the default AWS account, then prints them to the screen.

Before we begin, make sure to prepare your development environment as described in Set up. To make sure you're set up properly, use the following command. This makes sure that Swift is available and which version it is.

```
$ swift --version
```

On macOS, you should see output that looks like the following (with possibly different version and build numbers):

```
swift-driver version: 1.115 Apple Swift version 6.0 (swiftlang-6.0.0.9.10
 clang-1600.0.26.2)
Target: arm64-apple-macosx15.0
```

On Linux, the output should look something like the following:

```
Swift version 6.0 (swift-6.0-RELEASE)
Target: x86_64-unknown-linux-gnu
```

If Swift is not installed, or is older than version 6.0, follow the instructions in ??? to install or reinstall the tools.

> **Get this example on GitHub**
>
> You can fork or download this example from the AWS SDK for Swift code examples repository.

Creating the project

With the Swift tools installed, open a terminal session using your favorite terminal application (such as Terminal, iTerm, or the integrated terminal in your editor).

At the terminal prompt, go to the directory where you want to create the project. Then, enter the following series of commands to create an empty Swift project for a standard executable program.

```
$ mkdir ListBuckets
$ cd ListBuckets
$ swift package init --type executable
$ mv Sources/main.swift Sources/entry.swift
```

This creates the directory for the examples project, moves into that directory, and initializes that directory with the Swift Package Manager (SwiftPM). The result is the basic file system structure for a simple executable program. The Swift source code file `main.swift` is also renamed to `entry.swift` to work around a bug in the Swift tools that involves the use of asynchronous code in the main function of a source file named `main.swift`.

```
ListBuckets/
### Package.swift
### Sources/
   ### entry.swift
```

Open your created project in your preferred text editor or IDE. On macOS, you can open the project in Xcode with the following:

```
$ xed .
```

As another example, you can open the project in Visual Studio Code with the following:

```
$ code .
```

Configuring the package

After opening the project in your editor, open the `Package.swift` file. This is a Swift file that defines a SwiftPM `Package` object that describes the project, its dependencies, and its build rules.

The first line of every `Package.swift` file must be a comment specifying the minimum version of the Swift toolchain needed to build the project. This isn't only informational. The version specified here can change the behavior of the tools for compatibility purposes. The AWS SDK for Swift requires at least version 6.0 of the Swift tools.

```
// swift-tools-version:6.0
```

Specifying supported platforms

For projects whose target operating systems include any Apple platform, add or update the `platforms` property to include a list of the supported Apple platforms and minimum required versions. This list only specifies support for Apple platforms and doesn't preclude building for other platforms.

```
    // Let Xcode know the minimum Apple platforms supported.
    platforms: [
        .macOS(.v11),
        .iOS(.v13)
    ],
```

In this excerpt, the supported Apple platforms are macOS (version 11.0 and higher) and iOS/iPadOS (version 13 and higher).

Setting up dependencies

The package dependency list needs to include the AWS SDK for Swift. This tells the Swift compiler to fetch the SDK and its dependencies before attempting to build the project.

```
    dependencies: [
        // Dependencies declare other packages that this package depends on.
        .package(
            url: "https://github.com/awslabs/aws-sdk-swift",
            from: "1.0.0"
        )
    ],
```

Configuring the target

Now that the package depends on the AWS SDK for Swift, add a dependency to the executable program's target. Indicate that it relies on Amazon S3, which is offered by the SDK's `AWSS3` product.

```
    targets: [
        // Targets are the basic building blocks of a package, defining a module or a
 test suite.
        // Targets can depend on other targets in this package and products from
 dependencies.
        .executableTarget(
            name: "ListBuckets-Simple",
            dependencies: [
                .product(name: "AWSS3", package: "aws-sdk-swift")
            ],
            path: "Sources")
    ]
```

Accessing AWS services using Swift

The example program's Swift code is found in the `Source/entry.swift` file. This file begins by importing the needed Swift modules, using the `import` statement.

```
import AWSClientRuntime
```

```
import AWSS3
import Foundation
```

- `Foundation` is the standard Apple Foundation package.
- `AWSS3` is the SDK module that's used to access Amazon S3.
- `AWSClientRuntime` contains runtime components used by the SDK.

After you add the SDK for Swift to your project and import the service you want to use into your source code, you can create an instance of the client representing the service and use it to issue AWS service requests.

Creating a service client object

Each AWS service is represented by a specific client class in the AWS SDK for Swift. For example, while the Amazon DynamoDB client class is called `DynamoDBClient`, the class for Amazon Simple Storage Service is `S3Client`. To use Amazon S3 in this example, first create an `S3Client` object on which to call the SDK's Amazon S3 functions.

```
        // To set the Region programmatically, pass it to the initializer:
        //   let configuration = try await S3Client.S3ClientConfig(region: "us-east-2")
        let configuration = try await S3Client.S3ClientConfig()
        let client = S3Client(config: configuration)
```

Issuing AWS service requests

To issue a request to an AWS service, call the corresponding function on the service's client object. Each function's inputs are specified using a function-specific input structure as the value of the function's `input` parameter. For example, when calling `S3Client.listBuckets(input:)`, `input` is a structure of type `ListBucketsInput`.

```
        // Use "Paginated" to get all the buckets.
        // This lets the SDK handle the 'continuationToken' in "ListBucketsOutput".
        let pages = client.listBucketsPaginated(
            input: input
        )
```

Functions defined by the AWS SDK for Swift run asynchronously, so the example uses `await` to block the program's execution until the result is available. If SDK functions encounter errors, they

throw them so your code can handle them using a `do-catch` statement or by propagating them back to the caller.

Getting all bucket names

This example's main program calls `getBucketNames()` to get an array containing all of the bucket names. That function is defined as follows.

```
// Return an array containing the names of all available buckets.
//
// - Returns: An array of strings listing the buckets.
func getBucketNames() async throws -> [String] {
    do {
        // Get an S3Client with which to access Amazon S3.
        // To set the Region programmatically, pass it to the initializer:
        //   let configuration = try await S3Client.S3ClientConfig(region: "us-east-2")
        let configuration = try await S3Client.S3ClientConfig()
        let client = S3Client(config: configuration)

        // Use "Paginated" to get all the buckets.
        // This lets the SDK handle the 'continuationToken' in "ListBucketsOutput".
        let pages = client.listBucketsPaginated(
            input: ListBucketsInput( maxBuckets: 10)
        )

        // Get the bucket names.
        var bucketNames: [String] = []

        do {
            for try await page in pages {
                guard let buckets = page.buckets else {
                    print("Error: no buckets returned.")
                    continue
                }

                for bucket in buckets {
                    bucketNames.append(bucket.name ?? "<unknown>")
                }
            }

            return bucketNames
        } catch {
            print("ERROR: listBuckets:", dump(error))
```

```
            throw error
        }
    }
}
```

This function starts by creating an Amazon S3 client and calling its `listBuckets(input:)` function to request a list of all of the available buckets. The list is returned asynchronously. After it's returned, the bucket list is fetched from the output structure's `buckets` property. If it's `nil`, an empty array is immediately returned to the caller. Otherwise, each bucket name is added to the array of bucket name strings, which is then returned to the caller.

Adding the example entry point

To allow the use of asynchronous functions from within `main()`, use Swift's `@main` attribute to create an object that contains a static `async` function called `main()`. Swift will use this as the program entry point.

```
@main
struct Main {
    static func main() async {
        do {
            let names = try await getBucketNames()

            print("Found \(names.count) buckets:")
            for name in names {
                print("  \(name)")
            }
        } catch let error as AWSServiceError {
            print("An Amazon S3 service error occurred: \(error.message ?? "No details
 available")")
        } catch {
            print("An unknown error occurred: \(dump(error))")
        }
    }
}
```

`main()` calls `getBucketNames()`, then outputs the returned list of names. Errors thrown by `getBucketNames()` are caught and handled. Errors of type `ServiceError`, which represent errors reported by the AWS service, are handled specially.

Additional information

- the section called “Setting up”
- *Using the SDK*
- SDK for Swift code examples

Configuring service clients in the AWS SDK for Swift

To programmatically access AWS services, SDKs use a client class/object for each AWS 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 AWS service.

To make a request to an AWS service, you must first create a service client. For each AWS service your code uses, it has its own crate/library/gem/package and its own dedicated type for interacting with it. The client exposes one method for each API operation exposed by the service.

There are many alternative ways to configure SDK behavior, but ultimately everything has to do with the behavior of service clients. Any configuration has no effect until a service client that is created from them is used.

You must establish how your code authenticates with AWS when you develop with AWS services. You must also set the AWS Region you want to use.

The AWS SDKs and Tools Reference Guide also contains settings, features, and other foundational concepts common among many of the AWS SDKs.

Topics

- Configuring AWS SDK for Swift service clients externally
- Configuring AWS SDK for Swift service clients in code
- Setting the AWS Region for the AWS SDK for Swift
- Using AWS SDK for Swift credential providers
- Configuring retry using the AWS SDK for Swift
- Configuring HTTP-level settings within the AWS SDK for Swift

Configuring AWS SDK for Swift service clients externally

Many configuration settings can be handled outside of your code. When configuration is handled externally, the configuration is applied across all of your applications. Most configuration settings can be set as either environment variables or in a separate shared AWS `config` file. The shared `config` file can maintain separate sets of settings, called profiles, to provide different configurations for different environments or tests.

Environment variables and shared `config` file settings are standardized and shared across AWS SDKs and tools to support consistent functionality across different programming languages and applications.

See the *AWS SDKs and Tools Reference Guide* to learn about configuring your application through these methods, plus details on each cross-sdk setting. To see all the all settings that the SDK can resolve from the environment variables or configuration files, see the Settings reference in the *AWS SDKs and Tools Reference Guide*.

To make a request to an AWS service, you first instantiate a client for that service. You can configure common settings for service clients such as timeouts, the HTTP client, and retry configuration.

Each service client requires an AWS Region and a credential provider. The SDK uses these values to send requests to the correct Region for your resources and to sign requests with the correct credentials. You can specify these values programmatically in code or have them automatically loaded from the environment.

The SDK has a series of places (or sources) that it checks in order to find a value for configuration settings.

1. Any explicit setting set in the code or on a service client itself takes precedence over anything else.
2. Environment variables
   - For details on setting environment variables, see environment variables in the *AWS SDKs and Tools Reference Guide*.
   - Note that you can configure environment variables for a shell at different levels of scope: system-wide, user-wide, and for a specific terminal session.
3. Shared `config` and `credentials` files
   - For details on setting up these files, see the Shared `config` and `credentials` files in the *AWS SDKs and Tools Reference Guide*.
4. Optionally, the configuration is looked for on Amazon Elastic Container Service.
5. Optionally, the configuration can be taken from Amazon EC2 instance metadata.
6. Any default value provided by the SDK source code itself is used last.
   - Some properties, such as Region, don't have a default. You must specify them either explicitly in code, in an environment setting, or in the shared `config` file. If the SDK can't resolve required configuration, API requests can fail at runtime.

> **Note**
>
> The AWS SDK for Swift's default credential resolver chain does not include the SSO credential resolver. You must explicitly configure service instances to use it if desired.

Configuring AWS SDK for Swift service clients in code

When configuration is handled directly in code, the configuration scope is limited to the application that uses that code. Within that application, there are options for the global configuration of all service clients, the configuration to all clients of a certain AWS service type, or the configuration to a specific service client instance.

Configuration data types

The AWS SDK for Swift provides a configuration structure for each AWS service. Use it to control how a client connects to and interacts with that service. Customize a client's configuration to control the reliability, security, and performance of your application's connection. You can set the Region it reaches, the credentials it presents, and how it retries failed calls. For example, you might raise the maximum number of retry attempts for an application that runs on an unreliable network.

Each AWS service has its own configuration structure that you use to adjust options for that client type.

All of these structures are based on a number of protocols that together describe all the configuration options available for that client. This includes options such as the Region and credentials, which you always need to access an AWS service.

Amazon Simple Storage Service (Amazon S3) uses the S3Client.S3ClientConfig structure, described in the *AWS SDK for Swift API Reference*. The AWS SDK for Swift defines it like this:

```
extension S3Client {
    public struct S3ClientConfig: AWSClientRuntime.AWSDefaultClientConfiguration &
            AWSClientRuntime.AWSRegionClientConfiguration &
            ClientRuntime.DefaultClientConfiguration &
            ClientRuntime.DefaultHttpClientConfiguration, Swift.Sendable {

        // Service-specific properties are defined here, including:
```

```
        public var region: Swift.String?

        // ... and so on.
    }
}
```

As a result, `S3ClientConfig` includes its own properties and also all the properties defined in those protocols, making it a complete description of an Amazon S3 client's configuration. Every AWS service defines a corresponding configuration structure.

Because the configuration is a `struct`, it's a value type. Declare it using `var` if you need to change its properties after creating it. Remember that assigning it to another variable makes a copy rather than sharing a reference.

> **Note**
>
> Earlier versions of the AWS SDK for Swift used a configuration class named `S3Client.S3ClientConfiguration`, with similarly named classes for the other AWS services. We deprecated those classes and plan to remove them in a future release. Use the `S3ClientConfig` structure, or the equivalent structure for your service, instead.

By creating a custom configuration object and using it when creating a service client object, you can customize the client by specifying a configuration source from which credentials and other options are taken. Otherwise, you can directly specify the credentials instead of letting the SDK obtain them automatically.

Configuring a client

When only changing common options, you can often specify the custom values for your configuration when instantiating the service’s client object. If the client class constructor doesn't support the option you need to change, then create and specify a configuration object with the type that corresponds to the client class.

Creating a client with only a custom Region

Most services let you directly specify the Region when you call their constructors. For example, to create an Amazon S3 client configured for the Region `af-south-1`, specify the `region` parameter when creating the client, as shown.

```
do {
    let s3 = try S3Client(region: "af-south-1")

    // Use the client.
} catch {
    // Handle the error.
    dump(error, name: "Error accessing S3 service")
}
```

This lets you handle a common client configuration scenario (specifying a Region while using the default values for all other options) without going through the full configuration process. That process is covered in the next topic.

Creating and using a custom configuration

To customize the configuration of an AWS service, create a configuration object of the appropriate type for the service. Then pass that configuration object into the service client's constructor as the value of its `config` parameter.

For example, to configure an Amazon S3 client, create a value of type S3Client.S3ClientConfig. Set the properties to change, and then pass the configuration into S3Client(config:). The *AWS SDK for Swift API Reference* describes both.

The following example creates a new Amazon S3 client configured with the following options:

- The AWS Region is set to `us-east-1`.
- An exponential backoff strategy with default options is selected instead of the default backoff strategy.
- The retry mode is set to `RetryStrategyOptions.RateLimitingMode.adaptive`. See Retry behavior in the *AWS SDKs and Tools Reference Guide* for details.
- The maximum number of retries is set to 5.

```
        // Create an Amazon S3 client configuration object that specifies the
        // region as "us-east-1", an exponential backoff strategy, the
        // adaptive retry mode, and the maximum number of retries as 5.

        await SDKLoggingSystem().initialize(logLevel: .debug)
```

```
        let config: S3Client.S3ClientConfig

        do {
            config = try await S3Client.S3ClientConfig(
                awsRetryMode: .standard,
                maxAttempts: 3,
                region: "us-east-1"
            )
        } catch {
            print("Error: Unable to create configuration")
            dump(error)
            exit(1)
        }

        // Create an Amazon S3 client using the configuration created above.

        let client = S3Client(config: config)
```

If an error occurs creating the configuration, an error message is displayed and the details are dumped to the console. The program then exits. In a real world application, a more constructive approach should be taken when handling this error, such as falling back to a different configuration or using the default configuration, if doing so is appropriate for your application.

Setting the AWS Region for the AWS SDK for Swift

You can access AWS services that operate in a specific geographic area by using AWS Regions. This can be useful both for redundancy and to keep your data and applications running close to where you and your users access them.

> **Important**
>
> Most resources reside in a specific AWS Region and you must supply the correct Region for the resource when using the SDK.

You must set a default AWS Region for the SDK for Swift to use for AWS requests. This default is used for any SDK service method calls that aren't specified with a Region.

For examples on how to set the default region through the shared AWS `config` file or environment variables, see AWS Region in the *AWS SDKs and Tools Reference Guide*.

Specifying the Region programmatically

Most services let you directly specify the Region when you call their constructors. For example, to create an Amazon S3 client configured for the Region `af-south-1`, specify the `region` parameter when creating the client, as shown.

```
do {
    let s3 = try S3Client(region: "af-south-1")

    // Use the client.
} catch {
    // Handle the error.
    dump(error, name: "Error accessing S3 service")
}
```

This lets you handle a common client configuration scenario (specifying a Region while using the default values for all other options) without going through the full configuration process. That process is covered in the section called "Creating and using a custom configuration".

Using AWS SDK for Swift credential providers

All requests to AWS must be cryptographically signed by using credentials issued by AWS. At runtime, the SDK retrieves configuration values for credentials by checking several locations.

If the retrieved configuration includes AWS IAM Identity Center single sign-on access settings, the SDK works with the IAM Identity Center to retrieve temporary credentials that it uses to make request to AWS services.

If the retrieved configuration includes temporary credentials, the SDK uses them to make AWS service calls. Temporary credentials consist of access keys and a session token.

Authentication with AWS can be handled outside of your codebase. Many authentication methods can be automatically detected, used, and refreshed by the SDK using the credential provider chain.

For guided options for getting started on AWS authentication for your project, see Authentication and access in the *AWS SDKs and Tools Reference Guide*.

The credential provider chain

If you don't explicitly specify a credential provider when constructing a client, the AWS SDK for Swift uses a credential provider chain that checks a series of places where you can supply credentials. Once the SDK finds credentials in one of these locations, the search stops.

Credential retrieval order

The credential provider chain searches for credentials using the following predefined sequence:

1. **Access key environment variables**

   The SDK attempts to load credentials from the `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY`, and `AWS_SESSION_TOKEN` environment variables.

2. **The shared AWS `config` and `credentials` files**

   The SDK attempts to load credentials from the `[default]` profile in the shared AWS `config` and `credentials` files. You can use the `AWS_PROFILE` environment variable to choose a named profile you want the SDK to load instead of using `[default]`. The `config` and `credentials` files are shared by various AWS SDKs and tools. For more information on these files, see the Shared `config` and `credentials` files in the *AWS SDKs and Tools Reference Guide*.

   If you use console credentials to authenticate, this is when the SDK uses the login token that was set up by running AWS CLI command `aws login`. The SDK uses the temporary credentials from the cached token when it calls AWS services. For detailed information about this process, see Understand SDK credential resolution for AWS services in the *AWS SDKs and Tools Reference Guide*.

   If you use IAM Identity Center to authenticate, this is when the SDK uses the single sign-on token that was set up by running AWS CLI command `aws sso login`. The SDK uses the temporary credentials that the IAM Identity Center exchanged for a valid token. The SDK then uses the temporary credentials when it calls AWS services. For detailed information about this process, see Understand SDK credential resolution for AWS services in the *AWS SDKs and Tools Reference Guide*.

   - For guidance on configuring this provider, see IAM Identity Center authentication in the *AWS SDKs and Tools Reference Guide*.
   - For details on SDK configuration properties for this provider, see IAM Identity Center credential provider in the *AWS SDKs and Tools Reference Guide*.

3. **AWS STS web identity**

   When creating mobile applications or client-based web applications that require access to AWS, AWS Security Token Service (AWS STS) returns a set of temporary security credentials for federated users who are authenticated through a public identity provider (IdP).

   - When you specify this in a profile, the SDK or tool attempts to retrieve temporary credentials using AWS STS `AssumeRoleWithWebIdentity` API method. For details on this method, see AssumeRoleWithWebIdentity in the *AWS Security Token Service API Reference*.
   - For guidance on configuring this provider, see Federate with web identity or OpenID Connect in the *AWS SDKs and Tools Reference Guide*.
   - For details on SDK configuration properties for this provider, see Assume role credential provider in the *AWS SDKs and Tools Reference Guide*.

4. **Amazon ECS and Amazon EKS container credentials**

   Your Amazon Elastic Container Service tasks and Kubernetes service accounts can have an IAM role associated with them. The permissions granted in the IAM role are assumed by the containers running in the task or containers of the pod. This role allows your application code (on the container) to use other AWS services.

   The SDK attempts to retrieve credentials from the `AWS_CONTAINER_CREDENTIALS_RELATIVE_URI` or `AWS_CONTAINER_CREDENTIALS_FULL_URI` environment variables, which can be set automatically by Amazon ECS and Amazon EKS.

   - For details on setting up this role for Amazon ECS, see Amazon ECS task IAM role in the *Amazon Elastic Container Service Developer Guide*.
   - For Amazon EKS setup information, see Setting up the Amazon EKS Pod Identity Agent in the Amazon EKS User Guide.
   - For details on SDK configuration properties for this provider, see Container credential provider in the *AWS SDKs and Tools Reference Guide*.

5. **Amazon EC2 Instance Metadata Service**

   Create an IAM role and attach it to your instance. The SDK application on the instance attempts to retrieve the credentials provided by the role from the instance metadata.

   - For details on setting up this role and using metadata, IAM roles for Amazon EC2 and Work with instance metadata in the *Amazon EC2 User Guide*.

- For details on SDK configuration properties for this provider, see IMDS credential provider in the *AWS SDKs and Tools Reference Guide*.

For details on AWS credential provider configuration settings, see Standardized credential providers in the *Settings reference* of the *AWS SDKs and Tools Reference Guide*.

Credential identity resolvers

A **credential identity resolver** is an object that takes some form of identity, verifies that it's valid for use by the application, and returns credentials that can be used when using an AWS service . There are several supported ways to obtain a valid identity, and each has a corresponding credential identity resolver type available for you to use, depending on which authorization methods you want to use.

The credential identity resolver acts as an adaptor between the identity and the AWS service. By providing a credential identity resolver to the service instead of directly providing the user's credentials, the service is able to fetch currently-valid credentials for the identity at any time, as long as the identity provider allows it.

Identity features in the AWS SDK for Swift are defined in the module `AWSSDKIdentity`. In the `AWSSDKIdentity` module, credentials are represented by the struct `AWSCredentialIdentity`. See AWS security credentials in the IAM User Guide for further information about AWS credentials.

There are several credential identity resolver types available as a means of obtaining an identity to use for authentication. Some credential identity resolvers are specific to a given source while others encompass an assortment of identity sources that share similar technologies. For example, the `STSWebIdentityCredentialIdentityResolver`, which uses a JSON Web Token (JWT) as the source identity for which to return AWS credentials. The JWT can come from a number of different services, including Amazon Cognito federated identities, Sign In With Apple, Google, or Facebook. See Identity pools third-party identity providers for information on third-party identity providers.

`CachedAWSCredentialIdentityResolver`

A credential identity resolver that is chained with another one so it can cache the resolved identity for re-use until an expiration time elapses.

`CognitoAWSCredentialIdentityResolver`

A credential identity resolver that uses an Amazon Cognito Identity Pool or a Cognito Identity ID to resolve credentials.

`CustomAWSCredentialIdentityResolver`

A credential identity resolver that uses another credential identity resolver's output to resolve the credentials in a custom way.

`DefaultAWSCredentialIdentityResolverChain`

Represents a chain of credential identity resolvers that attempt to resolve the identity following the standard search order. See Credential provider chain in the AWS SDKs and Tools Reference Guide for details on the credential provider chain.

`ECSAWSCredentialIdentityResolver`

Obtains credentials from an Amazon Elastic Container Service container's metadata.

`EnvironmentAWSCredentialIdentityResolver`

Resolves credentials using the environment variables `AWS_ACCESS_KEY_ID`, `AWS_SECRET_ACCESS_KEY`, and `AWS_SESSION_TOKEN`.

`IMDSAWSCredentialIdentityResolver`

Uses IMDSv2 to fetch credentials within an Amazon Elastic Compute Cloud instance.

`LoginAWSCredentialIdentityResolver`

Resolves credentials using console credentials with AWS Signin.

`ProcessAWSCredentialIdentityResolver`

Resolves credentials by running a command or process. The process to run is sourced from a profile in the AWS `config` file. The profile key that identifies the process to use is `credential_process`.

`ProfileAWSCredentialIdentityResolver`

Uses the specified profile from an AWS `config` file to resolve credentials.

`SSOAWSCredentialIdentityResolver`

Resolves credentials using a single-sign-on login with AWS IAM Identity Center.

`StaticAWSCredentialIdentityResolver`

A credential resolver that uses specified credentials in the form of an `AWSCredentialIdentity` object.

`STSAssumeRoleAWSCredentialIdentityResolver`

Uses another credential identity resolver to assume a specified AWS Identity and Access Management role, then fetch the assumed credentials using AWS Security Token Service.

`STSWebIdentityAWSCredentialIdentityResolver`

Exchanges a JSON Web Token (JWT) for credentials using AWS Security Token Service.

Getting credentials from an identity

The process of using a credential identity resolver involves four primary steps:

1. Use an appropriate sign-in service to obtain an identity in a form supported by AWS.
2. Create a credential identity resolver of the type that corresponds to the given identity.
3. When creating an AWS service client object, provide the credential identity resolver as the value of its configuration's `awsCredentialIdentityResolver` property.
4. Call service functions using the service client object.

The following sections provide examples using some of the credential identity providers supported by AWS.

Login credential identity resolver with AWS Signin

Authenticating for an AWS service using console credentials requires first creating login token with AWS Command Line Interface. See Login for AWS local development using console credentials for more details.

Once the login token is created and cached by AWS CLI command **aws login** or **aws login --profile *profile-name***, your application can use `LoginAWSCredentialIdentityResolver` to obtain credentials using the established login token.

To create a Login credential identity resolver, create a new `LoginAWSCredentialIdentityResolver` that uses the desired settings for the profile name, `config` file path, and `credentials` file path. Any of these can be `nil` to use the same default value the AWS CLI would use.

> **Note**
>
> To use credential identity resolvers, you must import the `AWSSDKIdentity` module:
>
> ```
> import AWSSDKIdentity
> ```

```
let identityResolver = try LoginAWSCredentialIdentityResolver(
    profileName: profile,
    configFilePath: config,
    credentialsFilePath: credentials
)
```

To use this resolver to provide credentials to an AWS service, set the service configuration's `awsCredentialIdentityResolver` to the created credential identity resolver.

```
// Get an S3Client with which to access Amazon S3.
let configuration = try await S3Client.S3ClientConfig(
    awsCredentialIdentityResolver: identityResolver
)
let client = S3Client(config: configuration)

// Use "Paginated" to get all the buckets. This lets the SDK handle
// the 'continuationToken' in "ListBucketsOutput".
let pages = client.listBucketsPaginated(
    input: ListBucketsInput(maxBuckets: 10)
)
```

SSO credential identity resolvers with AWS IAM Identity Center

Authenticating for an AWS service using SSO requires first configuring SSO access using AWS IAM Identity Center. See IAM Identity Center authentication for your SDK or tool in the AWS SDKs and Tools Reference Guide for instructions on setting up IAM Identity Center and configuring SSO access on computers that will use your application.

Once a user has authenticated with the AWS Command Line Interface (AWS CLI) command **aws sso login** or **aws sso login --profile *profile-name***, your application can use an `SSOAWSCredentialIdentityResolver` to obtain credentials using the established IAM Identity Center identity.

To create an SSO credential identity resolver, create a new `SSOAWSCredentialIdentityResolver` that uses the desired settings for the profile name, `config` file path, and `credentials` file path. Any of these can be `nil` to use the same default value the AWS CLI would use.

> **Note**
>
> To use credential identity resolvers, you must import the `AWSSDKIdentity` module:
>
> ```
> import AWSSDKIdentity
> ```

```
        let identityResolver = try SSOAWSCredentialIdentityResolver(
            profileName: profile,
            configFilePath: config,
            credentialsFilePath: credentials
        )
```

To use the IAM Identity Center identity resolver to provide credentials to an AWS service, set the service configuration's `awsCredentialIdentityResolver` to the created credential identity resolver.

```
    // Get an S3Client with which to access Amazon S3.
    let configuration = try await S3Client.S3ClientConfig(
        awsCredentialIdentityResolver: identityResolver
    )
    let client = S3Client(config: configuration)

    // Use "Paginated" to get all the buckets. This lets the SDK handle
    // the 'continuationToken' in "ListBucketsOutput".
    let pages = client.listBucketsPaginated(
        input: ListBucketsInput(maxBuckets: 10)
    )
```

With the service configured this way, each time the SDK accesses the AWS service, it uses the credentials returned by the SSO credential identity resolver to authenticate the request.

Static credential identity resolvers

> **Warning**
>
> Static credential identity resolvers are highly unsafe unless used with care. They can return hard-coded credentials, which are inherently unsafe to use. Only use static credential identity resolvers when experimenting, testing, or generating safe static credentials from another source before using them.

Static credential identity resolvers use AWS credentials as an identity. To create a static credential identity resolver, create an `AWSCredentialIdentity` object with the static credentials, then create a new `StaticAWSCredentialIdentityResolver` that uses that identity.

> **Note**
>
> To use credential identity resolvers, you must import the `AWSSDKIdentity` module:
>
> ```
> import AWSSDKIdentity
> ```

```
        let credentials = AWSCredentialIdentity(
            accessKey: accessKey,
            secret: secretKey,
            sessionToken: sessionToken
        )

        let identityResolver = try StaticAWSCredentialIdentityResolver(credentials)
```

To use the static credential identity resolver to provide credentials to an AWS service, use it as the service configuration's `awsCredentialIdentityResolver`.:

```
            let s3Configuration = try await S3Client.S3ClientConfig(
                awsCredentialIdentityResolver: identityResolver,
                region: region
            )
            let client = S3Client(config: s3Configuration)
```

```
            // Use "Paginated" to get all the buckets. This lets the SDK handle
            // the 'continuationToken' in "ListBucketsOutput".
            let pages = client.listBucketsPaginated(
                input: ListBucketsInput( maxBuckets: 10)
            )
```

When the service client asks the credential identity resolver for its credentials, the resolver returns the `AWSCredentialIdentity` struct's access key, secret, and session token.

The complete example is available on GitHub.

Amazon Cognito credential identity resolvers

To use Amazon Cognito to provide credentials to an AWS service, first create an object of type `CognitoAWSCredentialIdentityResolver`, specifying the Amazon Cognito Identity Pool ID of the identity pool to use for credentials:

```
        // Create a Cognito credential resolver that uses the Cognito Identity
        // Pool.
        let cognitoCredentialResolver = try CognitoAWSCredentialIdentityResolver(
            identityPoolId: identityPoolID,
            identityPoolRegion: region
        )
```

Once the credential identity resolver has been created, you can instruct service clients to use it to resolve credentials by including the `awsCredentialIdentityResolver` property in the client's configuration with its value set to the `CognitoAWSCredentialIdentityResolver`.

```
        let s3Config = try await S3Client.S3ClientConfig(
            awsCredentialIdentityResolver: cognitoCredentialResolver,
            region: region
        )
        let s3Client = S3Client(config: s3Config)

        let listBucketsOutput = try await s3Client.listBuckets(
            input: ListBucketsInput()
        )
```

This example creates an Amazon S3 client that uses the Amazon Cognito credential identity resolver to authenticate when using the S3 client. All AWS actions performed

using the client will authenticate using the resolved credentials, as shown here by calling `S3Client.listBuckets(input:)`.

Additional information

- AWS SDKs and Tools Reference Guide: SSO token provider configuration

Configuring retry using the AWS SDK for Swift

Calls to AWS services occasionally encounter problems or unexpected situations. Certain types of errors, such as throttling or transient errors, might be successful if the call is retried.

This page describes how to configure automatic retries with the AWS SDK for Swift.

Default retry configuration

By default, every service client is automatically configured with a standard retry strategy. The default configuration tries each action up to three times (the initial attempt plus two retries). The intervening delay between each call is configured with exponential backoff and random jitter to avoid retry storms. This configuration works for the majority of use cases but may be unsuitable in some circumstances, such as high-throughput systems.

The SDK attempts retries only on retryable errors. Examples of retryable errors are socket timeouts, service-side throttling, concurrency or optimistic lock failures, and transient service errors. Missing or invalid parameters, authentication/security errors, and misconfiguration exceptions are not considered retryable.

Configuring retry

You can customize the standard retry strategy by setting the maximum number of attempts and the rate limiting strategy to use. To change the retry configuration, create a structure of type `S3ClientConfig` with the properties you wish to customize. For details on how to configure a service client, see Configuring AWS SDK for Swift service clients in code.

Maximum number of attempts

You can customize the maximum number of attempts by specifying a value for the `S3ClientConfig` structure's `maxAttempts` property. The default value is 3.

Retry mode

To set the retry mode, change the `S3ClientConfig` structure's `awsRetryMode` property. The enum `AWSRetryMode` provides the available values.

Legacy

The `legacy` retry mode—the default—is the original standard retry mode. In this mode, requests may be sent immediately, and are not delayed for rate limiting when throttling is detected. Requests are only delayed according to the backoff strategy in use, which is exponentially by default.

Standard

In the AWS SDK for Swift, the `standard` mode is the same as legacy mode.

Adaptive

In the `adaptive` retry mode, initial and retry requests may be delayed by an additional amount when throttling is detected. This is intended to reduce congestion when throttling is in effect. This is sometimes called "client-side rate limiting" mode, and is available opt-in. You should only use adaptive mode when advised to do so by an AWS representative.

Example

First, import the modules needed to configure retry:

```
import AWSS3
import SmithyRetries
import SmithyRetriesAPI
```

Then create the custom configuration. This example's `S3ClientConfig` asks for the client to make up to three attempts for each action, using the `adaptive` retry mode.

```
        let config: S3Client.S3ClientConfig

        // Create an Amazon S3 client configuration object that specifies the
        // adaptive retry mode and sets the maximum number of attempts to 3.
        // If that fails, create a default configuration instead.

        do {
```

```
            config = try await S3Client.S3ClientConfig(
                awsRetryMode: .adaptive,
                maxAttempts: 3
            )
        } catch {
            do {
                config = try await S3Client.S3ClientConfig()
            } catch {
                print("Error: Unable to configure Amazon S3.")
                dump(error)
                return
            }
        }

        // Create an Amazon S3 client using the configuration created above.

        let client = S3Client(config: config)
```

It first attempts to create the configuration using the application's preferred settings. If that fails, a default configuration is created instead. If that also fails, the example outputs an error message. A real-world application might instead present an error message and let the user decide what to do.

Configuring HTTP-level settings within the AWS SDK for Swift

Configuring HTTP timeouts

The AWS SDK for Swift supports two types of timeout on HTTP clients. These are configured by setting the corresponding properties on the service client's `httpClientConfiguration` structure to the number of seconds to allow before timing out:

- **`connectTimeout`** specifies how much time to allow before an attempt to open an HTTP connection times out.
- **`socketTimeout`** specifies how long to wait for a socket to respond to a request before timing out.

```
        do {
            let config = try await S3Client.S3ClientConfig(
                region: region,
                httpClientConfiguration: HttpClientConfiguration(
                    connectTimeout: 2,
```

```
                    socketTimeout: 5
                )
            )
            let s3Client = S3Client(config: config)
            _ = try await s3Client.listBuckets(input: ListBucketsInput())
            print("*** Success!")
        } catch CommonRunTimeError.crtError(let crtError) {
            print("*** An error occurred accessing the bucket list:
 \(crtError.message)")
        } catch {
            print("*** Unexpected error occurred requesting the bucket list.")
        }
```

This example creates an Amazon S3 client that times out HTTP connection attempts after 2 seconds, while allowing up to 5 seconds for sockets to time out.

Customizing HTTP headers

You can add default headers to your HTTP requests by setting the service client's `httpConfiguration.defaultHeaders` property to a `Headers` structure header which is initialized with a dictionary that maps header names to their values.

> **Important**
>
> Setting default headers that interfere with headers the AWS service expects may result in unpredictable results.

This example creates an Amazon S3 client whose HTTP configuration adds two custom headers to HTTP requests sent by the client:

```
import ClientRuntime
import AWSS3
import SmithyHTTPAPI
import AwsCommonRuntimeKit

        let config = try await S3Client.S3ClientConfig(
            region: region,
            httpClientConfiguration: HttpClientConfiguration(
                defaultHeaders: Headers(
                    [
```

```
                        "X-My-Custom-Header": "CustomHeaderValue",
                        "X-Another-Custom-Header": "AnotherCustomValue"
                    ]
                )
            )
        )
        let s3Client = S3Client(config: config)
```

HTTP/1 vs. HTTP/2

HTTP version 1.1 is adequate for most tasks when using AWS services, and is used automatically for these. However, there are times when HTTP version 2 is required. While some AWS SDKs require you to specifically customize the service client to use HTTP/2, the SDK for Swift detects when a function uses event streaming and automatically uses HTTP/2 for these actions. Your code doesn't need to do anything to handle this.

Using a custom HTTP client

To use a custom HTTP client, create a client class which conforms to the `HTTPClient` class in the `SmithyHTTPAPI` module. This protocol requires a single function: `send(request:)`, which returns a `HTTPResponse` object. You can then specify your custom HTTP client by setting the client configuration's `httpClientEngine` property to an instance of your custom HTTP client class.

Using the SDK for Swift

After completing the steps in the section called "Setting up", you're ready to make requests to AWS services such as Amazon S3, DynamoDB, IAM, Amazon EC2, and more by instantiating service objects and making calls on them using the AWS SDK for Swift. The guides below cover setting up and using AWS services for common use cases.

Topics

- Making AWS service requests using the AWS SDK for Swift
- Handling errors in the AWS SDK for Swift
- Using paginated results in the AWS SDK for Swift
- Testing and debugging with the AWS SDK for Swift
- Using waiters with the AWS SDK for Swift
- Creating AWS Lambda functions using the AWS SDK for Swift
- Event streaming with AWS SDK for Swift
- Presigned URLs and requests with AWS SDK for Swift
- Integrating with Apple platforms using the AWS SDK for Swift

Making AWS service requests using the AWS SDK for Swift

To programmatically access AWS services, SDKs use a client class/object for each AWS 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 AWS service.

To make a request to an AWS service, you must first create and configure a service client. For each AWS service your code uses, it has its own crate/library/gem/package and its own dedicated type for interacting with it. The client exposes one method for each API operation exposed by the service.

The process of sending requests to AWS services is as follows:

1. Create a service client object with the desired configuration, such as the specific AWS Region.
2. Create an input object with the values and data needed to make the request. For example, when sending a `GetObject` request to Amazon S3, you need to specify the bucket name and the key

of the Amazon S3 object that you want to access. For a service action named `SomeOperation`, the input parameters object is created using a function called `SomeOperationInput()`.

3. Call the service object method on the client object that sends the desired request, with the input object created in the previous step.
4. Use `await` to wait for the response, and handle thrown exceptions to appropriately handle error conditions.
5. Examine the contents of the returned structure for the results you need. Every SDK function returns a structure with a type whose name is the same as the service action performed by the function, followed by the word `Output`. For example, when calling the Amazon S3 function `S3Client.createBucket(input:)`, the return type is `CreateBucketOutput`.

The following sections demonstrate this process in more detail.

Creating and using AWS client objects

Before you can send requests to an AWS service, you must first instantiate a client object corresponding to the service. These client classes are helpfully named using the service name and the word `Client`. Examples include `S3Client` and `IAMClient`.

After creating the client object, use it to make your requests. When you're done, release the object. If the service connection is open, it is closed for you automatically.

```
do {
    let s3 = try await S3Client()

    // ...
} catch {
    dump(error)
}
```

If an error occurs while trying to instantiate an AWS service — or at any time while using the service — an exception is thrown. Your `catch` block should handle the error appropriately.

Unless you are in a testing environment in which you expect a knowledgeable user to have configured reasonable default options, specify the appropriate service configuration when instantiating the client object. This is described in the section called “Client configuration in code”.

Specifying service client function parameters

When calling service client methods, you pass the input object corresponding to that operation. For example, before calling the `getObject()` method on the Amazon S3 service class `S3Client`, you need to create the input parameter object using the initializer `GetObjectInput()`.

```
do {
    let s3 = try S3Client()
    let inputObject = GetObjectInput(bucket: "amzn-s3-demo-bucket", key: "keyName")
    let output = try await s3.getObject(input: inputObject)

    // ...
} catch {
    dump(error)
}
```

In this example, `GetObjectInput()` is used to create an input object for the `getObject(input:)` method. The resulting input object specifies that the desired data has the `keyName` key and should be fetched from the Amazon S3 bucket named `amzn-s3-demo-bucket`.

Calling SDK functions

Nearly all AWS SDK for Swift functions are asynchronous and can be called using Swift's `async/await` model.

To call one of the SDK's asynchronous functions from synchronous code, call the function from a Swift `Task` created and run from your synchronous code.

Calling SDK functions asynchronously

The following function fetches and returns the content of a file named `text/motd.txt` from a bucket named amzn-s3-demo-bucket, using Amazon S3.

```
func getMOTD() async throws -> String? {
    let s3 = try S3Client()
    let motdInput = GetObjectInput(bucket: "amzn-s3-demo-bucket",
                                   key: "text/motd.txt")
    let output = try await s3.getObject(input: motdInput)

    guard let data = output.body?.toBytes().toData() else {
        return nil
```

```
    }
    return String(decoding: data, as: UTF8.self)
}
```

The `getMOTD()` function can only be called from another `async` function, and returns a string that contains the text in the `MOTD` file or `nil` if the file is empty. It throws an exception on errors. Thus, you call the `getMOTD()` function.

```
do {
    let motd = try await getMOTD()
    // ...
} catch {
    dump(error)
}
```

Here, the fetched "message of the day" text is available in the variable `motd` immediately following the call to `getMOTD()`. If an error occurs attempting to fetch the text, an appropriate exception is delivered to the `catch` clause. The standard Swift variable `error` describes the problem that occurred.

Calling SDK functions from synchronous code

To call AWS SDK for Swift functions from synchronous code, enclose the code that needs to run asynchronously in a `Task`. The `Task` uses `await` for each SDK function call that returns its results asynchronously. You might need to use an atomic flag or other means to know that the operation has finished.

> **Note**
>
> It's important to properly manage asynchronous requests. Be sure that any operation that's dependent on a previous result waits until that result is available before it begins. When used properly, the `async/await` model handles most of this for you.

```
func updateMOTD() {
    Task() {
        var motd: String = ""

        do {
            let s3 = try S3Client()
```

```
            let motdInput = GetObjectInput(bucket: "amzn-s3-demo-bucket",
                                key: "text/motd.txt")
            let output = try await s3.getObject(input: motdInput)

            if let bytes = output.body?.toBytes() {
                motd = String(decoding: bytes.toData(), as: UTF8.self)
            }
        } catch {
            motd = ""
        }

        setMOTD(motd)
    }
}
```

In this example, the code inside the `Task` block runs asynchronously, returning no output value to the caller. It fetches the contents of a text file with the key `text/motd.txt` and calls a function named `setMOTD()`, with the contents of the file decoded into a UTF-8 string.

A `do`/`catch` block is used to capture any thrown exceptions and set the `motd` variable to an empty string, which indicates that no message is available.

A call to this `updateMOTD()` function will spawn the task and return immediately. In the background, the program continues to run while the asynchronous task code fetches and uses the text from the specified file on Amazon S3. When the task has completed, the `Task` automatically ends.

Handling errors in the AWS SDK for Swift

Overview

The AWS SDK for Swift uses Swift's standard error handling mechanism to report errors that occur while using AWS services. Errors are reported using Swift's `throw` statement.

To catch the errors that an AWS SDK for Swift function might return, use the Swift `do`/`catch` statement. Encapsulate the code that calls the SDK inside the `do` block, then use one or more `catch` blocks to capture and handle the errors. Each `catch` block can capture a specific error, all errors of a specific category, or all uncaught errors. This lets an application recover from errors it knows how to handle, notify the user of transient errors or errors that can't be recovered from but are non-fatal, and safely exit the program if the error is fatal.

> **Note**
>
> The APs for each service client are generated using models specified using the Smithy interface definition language (IDL). The Smithy models describe the underlying types and APIs of each AWS service. These models are used to generate the Swift types and classes that comprise the SDK. This is useful to understand both while reading this guide and while writing error handling code.

AWS SDK for Swift error protocols

SDK for Swift errors conform to one or more error protocols. The protocols implemented by the error depend on the type of error that occurred and the context in which it occurred.

The `Error` protocol

Every error thrown by the AWS SDK for Swift conforms to the standard Swift `Error` protocol. As such, every error has a `localizedDescription` property that returns a string containing a useful description of the error.

When the underlying AWS service provides an error message, that string is used as the `localizedDescription`. These are usually in English. Otherwise, the SDK generates an appropriate message, which may or may not be localized.

The `AWSServiceError` protocol

When the AWS service responds with a service error, the error object conforms to the `AWSClientRuntime.AWSServiceError` protocol.

> **Note**
>
> If an `AWSServiceError` occurs while performing a service action over an HTTP connection, the error also implements the `HTTPError` protocol. Currently, all AWS protocols use HTTP, but if this were to change, an appropriate error protocol would be added.

Errors that conform to `AWSServiceError` include these additional properties:

`errorCode`

An `optional` string identifying the error type.

`requestID`

An `optional` string that gives the request ID of the request that resulted in the error.

The `ModeledError` protocol

When an error occurs that matches a defined, modeled error type, the error object conforms to the protocol `ClientRuntime.ModeledError`, in addition to any other appropriate protocols such as `HTTPError`. This includes most of the errors defined by an AWS service.

`ModeledError` adds several useful properties to the error:

`fault`

A value from the `ClientRuntime.ErrorFault` enum. The value is `.client` if the source of the error is the client, or `.server` if the server is the source of the error.

`isRetryable`

A Boolean value indicating whether or not the model indicates that the failed operation can be retried.

`isThrottling`

A Boolean value indicating whether or not the model indicates that the error is due to throttling.

The `HTTPError` protocol

Errors that occur during an action that uses an HTTP connection conform to the `ClientRuntime.HTTPError` protocol. An error conforming to `HTTPError` contains an HTTP response whose status code is in either the 4xx range or the 5xx range.

`HTTPError` adds one property to the error:

`httpResponse`

An object of type `HttpResponse`, which describes the entire HTTP response from the AWS service. It has properties that include the response's headers, body, and the HTTP status code.

Handling errors

All errors returned by the SDK for Swift implement the standard Swift `Error` protocol. The error's type depends on the service and the error being reported, so it could be any Swift type including but not limited to `enum`, `struct`, or `class`, depending on what kind of error occurred. For example, an error reporting that an Amazon S3 bucket is missing may conform to `Error`, `AWSServiceError`, and `HTTPError`. This lets you know it's a service error that occurred while communicating using the HTTP protocol. In this case, the HTTP status code is 404 (Not Found), because of the missing bucket.

Even if no other information is provided, the error's `localizedDescription` property is always a string describing the error.

When catching errors thrown by the AWS SDK for Swift, follow these guidelines:

- If the error is modeled, the error is a `struct` describing the error. Catch these errors using that `struct`'s name. In many cases, you can find these modeled errors listed in the documentation of an action in the AWS SDK for Swift API Reference.
- If the error isn't modeled, but still originates from an AWS service, it will conform to the protocol `AWSServiceError`. Use `catch let error as AWSServiceError`, then look at the error's `errorCode` property to determine what error occurred.
- **Don't catch any concrete types that represent unknown errors**, such as `UnknownAWSHTTPServiceError`. These are reserved for internal use.

Service errors

An error thrown because of an AWS service response, whether it could be parsed or not, conforms to the `AWSServiceError` protocol. An error defined by the underlying Smithy model for service also conforms to `ModeledError` and has a concrete type. One example is the Amazon S3 error `CreateBucketOutputError`, which is thrown by the `S3Client.createBucket()` method.

Any `AWSServiceError` received over an HTTP connection also conforms to `HTTPError`. This is currently all service errors, but that could change in the future if a service adds support for other network protocols.

The following code tries to create an object on Amazon S3, with code to handle service errors. It features a `catch` clause that specifically handles the error code `NoSuchBucket`, which indicates that the bucket doesn't exist. This snippet assumes that the given bucket name doesn't exist.

```
        do {
            let client = try S3Client(region: "us-east-1")

            _ = try await client.putObject(input: PutObjectInput(
                body: ByteStream.data(Data(body.utf8)),
                bucket: bucketName,
                key: objectKey
            ))
            print("Done.")
        } catch let error as AWSServiceError {
            let errorCode = error.errorCode ?? "<none>"
            let message = error.message ?? "<no message>"

            switch errorCode {
            case "NoSuchBucket":
                print("   | The bucket \"\(bucketName)\" doesn't exist. This is the
expected result.")
                print("   | In a real app, you might ask the user whether to use a
different name or")
                print("   | create the bucket here.")
            default:
                print("   | Service error of type \(error.errorCode ?? "<unknown>"):
\(message)")
            }
        } catch {
            print("Some other error occurred.")
        }
```

HTTP errors

When the SDK encounters an error while communicating with an AWS service over HTTP, it throws an error that conforms to the protocol `ClientRuntime.HTTPError`. This kind of error represents an HTTP response whose status codes are in the 4xx and 5xx ranges.

> **Note**
>
> Currently, HTTP is the only network protocol used by AWS. If a future AWS product uses a non-HTTP network protocol, a corresponding error protocol would be added to the SDK. Errors that occur while using the new wire protocol would conform to that Swift protocol instead of `HTTPError`.

`HTTPError` includes an `httpResponse` property that contains an object of the class `HttpResponse`. The `httpResponse` provides information received in the response to the failed HTTP request. This provides access to the response headers, including the HTTP status code.

```
        do {
            let client = try S3Client(region: "us-east-1")

            _ = try await client.getObject(input: GetObjectInput(
                bucket: "not-a-real-bucket",
                key: "not-a-real-key"
            ))
            print("   | Found a matching bucket but shouldn't have!")
        } catch let error as HTTPError {
            print("   | HTTP error; status code:
 \(error.httpResponse.statusCode.rawValue). This is the")
            print("   | expected result.")
        } catch {
            dump(error, name: "   | An unexpected error occurred.")
        }
```

This example creates an Amazon S3 client, then calls its `getObject(input:)` function to fetch an object using a bucket name and key that don't exist. Two `catch` blocks are used. The first matches errors of type `HTTPError`. It retrieves the HTTP status code from the response. The status code can then be used to handle specific scenarios, recover from recoverable errors, or whatever the project requires.

The second `catch` block is a catch-all that just dumps the error to the console. In a full application, this block would ideally either clean up after the failed access attempt and return the application to a safe state, or perform as clean an application exit as possible.

Handling other errors

To catch any errors not already caught for a given do block, use the `catch` keyword with no qualifiers. The following snippet simply catches all errors.

```
do {
    let s3 = try await S3Client()

    // ...
} catch {
    // Handle the error here.
```

```
}
```

Within the context of the `catch` block, the caught error, reported in the constant with the default name `error`, conforms to at least the standard Swift `Error` type. It may also conform to a combination of the other AWS SDK for Swift error protocols.

If you use a catch-all like this in your code, it needs to safely stop whatever task it was trying to perform and clean up after itself. In extreme cases, it may be necessary to safely terminate the application and ideally provide diagnostic output to be relayed to the developer.

While developing a project, it can be helpful to temporarily output error details to the console. This can be useful when debugging, or to help determine which errors that occur may need special handling. The Swift `dump()` function can be used to do this.

```
do {
    let output = try await client.listBuckets(input: ListBucketsInput())

    // ...
} catch {
    dump(error, name: "Getting the bucket list")
}
```

The `dump()` function outputs the entire contents of the `error` to the console. The `name` argument is an `optional` string used as a label for the output, to help identify the source of the error in the program's output.

Using paginated results in the AWS SDK for Swift

Many AWS operations return truncated results when the payload is too large to return in a single response. Instead, the service returns a portion of the data and a token to retrieve the next set of items. This pattern is known as pagination.

The AWS SDK for Swift provides specialized versions of the functions that provide paginated results. These special functions end with the word **`Paginated`**. All your code needs to do is process the results as they arrive.

Each paginator is a function that returns an object of type `PaginatorSequence<`*`input-type`*`,` *`output-type`*`>`. The `PaginatorSequence<>` is an `AsyncSequence`. The `PaginatorSequence` type is a "lazy" sequence, so no AWS service requests are made until you start iterating over the

pages. This also means that any errors that occur during the operation don't reach you until iteration begins.

> **Note**
>
> The examples in this section of the developer guide use Amazon S3. However, the concept is the same for any service that has one or more paginated APIs.

For example, the paginated version of the `S3Client` function `listBuckets(input:)`, `listBucketsPaginated(input:)`, returns an object of type `PaginatorSequence<ListBucketsInput, ListBucketsOutput>`:

```
let pages = client.listBucketsPaginated(
    input: ListBucketsInput(maxBuckets: PAGE_SIZE)
)
```

In this example, the number of results in each page is specified by adding a `maxBuckets` property to the `ListBucketsInput` object. Each paginator uses an appropriate name for this property. As of the time `listBucketsPaginated(input:)` returns, no requests have been sent to the Amazon S3 service.

The `PaginatorSequence<>` is a sequence of pages which are asynchronously added to the sequence as the results are received. The type of each entry in the sequence is the `Output struct` corresponding to the function called. For example, if you call `S3Client.listBucketsPaginated(input:)`, each entry in the sequence is a `ListBucketsOutput` object. Each entry's buckets can be found in the its `ListBucketsOutput.buckets` property, which is an array of objects of type `S3ClientTypes.Bucket`.

To begin sending requests and receiving results, asynchronously iterate over each page, then iterate over each page's items:

```
var pageNumber = 0

do {
    for try await page in pages {
        pageNumber += 1
```

```
                guard let pageBuckets = page.buckets else {
                    print("ERROR: No buckets returned in page \(pageNumber)")
                    continue
                }

                print("\nPage \(pageNumber):")

                // Print this page's bucket names.

                for bucket in pageBuckets {
                    print("  " + (bucket.name ?? "<unknown>"))
                }
            }
        } catch {
            print("ERROR: Unable to process bucket list pages.")
        }
```

The outer `for` loop uses `await` to process pages of results as they're delivered, asynchronously. Once a page is received, the inner loop iterates over the buckets found in each entry's `buckets` property. The full example is available on GitHub.

Testing and debugging with the AWS SDK for Swift

Logging

The AWS SDK for Swift and the underlying Common RunTime (CRT) library use Apple's SwiftLog mechanism to log informational and debugging messages. The output from the logging system appears in the console provided by most IDEs, and also in the standard system console. This section covers how to control the output level for both the SDK for Swift and the CRT library.

Configuring SDK debugging

By default, the SDK's logging system emits logs containing trace and debug level information. The default log level is `info`. To see more informational text as the SDK works, you can change the log level to `error` as shown in the following example:

```
import ClientRuntime

await SDKLoggingSystem().initialize(logLevel: .error)
```

Call `SDKLoggingSystem.initialize(logLevel:)` no more than one time in your application, to set the log level cutoff.

The log levels supported by the AWS SDK for Swift are defined in the `SDKLogLevel enum`. These correspond to similarly-named log levels defined in SwiftLog. Each log level is inclusive of the messages at and more severe that level. For example, setting the log level to `warning` also causes log messages to be output for levels `error` and `critical`.

The SDK for Swift log levels (from least severe to most severe) are:

- `.trace`
- `.debug`
- `.info` (the default)
- `.notice`
- `.warning`
- `.error`
- `.critical`

Configuring Common RunTime logging

If the level of detail generated by the SDK logs isn't providing what you need, you can try configuring the Common RunTime (CRT) log level. CRT is responsible for the lower-level networking and other system interactions performed by the SDK. Its log output includes details about HTTP traffic, for example.

To set CRT's log level to `debug`:

```
import ClientRuntime

SDKDefaultIO.setLogLevel(level: .debug)
```

The CRT log levels (from least severe to most severe) are:

- `.none`
- `.trace`
- `.debug`
- `.info`

- `.warn`
- `.error`
- `.fatal`

Setting the CRT log level to `trace` provides an enormous level of detail that can take some effort to read, but it can be useful in demanding debugging situations.

Mocking the AWS SDK for Swift

When writing unit tests for your AWS SDK for Swift project, it's useful to be able to mock the SDK. Mocking is a technique for unit testing in which external dependencies — such as the SDK for Swift — are replaced with code that simulates those dependencies in a controlled and predictable way. Mocking the SDK removes network requests, which eliminates the chance that tests can be unreliable due to intermittent network issues.

In addition, well-written mocks are almost always faster than the operations they simulate, letting you test more thoroughly in less time.

The Swift language doesn't provide the read/write reflection needed for direct mocking. Instead, you adapt your code to allow an indirect form of mocking. This article describes how to do so with minimal changes to the main body of your code.

To mock the AWS SDK for Swift implementation of a service class, create a protocol. In the protocol, define each of that class's functions that you need to use. This serves as the abstraction layer that you need to implement mocking. It's up to you whether to use a separate protocol for each AWS service class used in your project. Alternatively, you can use a single protocol that encapsulates every SDK function that you call. The examples in this guide use the latter approach, but this is often the same as using a protocol for each service.

After you define the protocol, you need two classes that conform to the protocol: one class in which each function calls through to the corresponding SDK function, and one that mocks the results as if the SDK function was called. Because these two classes both conform to the same protocol, you can create functions that perform AWS actions by calling functions on an object conforming to the protocol.

Example: Mocking an Amazon S3 function

Consider a program that needs to use the `S3Client` function `listBuckets(input:)`. To support mocking, this project implements the following:

- `S3SessionProtocol`, a Swift protocol which declares the Amazon S3 functions used by the project. This example uses just one Amazon S3 function: `listBuckets(input:)`.
- `S3Session`, a class conforming to `S3SessionProtocol`, whose implementation of `listBuckets(input:)` calls `S3Client.listBuckets(input:)`. This is used when running the program normally.
- `MockS3Session`, a class conforming to `S3SessionProtocol`, whose implementation of `listBuckets(input:)` returns mocked results based on the input parameters. This is used when running tests.
- `BucketManager`, a class that implements access to Amazon S3. This class should accept an object conforming to the session protocol `S3SessionProtocol` during initialization, then perform all AWS requests by making calls through that object. This makes the code testable: the *application* initializes the class by using an `S3Session` object for AWS access, while *tests* use a `MockS3Session` object.

The rest of this section takes an in-depth look at this implementation of mocking. The complete example is available on GitHub.

Protocol

In this example, the `S3SessionProtocol` protocol declares the one `S3Client` function that it needs:

```
/// The S3SessionProtocol protocol describes the Amazon S3 functions this
/// program uses during an S3 session. It needs to be implemented once to call
/// through to the corresponding SDK for Swift functions, and a second time to
/// instead return mock results.
public protocol S3SessionProtocol {
    func listBuckets(input: ListBucketsInput) async throws
            -> ListBucketsOutput
}
```

This protocol describes the interface by which the pair of classes perform Amazon S3 actions.

Main program implementation

To let the main program make Amazon S3 requests using the session protocol, you need an implementation of the protocol in which each function calls the corresponding SDK

function. In this example, you create a class named `S3Session` with an implementation of `listBuckets(input:)` that calls `S3Client.listBuckets(input:)`:

```
public class S3Session: S3SessionProtocol {
    let client: S3Client
    let region: String

    /// Initialize the session to use the specified AWS Region.
    ///
    /// - Parameter region: The AWS Region to use. Default is `us-east-1`.
    init(region: String = "us-east-1") throws {
        self.region = region

        // Create an ``S3Client`` to use for AWS SDK for Swift calls.
        self.client = try S3Client(region: self.region)
    }

    /// Call through to the ``S3Client`` function `listBuckets()`.
    ///
    /// - Parameter input: The input to pass through to the SDK function
    ///   `listBuckets()`.
    ///
    /// - Returns: A ``ListBucketsOutput`` with the returned data.
    ///
    public func listBuckets(input: ListBucketsInput) async throws
            -> ListBucketsOutput {
        return try await self.client.listBuckets(input: input)
    }
}
```

The initializer creates the underlying `S3Client` through which the SDK for Swift is called. The only other function is `listBuckets(input:)`, which returns the result of calling the `S3Client` function of the same name. Calls to AWS services work the same way they do when calling the SDK directly.

Mock implementation

In this example, add support for mocking calls to Amazon S3 by using a second implementation of `S3SessionProtocol` called `MockS3Session`. In this class, the `listBuckets(input:)` function generates and returns mock results:

```
    /// An implementation of the Amazon S3 function `listBuckets()` that
```

```
    /// returns the mock data instead of accessing AWS.
    ///
    /// - Parameter input: The input to the `listBuckets()` function.
    ///
    /// - Returns: A `ListBucketsOutput` object containing the list of
    ///   buckets.
    public func listBuckets(input: ListBucketsInput) async throws
            -> ListBucketsOutput {
        let response = ListBucketsOutput(
            buckets: self.mockBuckets,
            owner: nil
        )
        return response
    }
```

This works by creating and returning a `ListBucketsOutput` object, like the actual `S3Client` function does. Unlike the SDK function, this makes no actual AWS service requests. Instead, it fills out the response object with data that simulates actual results. In this case, an array of `S3ClientTypes.Bucket` objects describing a number of mock buckets is returned in the `buckets` property.

Not every property of the returned response object is filled out in this example. The only properties that get values are those that always contain a value and those actually used by the application. Your project might require more detailed results in its mock implementations of functions.

Encapsulating access to AWS services

A convenient way to use this approach in your application design is to create an access manager class that encapsulates all your SDK calls. For example, when using Amazon DynamoDB (DynamoDB) to manage a product database, create a `ProductDatabase` class that has functions to perform needed activities. This might include adding products and searching for products. This Amazon S3 example has a class that handles bucket interactions, called `BucketManager`.

The `BucketManager` class initializer needs to accept an object conforming to `S3SessionProtocol` as an input. This lets the caller specify whether to interact with AWS by using actual SDK for Swift calls or by using a mock. Then, every other function in the class that uses AWS actions should use that session object to do so. This lets `BucketManager` use actual SDK calls or mocked ones based on whether testing is underway.

With this in mind, the `BucketManager` class can now be implemented. It needs an `init(session:)` initializer and a `getBucketNames(input:)` function:

```
public class BucketManager {
    /// The object based on the ``S3SessionProtocol`` protocol through which to
    /// call SDK for swift functions. This may be either ``S3Session`` or
    /// ``MockS3Session``.
    var session: S3SessionProtocol

    /// Initialize the ``BucketManager`` to call Amazon S3 functions using the
    /// specified object that implements ``S3SessionProtocol``.
    ///
    /// - Parameter session: The session object to use when calling Amazon S3.
    init(session: S3SessionProtocol) {
        self.session = session
    }

    /// Return an array listing all of the user's buckets by calling the
    /// ``S3SessionProtocol`` function `listBuckets()`.
    ///
    /// - Returns: An array of bucket name strings.
    ///
    public func getBucketNames() async throws -> [String] {
        let output = try await session.listBuckets(input: ListBucketsInput())

        guard let buckets = output.buckets else {
            return []
        }

        return buckets.map { $0.name ?? "<unknown>" }
    }
}
```

The `BucketManager` class in this example has an initializer that takes an object that conforms to `S3SessionProtocol` as an input. That session is used to access or simulate access to AWS actions instead of calling the SDK directly, as shown by the `getBucketNames()` function.

Using the access manager in the main program

The main program can now specify an `S3Session` when creating a `BucketManager` object, which directs its requests to AWS:

```
        /// An ``S3Session`` object that passes calls through to the SDK for
        /// Swift.
        let session: S3Session
```

```
        /// A ``BucketManager`` object that will be initialized to call the
        /// SDK using the session.
        let bucketMgr: BucketManager

        // Create the ``S3Session`` and a ``BucketManager`` that calls the SDK
        // using it.
        do {
            session = try S3Session(region: "us-east-1")
            bucketMgr = BucketManager(session: session)
        } catch {
            print("Unable to initialize access to Amazon S3.")
            return
        }
```

Writing tests using the protocol

Whether you write tests using Apple's XCTest framework or another framework, you must design the tests to use the mock implementation of the functions that access AWS services. In this example, tests use a class of type `XCTestCase` to implement a standard Swift test case:

```
final class MockingTests: XCTestCase {
    /// The session to use for Amazon S3 calls. In this case, it's a mock
    /// implementation.
    var session: MockS3Session? = nil
    /// The ``BucketManager`` that uses the session to perform Amazon S3
    /// operations.
    var bucketMgr: BucketManager? = nil

    /// Perform one-time initialization before executing any tests.
    override class func setUp() {
        super.setUp()
    }

    /// Set up things that need to be done just before each
    /// individual test function is called.
    override func setUp() {
        super.setUp()

        self.session = MockS3Session()
        self.bucketMgr = BucketManager(session: self.session!)
    }

    /// Test that `getBucketNames()` returns the expected results.
```

```
    func testGetBucketNames() async throws {
        let returnedNames = try await self.bucketMgr!.getBucketNames()
        XCTAssertTrue(self.session!.checkBucketNames(names: returnedNames),
                "Bucket names don't match")
    }
}
```

This `XCTestCase` example's per-test `setUp()` function creates a new `MockS3Session`. Then it uses the mock session to create a `BucketManager` that will return mock results. The `testGetBucketNames()` test function tests the `getBucketNames()` function in the bucket manager object. This way, the tests operate using known data, without needing to access the network, and without accessing AWS services at all.

Using waiters with the AWS SDK for Swift

A waiter is a client-side abstraction that automatically polls a resource until a desired state is reached, or until it's determined that the resource won't enter that state. This kind of polling is commonly used when working with services that typically reach a consistent state, such as Amazon Simple Storage Service (Amazon S3), or services that create resources asynchronously, like Amazon Elastic Compute Cloud (Amazon EC2).

Instead of writing logic to continuously poll an AWS resource, which can be cumbersome and error-prone, you can use a waiter to poll the resource. When the waiter returns, your code knows whether or not the desired state was reached and can act accordingly.

How to use waiters

Many service client classes in the AWS SDK for Swift provide waiters for common "poll and wait" situations that occur while using AWS services. These situations can include both waiting until a resource is available and waiting until the resource becomes unavailable. For example:

`S3Client.waitUntilBucketExists(options:input:)`, **`S3Client.waitUntilBucketNotExists(options:input:)`**

Wait until an Amazon S3 bucket exists (or no longer exists).

`DynamoDBClient.waitUntilTableExists(options:input:)`, **`DynamoDBClient.waitUntilTableNotExists(options:input:)`**

Wait until a specific Amazon DynamoDB table exists (or no longer exists).

Waiter functions in the AWS SDK for Swift take two parameters, `options` and `input`.

Waiter options

The first parameter for any waiter function, `options`, is a structure of type `WaiterOptions` (part of the `ClientRuntime` package) that describes the waiter's polling behavior. These options specify the maximum number of seconds to wait before polling times out, and let you optionally set the minimum and maximum delays between retries.

The following example shows how to configure a waiter to wait between a tenth of a second and a half second between retries. The maximum polling time is two seconds.

```
let options = WaiterOptions(maxWaitTime: 2, minDelay: 0.1, maxDelay: 0.5)
```

Waiter parameters

A waiter function's inputs are specified using its `input` parameter. The input's data type corresponds to that of the SDK for Swift function used internally by the waiter to poll the AWS service. For example, the type is `HeadBucketInput` for Amazon S3 waiters like `S3Client.waitUntilBucketExists(options:input:)` because this waiter uses the `S3Client.headBucket(input:)` function to poll the bucket. The DynamoDB waiter `DynamoDBClient.waitUntilTableExists(options:input:)` takes as its input a structure of type `DescribeTableInput` because it calls `DynamoDBClient.describeTable(input:)` internally.

Example: Waiting for an S3 bucket to exist

The AWS SDK for Swift offers several waiters for Amazon S3. One of them is `waitUntilBucketExists(options:input:)`, which polls the server until the specified bucket exists.

```
    /// Wait until a bucket with the specified name exists, then return
    /// to the caller. Times out after 60 seconds. Throws an error if the
    /// wait fails.
    ///
    /// - Parameter bucketName: A string giving the name of the bucket
    ///   to wait for.
    ///
    /// - Returns: `true` if the bucket was found or `false` if not.
```

```
    ///
    public func waitForBucket(name bucketName: String) async throws -> Bool {
        // Because `waitUntilBucketExists()` internally uses the Amazon S3
        // action `HeadBucket` to look for the bucket, the input is specified
        // with a `HeadBucketInput` structure.

        let output = try await client.waitUntilBucketExists(
            options: WaiterOptions(maxWaitTime: 60.0),
            input: HeadBucketInput(bucket: bucketName)
        )

        switch output.result {
            case .success:
                return true
            case .failure:
                return false
        }
    }
```

This example creates a function that waits until the specified bucket exists, then returns `true`. It returns `false` if polling fails, and might throw an exception if an error occurs. It works by calling the waiter function `S3Client.waitUntilBucketExists(options:input:)` to poll the server. The `options` specify that polling should time out after 60 seconds.

Because this polling is done using the Amazon S3 action `HeadBucket`, a `HeadBucketInput` object is created with the input parameters for that operation, including the name of the bucket to poll for. This is used as the `input` parameter's value.

`S3Client.waitUntilBucketExists(options:input:)` returns a `result` property whose value is either `.success` if the polling successfully found the bucket, or `.failure` if polling failed. The function returns the corresponding Boolean value to the caller.

Creating AWS Lambda functions using the AWS SDK for Swift

Overview

You can use the AWS SDK for Swift from within an AWS Lambda function by using the Swift AWS Lambda Runtime package in your project. See the documentation for the `swift-aws-lambda-runtime` repository on GitHub for more information about the runtime package.

Setting up a project to use AWS Lambda

If you're starting a new project, create the project in Xcode or open a shell session and use the following command to use Swift Package Manager (SwiftPM) to manage your project:

```
$ swift package init --type executable --name LambdaExample
```

Remove the file `Sources/main.swift`. The source code file will have be `Sources/lambda.swift` to work around a known Swift bug that can cause problems when the entry point is in a file named `main.swift`.

Add the `swift-aws-lambda-runtime` package to the project. There are two ways to accomplish this:

- If you're using Xcode, choose the **Add package dependencies...** option in the **File** menu, then provide the package URL: `https://github.com/awslabs/swift-aws-lambda-runtime.git`. Choose the `AWSLambdaRuntime` module.
- If you're using SwiftPM to manage your project dependencies, add the runtime package and its `AWSLambdaRuntime` module to your `Package.swift` file to make the module available to your project:

```
import PackageDescription

let package = Package(
    name: "LambdaExample",
    platforms: [
        .macOS(.v12)
    ],
    // The product is an executable named "LambdaExample", which is built
    // using the target "LambdaExample".
    products: [
        .executable(name: "LambdaExample", targets: ["LambdaExample"])
    ],
    // Add the dependencies: these are the packages that need to be fetched
    // before building the project.
    dependencies: [
        .package(
            url: "https://github.com/awslabs/swift-aws-lambda-runtime.git",
            from: "2.0.0"),
        .package(url: "https://github.com/awslabs/aws-sdk-swift.git",
            from: "1.0.0"),
```

```
    ],
    targets: [
        // Add the executable target for the main program. These are the
        // specific modules this project uses within the packages listed under
        // "dependencies."
        .executableTarget(
            name: "LambdaExample",
            dependencies: [
                .product(name: "AWSLambdaRuntime", package: "swift-aws-lambda-
runtime"),
                .product(name: "AWSS3", package: "aws-sdk-swift"),
            ]
        )
    ]
)
```

This example adds a dependency on the Amazon S3 module of the AWS SDK for Swift in addition to the Lambda runtime.

You may find it useful to build the project at this point. Doing so will pull the dependencies and may make them available for your editor or IDE to generate auto-completion or inline help:

```
$ swift build
```

Creating a Lambda function

To create a Lambda function in Swift, you generally need to define several components:

- A `struct` that represents the data your Lambda function will receive from the client. It must implement the `Decodable` protocol. The Swift AWS Lambda Runtime Events library contains a variety of struct definitions that represent common messages posted to a Lambda function by other AWS services.
- A Lambda function handler that performs the Lambda function's work.
- An optional `struct` that represents the data returned by your Lambda function. This is usually an `Encodable` `struct` describing the contents of a JSON document returned to the client.

Imports

Create the file `Sources/lambda.swift` and begin by importing the needed modules and types:

```
import Foundation
import AWSLambdaRuntime
@preconcurrency import AWSS3

import protocol AWSClientRuntime.AWSServiceError
import enum Smithy.ByteStream
```

These imports are:

1. The standard Apple `Foundation` API.
2. `AWSLambdaRuntime`is the Swift Lambda Runtime's main module.
3. `AWSS3` is the Amazon S3 module from the AWS SDK for Swift.
4. The `AWSClientRuntime.AWSServiceError` protocol describes service errors returned by the SDK.
5. The `Smithy.ByteStream` enum is a type that represents a stream of data. The Smithy library is one of the SDK's core modules.

Defining structs and enums

Next, define the structs that represent the incoming requests and the responses sent back by the Lambda function, along with the enum used to identify errors thrown by the handler function:

```
/// Represents the contents of the requests being received from the client.
/// This structure must be `Decodable` to indicate that its initializer
/// converts an external representation into this type.
struct Request: Decodable, Sendable {
    /// The request body.
    let body: String
}

/// The contents of the response sent back to the client. This must be
/// `Encodable`.
struct Response: Encodable, Sendable {
    /// The ID of the request this response corresponds to.
    let req_id: String
    /// The body of the response message.
    let body: String
}

/// The errors that the Lambda function can return.
```

```
enum S3ExampleLambdaErrors: Error {
    /// A required environment variable is missing. The missing variable is
    /// specified.
    case noEnvironmentVariable(String)
}
```

Function handler

To receive, process, and respond to incoming requests, implement the Swift Lambda Runtime's `LambdaRuntime` class like this:

```
let runtime = LambdaRuntime {
    (event: Request, context: LambdaContext) async throws -> Response in

    var responseMessage: String

    // Get the name of the bucket to write the new object into from the
    // environment variable `BUCKET_NAME`.
    guard let bucketName = ProcessInfo.processInfo.environment["BUCKET_NAME"] else {
        context.logger.error("Set the environment variable BUCKET_NAME to the name of
 the S3 bucket to write files to.")
        throw S3ExampleLambdaErrors.noEnvironmentVariable("BUCKET_NAME")
    }

    do {
        let filename = try await putObject(body: event.body, bucketName: bucketName)

        // Generate the response text and update the log.
        responseMessage = "The Lambda function has successfully stored your data in S3
 with name '\(filename)'"
        context.logger.info("Data successfully stored in S3.")
    } catch let error as AWSServiceError {
        // Generate the error message and update the log.
        responseMessage = "The Lambda function encountered an error and your data was
 not saved. Root cause: \(error.errorCode ?? "") - \(error.message ?? "")"
        context.logger.error("Failed to upload data to Amazon S3.")
    }

    return Response(req_id: context.requestID, body: responseMessage)
}
```

The name of the bucket to use is first fetched from the environment variable `BUCKET_NAME`. Then the `putObject(body:bucketName:)` function is called to write the text into an Amazon S3

object, and the text of the response is set depending on whether or not the object is successfully written to storage. The response is created with the response message and the request ID string that was in the original Lambda request.

Helper function

This example uses a helper function, `putObject(body:bucketName:)`, to write strings to Amazon S3. The string is stored in the specified bucket by creating an object whose name is based on the current number of seconds since the start of the year 1970:

```
/// Create a new object on Amazon S3 whose name is based on the current
/// timestamp, containing the text specified.
///
/// - Parameters:
///   - body: The text to store in the new S3 object.
///   - bucketName: The name of the Amazon S3 bucket to put the new object
///     into.
///
/// - Throws: Errors from `PutObject`.
///
/// - Returns: The name of the new Amazon S3 object that contains the
///   specified body text.
func putObject(body: String, bucketName: String) async throws -> String {
    // Generate an almost certainly unique object name based on the current
    // timestamp.

    let objectName = "\(Int(Date().timeIntervalSince1970*1_000_000)).txt"

    // Create a Smithy `ByteStream` that represents the string to write into
    // the bucket.

    let inputStream = Smithy.ByteStream.data(body.data(using: .utf8))

    // Store the text into an object in the Amazon S3 bucket.

    _ = try await s3Client.putObject(
        input: PutObjectInput(
            body: inputStream,
            bucket: bucketName,
            key: objectName
        )
    )
```

```
    // Return the name of the file

    return objectName
}
```

Start the runtime

The final step is to have a global statement that calls the example runtime's `run()` function:

```
try await runtime.run()
```

Build and test locally

While you can test your Lambda function by adding it in the Lambda console, the Swift AWS Lambda Runtime provides an integrated Lambda server you can use for testing. This server accepts requests and dispatches them to your Lambda function. This integrated server is started automatically when you run your program locally.

In this example, the program is built and run with the Region set to `eu-west-1`, the bucket name set to `amzn-s3-demo-bucket`, and the local Lambda server enabled:

```
$ AWS_REGION=eu-west-1            \
BUCKET_NAME=amzn-s3-demo-bucket    \
swift run
```

After running this command, the Lambda function is available on the local server.
Test it by opening another terminal session and using it to send a Lambda request to `http://127.0.0.1:7000/invoke`, or to port 7000 on `localhost`:

```
$ curl -X POST                          \
     --data '{"body":"This is the message to store on Amazon S3."}' \
   http://127.0.0.1:7000/invoke
```

Upon success, a JSON object similar to this is returned:

```
{
  "req_id": "290935198005708",
  "body": "The Lambda function has successfully stored your data in S3 with name
 '1720098625801368.txt'"
}
```

You can remove the created object from your bucket using this AWS CLI command:

```
$ aws s3 rm s3://amzn-s3-demo-bucket/file-name
```

Packaging and uploading the app

To use a Swift app as a Lambda function, compile it for an x86_64 or ARM Linux target depending on the build machine's architecture. This may involve cross-compiling, so you may need to resolve dependency issues, even if they don't happen when building for your build system.

The Swift Lambda Runtime includes an `archive` command as a plugin for the Swift compiler. This plugin lets you cross-compile from macOS to Linux just using the standard `swift` command. The plugin uses a Docker container to build the Linux executable, so you'll need Docker installed.

To build your app for use as a Lambda function:

1. Build the app using the SwiftPM `archive` plugin. This automatically selects the architecture based on that of your build machine (x86_64 or ARM).

   ```
   $ swift package archive --allow-network-connections docker
   ```

   This creates a ZIP file containing the function executable, placing the output in `.build/plugins/AWSLambdaPackager/outputs/AWSLambdaPackager/`*`target-name`*`/`*`executable-name`*`.zip`

2. Create a Lambda function using the method appropriate for your needs, such as:

   - AWS Lambda Developer Guide
   - AWS Command Line Interface User Guide
   - AWS Cloud Development Kit (AWS CDK) Developer Guide

   > **Warning**
   >
   > Things to keep in mind when deploying the Lambda function:
   >
   > - Use the same architecture (x86_64 or ARM64) for your function and your binary.
   > - Use the Amazon Linux 2 runtime.
   > - Define any environment variables required by the function. In this example, the `BUCKET_NAME` variable needs to be set to the name of the bucket to write objects into.

- Give your function the needed permissions to access AWS resources. For this example, the function needs IAM permission to use `PutObject` on the bucket specified by `BUCKET_NAME`.

3. Once you've created and deployed the Swift-based Lambda function, it should be ready to accept requests. You can invoke the function using the `Invoke` Lambda API.

```
$ aws lambda invoke                    \
--region eu-west-1                     \
--function-name LambdaExample          \
--cli-binary-format raw-in-base64-out  \
--payload '{"body":"test message"}'    \
output.json
```

The file `output.json` contains the results of the invocation (or the error message injected by our code).

Additional information

- Swift AWS Lambda Runtime package on GitHub
- Swift AWS Lambda Events package on GitHub
- AWS Lambda Developer Guide
- AWS Lambda API Reference

Event streaming with AWS SDK for Swift

Overview

Some AWS services provide timely feedback about the state of a system or process by streaming to your application a series of events describing that state, if your application wants to receive them. Likewise, other services may be able to receive a stream of events from your application to provide needed data as it becomes available. The AWS SDK for Swift provides support for sending and receiving streams of events with services that support this feature.

This section of the guide demonstrates how to stream events to a service and receive events from a service, with an example that uses Amazon Transcribe to transcribe voice content from an audio file into text displayed on the screen.

Event streaming example

This Amazon Transcribe example uses the swift-argument-parser package from Apple to parse the command line neatly, as well as the `AWSTranscribeStreaming` module from the AWS SDK for Swift.

The example's complete source code is available on GitHub.

Importing modules

The example begins by importing the modules it needs:

```
import ArgumentParser
import AWSClientRuntime
import AWSTranscribeStreaming
import Foundation
```

Enum definition

Then an enum is defined to represent the three audio formats Amazon Transcribe supports for streaming. These are used to match against the format specified on the command line using the **--format** option.:

```
/// Identify one of the media file formats supported by Amazon Transcribe.
enum TranscribeFormat: String, ExpressibleByArgument {
    case ogg = "ogg"
    case pcm = "pcm"
    case flac = "flac"
}
```

Creating the audio stream

A function named `createAudioStream()` returns an `AsyncThrowingStream` that contains the audio file's contents, broken into 125ms chunks. The `AsyncThrowingStream` supplies audio data to Amazon Transcribe. The stream is specified as an input property when calling the client's `startStreamTranscription(input:)` function.

```
    /// Create and return an Amazon Transcribe audio stream from the file
    /// specified in the arguments.
    ///
    /// - Throws: Errors from `TranscribeError`.
    ///
    /// - Returns: `AsyncThrowingStream<TranscribeStreamingClientTypes.AudioStream,
Error>`
    func createAudioStream() async throws
                -> AsyncThrowingStream<TranscribeStreamingClientTypes.AudioStream,
Error> {

        let fileURL: URL = URL(fileURLWithPath: path)
        let audioData = try Data(contentsOf: fileURL)

        // Properties defining the size of audio chunks and the total size of
        // the audio file in bytes. You should try to send chunks that last on
        // average 125 milliseconds.

        let chunkSizeInMilliseconds = 125.0
        let chunkSize = Int(chunkSizeInMilliseconds  / 1000.0 * Double(sampleRate) *
2.0)
        let audioDataSize = audioData.count

        // Create an audio stream from the source data. The stream's job is
        // to send the audio in chunks to Amazon Transcribe as
        // `AudioStream.audioevent` events.

        let audioStream =
AsyncThrowingStream<TranscribeStreamingClientTypes.AudioStream,
                                Error> { continuation in
            Task {
                var currentStart = 0
                var currentEnd = min(chunkSize, audioDataSize - currentStart)

                // Generate and send chunks of audio data as `audioevent`
                // events until the entire file has been sent. Each event is
                // yielded to the SDK after being created.

                while currentStart < audioDataSize {
                    let dataChunk = audioData[currentStart ..< currentEnd]

                    let audioEvent =
TranscribeStreamingClientTypes.AudioStream.audioevent(
```

```
                        .init(audioChunk: dataChunk)
                    )
                    let yieldResult = continuation.yield(audioEvent)
                    switch yieldResult {
                        case .enqueued(_):
                            // The chunk was successfully enqueued into the
                            // stream. The `remaining` parameter estimates how
                            // much room is left in the queue, but is ignored here.
                            break
                        case .dropped(_):
                            // The chunk was dropped because the queue buffer
                            // is full. This will cause transcription errors.
                            print("Warning: Dropped audio! The transcription will be
incomplete.")
                        case .terminated:
                            print("Audio stream terminated.")
                            continuation.finish()
                            return
                        default:
                            print("Warning: Unrecognized response during audio
streaming.")
                    }

                    currentStart = currentEnd
                    currentEnd = min(currentStart + chunkSize, audioDataSize)
                }

                // Let the SDK's continuation block know the stream is over.

                continuation.finish()
            }
        }

        return audioStream
    }
```

This function returns an `AsyncThrowingStream<TranscribeStreamingClientTypes.AudioStream, Error>`. This is a function that asynchronously generates chunks of audio data, yielding them to the caller, until there's no audio left to process.

The function begins by creating a Foundation `URL` from the path of the audio file. Then it reads the audio into a `Data` object (to support larger audio files, this would need to be changed to load the

audio from disk in chunks). The size of each audio chunk to send to the SDK is calculated so it will hold 125 milliseconds of audio, and the total size of the audio file in bytes is obtained.

The audio stream is generated by iterating over the audio data, taking the next chunk of audio and creating a `TranscribeStreamingClientTypes.AudioStream.audioevent` that represents it. The event is sent to the SDK using the continuation object's `yield()` function. The yield result is checked to see if any problems occurred, such as the event being dropped because the event queue is full.

This continues until the last chunk of audio is sent; then the continuation's `finish()` function is executed to let the SDK know the file has been fully transmitted.

Transcribing audio

Transcription is handled by the `transcribe()` function:

```
   /// Run the transcription process.
   ///
   /// - Throws: An error from `TranscribeError`.
   func transcribe(encoding: TranscribeStreamingClientTypes.MediaEncoding) async
throws {
       // Create the Transcribe Streaming client.

       let client = TranscribeStreamingClient(
           config: try await
TranscribeStreamingClient.TranscribeStreamingClientConfig(
               region: region
           )
       )

       // Start the transcription running on the audio stream.

       let output = try await client.startStreamTranscription(
           input: StartStreamTranscriptionInput(
               audioStream: try await createAudioStream(),
               languageCode: TranscribeStreamingClientTypes.LanguageCode(rawValue:
lang),
               mediaEncoding: encoding,
               mediaSampleRateHertz: sampleRate
           )
       )
```

```
        // Iterate over the events in the returned transcript result stream.
        // Each `transcriptevent` contains a list of result fragments which
        // need to be concatenated together to build the final transcript.
        for try await event in output.transcriptResultStream! {
            switch event {
            case .transcriptevent(let event):
            for result in event.transcript?.results ?? [] {
                guard let transcript = result.alternatives?.first?.transcript else {
                    continue
                }

                // If showing partial results is enabled and the result is
                // partial, show it. Partial results may be incomplete, and
                // may be inaccurate, with upcoming audio making the
                // transcription complete or by giving more context to make
                // transcription make more sense.

                if (result.isPartial && showPartial) {
                    print("[Partial] \(transcript)")
                }

                // When the complete fragment of transcribed text is ready,
                // print it. This could just as easily be used to draw the
                // text as a subtitle over a playing video, though timing
                // would need to be managed.

                if !result.isPartial {
                    if (showPartial) {
                        print("[Final  ] ", terminator: "")
                    }
                    print(transcript)
                }
            }
            default:
                print("Error: Unexpected message from Amazon Transcribe:")
            }
        }
    }
```

This function first looks at the value of the **--format** option passed into the program on the command line and prepares a constant of type `TranscribeStreamingClientTypes.MediaEncoding` that indicates the format of the incoming audio. Then it calls `client.startStreamTranscription(input:)` to start the

transcription process. The audio stream is specified by a function named `createAudioStream()`, which is described below.

The event stream returned by `startStreamTranscription(input:)` is monitored using a `for await` loop. Each `transcriptevent` is handled by pulling the first available transcription from the result stream. If the transcript is flagged as partial and the user specified the **--show-partial** option on the command line, the partial output is printed to the console.

If it's a completed transcription of a section of the audio, the transcription is output to the screen. The importance of checking the value of the result's `isPartial` property is simple: as chunks of audio are processed, they may contain partial words that need to be completed by referring to other chunks. Similarly, if a transcription's certainty is low, it might be higher if subsequent chunks provide additional context. For example, if the transcription includes the word "its," the following chunk may help determine if the word should actually be "it's" instead.

Running the example

If you download and build the complete example, you can run it using the **tsevents** executable. For example, if you have a 44,100Hz audio file named `audio-sample.flac`, you can process it with the command:

```
$ tsevents --path audio-sample.flac --format flac --sample-rate 44100
```

If the language of the audio file isn't US English, you can specify the file's language using the **--lang** option. For example, for modern Arabic, you can use:

```
$ tsevents --path audio-sample.flac --format flac --sample-rate 44100 --lang ar-SA
```

For complete usage information, simply run the command **tsevents --help**.

Additional information

- Amazon Transcribe Developer Guide
- Amazon Transcribe API Reference

Presigned URLs and requests with AWS SDK for Swift

Overview

You can presign certain AWS API operations in advance of their use to let the request be used at a later time without the need to provide credentials. With a presigned URL, the owner of a resource can grant an unauthorized person access to a resource for a limited time by simply sending the other user a presigned URL to use the resource.

Presigning basics

The AWS SDK for Swift provides functions that create presigned URLs or requests for each of the service actions that support presigning. These functions take as their `input` parameter the same input struct used by the unsigned action, plus an expiration time that restricts how long the presigned request will be valid and usable.

For example, to create a presigned request for the Amazon S3 action `GetObject`:

```
let getInput = GetObjectInput(
    bucket: bucket,
    key: key
)

let presignedRequest: URLRequest
do {
    presignedRequest = try await s3Client.presignedRequestForGetObject(
        input: getInput,
        expiration: TimeInterval(5 * 60)
    )
} catch {
    throw TransferError.signingError
}
```

This first creates a `GetObjectInput` struct to identify the object to retrieve, then creates a Foundation `URLRequest` with the presigned request by calling the SDK for Swift function `presignedRequestForGetObject(input: expiration:)`, specifying the input struct and a five-minute expiration period. The resulting request can then be sent by anyone, up to five minutes after the request was created. The codebase that sends the request can be in a different application, and even written in a different programming language.

Advanced presigning configuration

The SDK's input structs used to pass options into presignable operations have two methods you can call to generate presigned requests or URLs. For example, the `PutObjectInput` struct has the methods `presign(config: expiration:)` and `presignURL(config: expiration:)`. These are useful if you need to apply to the operation a configuration other than the one used when initializing the service client.

In this example, the `AsyncHTTPClient` package from Apple's `swift-server` project is used to create an `HTTPClientRequest`, which in turn is used to send the file to Amazon S3. A custom configuration is created that lets the SDK make up to six attempts to send an object to Amazon S3:

```
        let fileData = try Data(contentsOf: fileURL)
        let dataStream = ByteStream.data(fileData)
        let presignedURL: URL

        // Create a presigned URL representing the `PutObject` request that
        // will upload the file to Amazon S3. If no URL is generated, a
        // `TransferError.signingError` is thrown.

        let putConfig = try await S3Client.S3ClientConfig(
            maxAttempts: 6,
            region: region
        )

        do {
            let url = try await PutObjectInput(
                body: dataStream,
                bucket: bucket,
                key: fileName
            ).presignURL(
                config: putConfig,
                expiration: TimeInterval(10 * 60)
            )

            guard let url = url else {
                throw TransferError.signingError
            }
            presignedURL = url
        } catch {
            throw TransferError.signingError
        }
```

```
        // Send the HTTP request and upload the file to Amazon S3.

        var request = HTTPClientRequest(url: presignedURL.absoluteString)
        request.method = .PUT
        request.body = .bytes(fileData)

        _ = try await HTTPClient.shared.execute(request, timeout: .seconds(5*60))
```

This creates a new `S3ClientConfig` struct with the value of `maxAttempts` set to 6, then creates a `PutObjectInput` struct which is used to generate an URL that's presigned using the custom configuration. The presigned URL is a standard Foundation `URL` object.

To use the `AsyncHTTPClient` package to send the data to Amazon S3, an `HTTPClientRequest` is created using the presigned URL as the destination URL. The request's method is set to `.PUT`, indicating that it's an upload request. Then the request body is set to the contents of `fileData`, which contains the `Data` to be sent to Amazon S3.

Finally, the request is executed using the shared `HTTPClient` managed by `AsyncHTTPClient`.

Integrating with Apple platforms using the AWS SDK for Swift

Software running on any of Apple's platforms (macOS, iOS, iPadOS, tvOS, visionOS, or watchOS) may wish to integrate into the Apple ecosystem. This may include, for example, letting the user authenticate to access AWS services using Sign In With Apple.

Adding the AWS SDK for Swift to an existing Xcode project

If you have an existing Xcode project, you can add the SDK for Swift to it. Open your project's main configuration pane and choose the **Swift Packages** tab at the right end of the tab bar. The following image shows how to do this for an Xcode project called "Supergame," which will use Amazon S3 to get game data from a server.

Swift Packages tab in Xcode

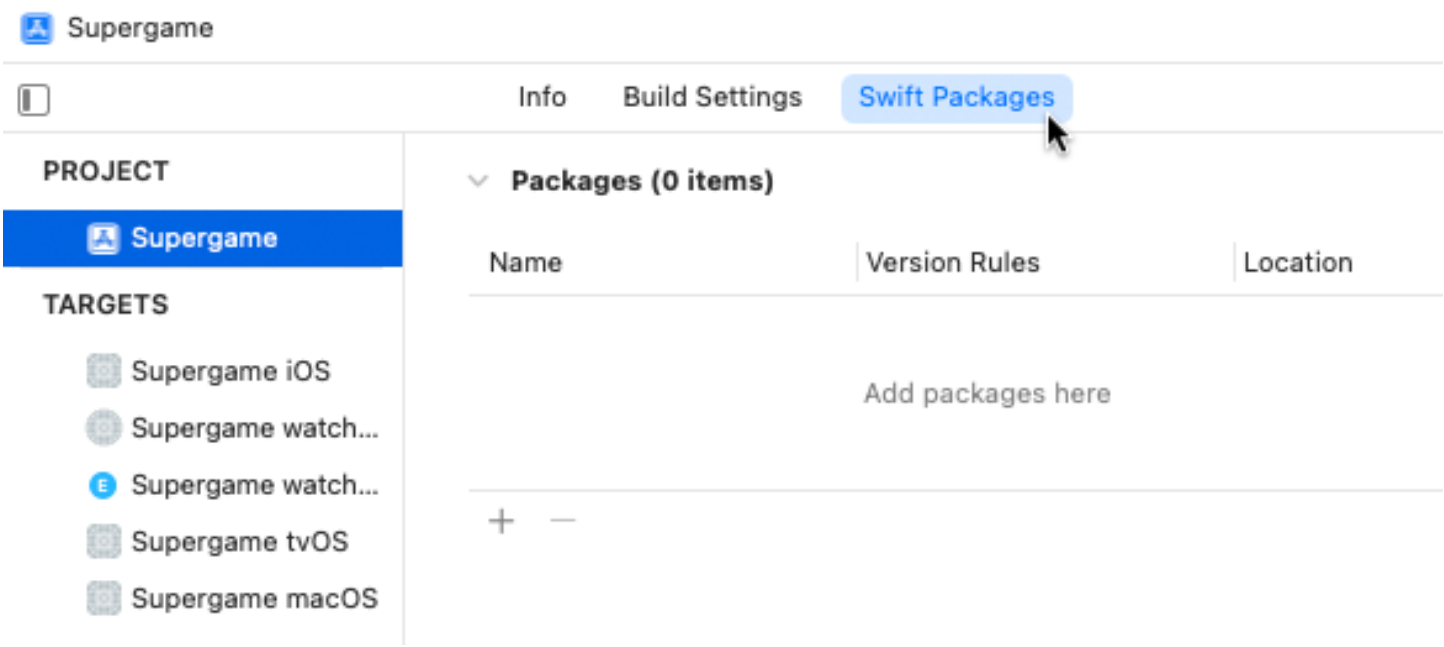

This shows a list of the Swift packages currently in use by your project. If you haven't added any Swift packages, the list will be empty, as shown in the preceding image. To add the AWS SDK for Swift package to your project, choose the **+** button under the package list.

Find and select packages to import

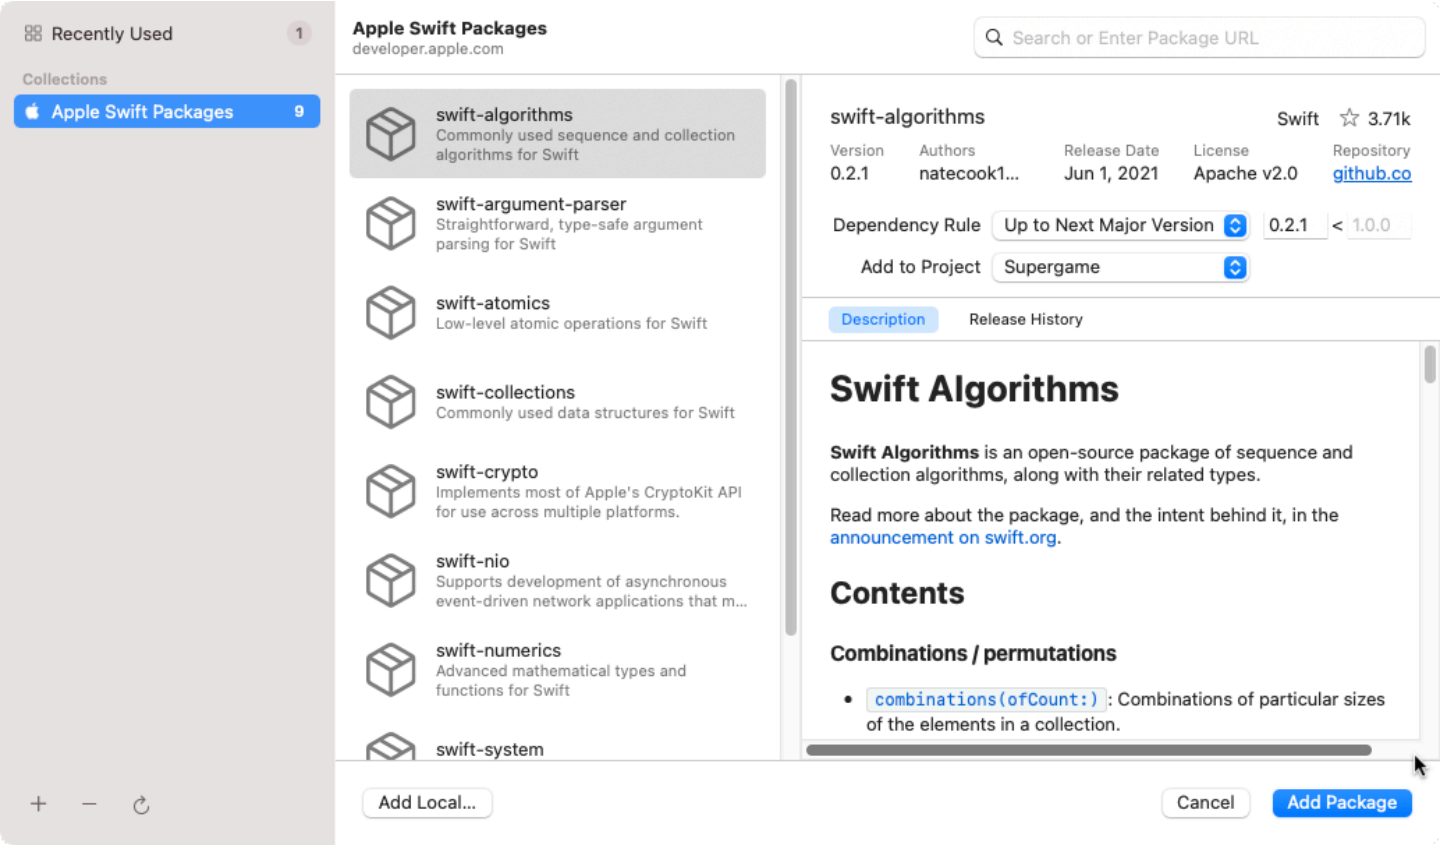

Next, specify the package or packages to add to your project. You can choose from standard Apple-provided packages or enter the URL of a custom package in the search box at the top of the window. Enter the URL of the AWS SDK for Swift as follows: `https://github.com/awslabs/aws-sdk-swift.git` .

After you enter the SDK URL, you can configure version requirements and other options for the SDK package import.

Configure dependency rules for the SDK for Swift package

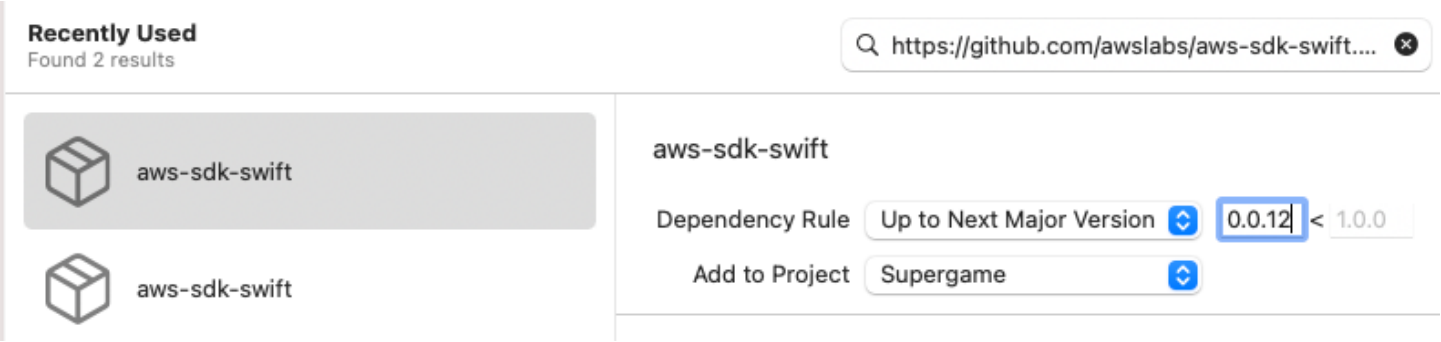

Configure the dependency rule. Make sure that **Add to Project** is set to your project — "Supergame" in this case — and choose **Add Package**. You will see a progress bar while the SDK and all its dependencies are processed and retrieved.

Fetching the AWS SDK for Swift package and its product list

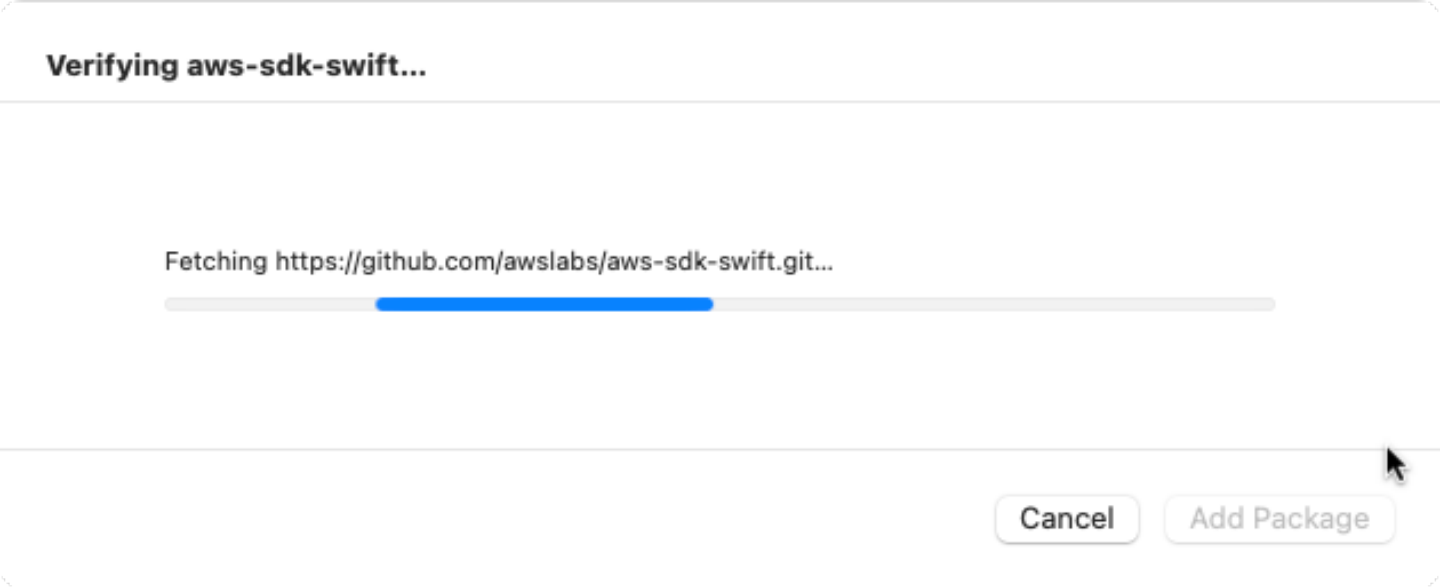

Next, select specific *products* from the AWS SDK for Swift package to include in your project. Each product is generally one AWS API or service. Each package is listed by package name, starting with `AWS` and followed by the shorthand name of the service or toolkit.

For the Supergame project, select `AWSS3`, `AWSDynamoDB`, and `AWSGameLift`. Assign them to the correct target (iOS in this example), and choose **Add Package**.

Choose package products for specific AWS services and toolkits

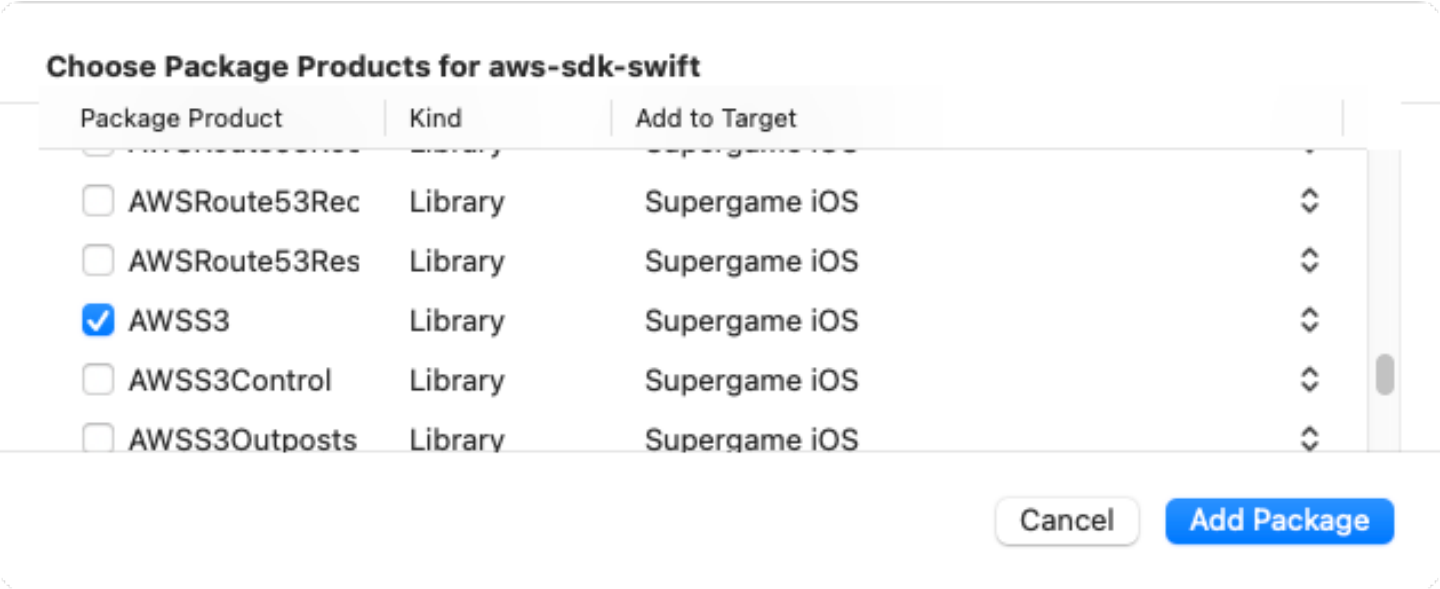

Your project is now configured to import the AWS SDK for Swift package and to include the desired APIs in the build for that target. To see a list of the AWS libraries, open the target's **General** tab and scroll down to **Frameworks, Libraries, and Embedded Content**.

AWS SDK for Swift libraries in the Xcode target

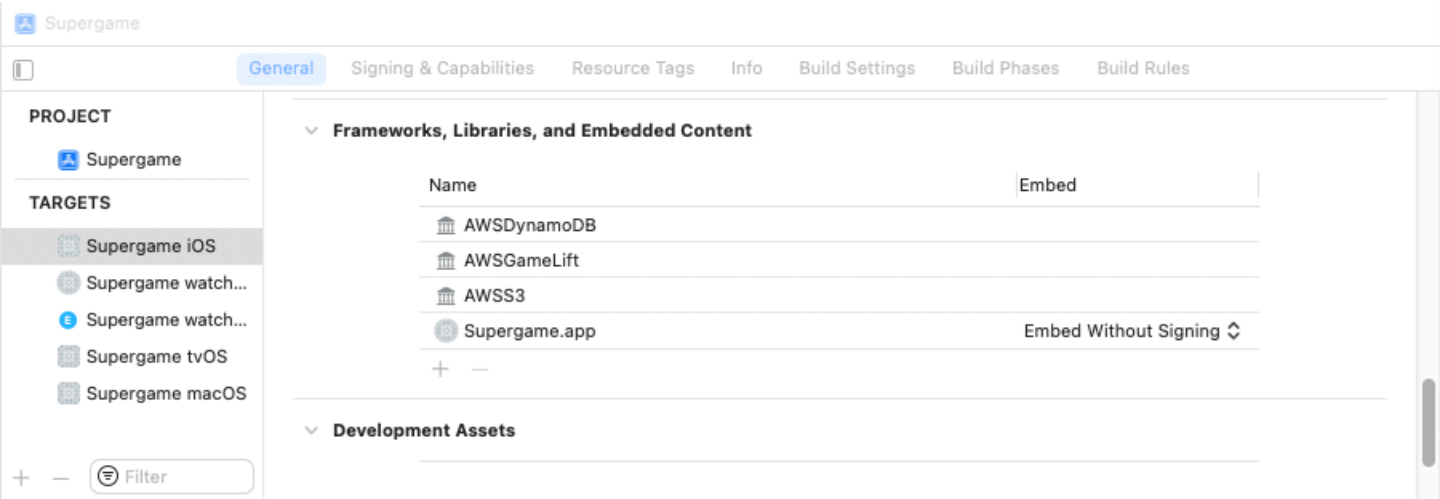

If your project is a multi-platform project, you also need to add the AWS libraries to the other targets in your project. For each platform's target, navigate to the **Frameworks, Libraries, and Embedded Content** section under the **General** tab and choose **+** to open the library picker window.

Then, you can scroll to find and select all of the needed libraries and choose **Add** to add them all at once. Alternatively, you can search for and select each library, then choose **Add** to add them to the target one at a time.

Find and add SDK for Swift libraries using the Xcode library picker window

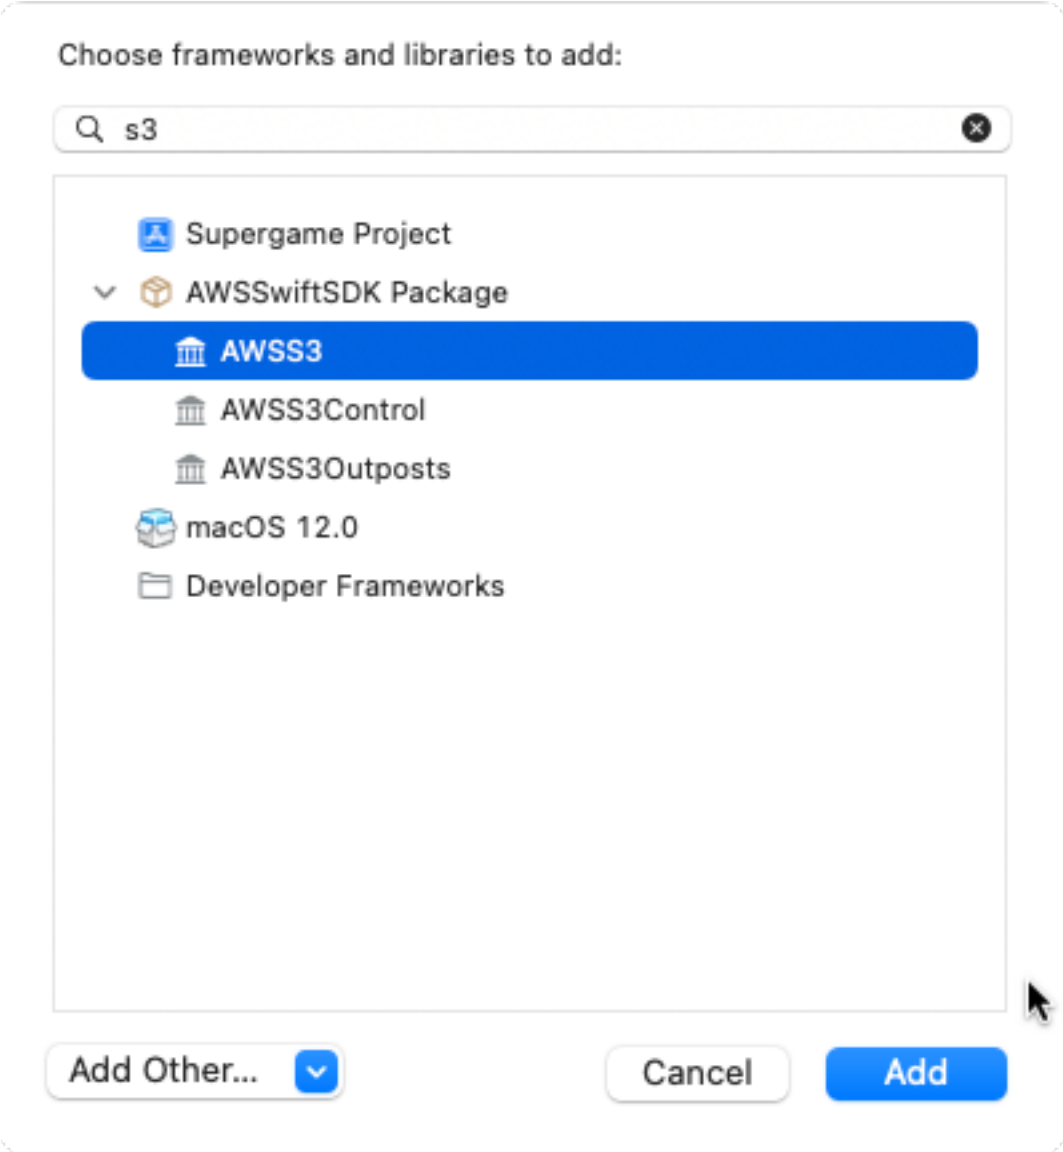

You're now ready to import the libraries and any needed dependencies into individual Swift source code files and start using the AWS services in your project. Build your project by using the Xcode **Build** option in the **Product** menu.

Authenticating using Sign In With Apple

One convenient way for users to sign into AWS while using your application by adding support for Sign In With Apple. This allows your users to access AWS services using their Apple ID and either TouchID or FaceID.

Adding Sign In With Apple support to your app requires planning and configuration of services:

1. Set up your application in Apple's "Certificates, Identifiers & Profiles" dashboard, or configure Sign In With Apple using Xcode.
2. Set up the application's security and authentication configuration in the AWS Management Console.
3. Add the Sign In With Apple authentication method to your application.

> **Note**
>
> When setting up Sign In With Apple as a way to authenticate to use AWS services, the JWT audience should always be the same as your application's bundle ID as configured in Xcode (or the `Info.plist` file) and the Apple developer portal.

The rest of this section discusses the process of setting up and using Sign In With Apple for AWS authentication. The example authenticates to AWS using Sign In With Apple, then lists the user's Amazon Simple Storage Service (Amazon S3) buckets in a SwiftUI view. The example works on macOS, iOS, and iPadOS and is shown in part during the walkthrough below. The complete example is on GitHub.

Configuring the app for Sign in With Apple

If you use Xcode, you can enable Sign In With Apple in the options for your application's main target.

1. Verify that your team name and bundle identifier are correct for the application. Be sure to use Apple's standard reverse-URI notation, such as `com.example.buckets`.
2. Click the **+ Capability** button to open a capability picker. Add Sign In With Apple to your target's capabilities.
3. Click the **Automatically manage signing** checkbox to enable automatic signing of your application. This also configures your application on the Apple developer portal.
4. Under **App Sandbox**, enable **Outgoing Connections (Client)** to have your signed application request network access permission.

If you don't use Xcode, or you prefer to configure the app by hand, you can do so by following the instructions on Apple's developer website.

Configuring AWS services for your application

Next, open the IAM Management Console to configure IAM to support authentication using the JSON Web Token (JWT) returned by Sign In With Apple after a sign in request.

1. In the sidebar, click **Identity providers**, then look to see if there's already a provider for `appleid.apple.com`. If there is, click on it to add your new application's bundle ID as a new audience. Otherwise, click **Add provider** to create a new identity provider with the following configuration:
   - Set the provider type to **OpenID Connect**.
   - Set the **Provider URL** to `https://appleid.apple.com`.
   - Set the audience to match your application's bundle ID, such as `com.example.buckets`.
2. If you don't already have a permissions policy set up to grant the permissions needed by your application (and no permissions it doesn't need):
   a. Click **Policies** in the IAM Management Console sidebar, then click **Create policy**.
   b. Add the permissions your application needs, or directly edit the JSON. The JSON for this example application looks like this:

      JSON

      ```
      {
          "Version":"2012-10-17",
          "Statement": [
              {
                  "Effect": "Allow",
                  "Action": "s3:ListAllMyBuckets",
                  "Resource": ["arn:aws-eusc:s3:::*"]
              }
          ]
      }
      ```

      This policy allows your application to get a list of your Amazon S3 buckets.

   c. On the **Review and create** page, set the **Policy name** to be a unique name to identify your policy, such as `example-buckets-app-policy`.
   d. Enter a helpful description for your policy, such as "Application permissions for the Sign In With Apple buckets example".
   e. Click **Create policy**.

3. Create a new role for your application:
   a. Click **Roles** in the IAM Management Console sidebar, then click **Create role**.
   b. Set the trusted entity type to **Web identity**.
   c. Under **Web identity**, select the Apple ID identity provider you created above.
   d. Select your application's bundle ID as the **Audience**.
   e. Under **Permissions policies**, choose the policy you created above.
   f. On the **Name review, and create** page, set a unique name for your new role, such as `example-buckets-app-role`, and enter a description for the role, such as "Permissions for the Buckets example."
   g. Review the **Trust policy** generated by the console. The JSON should resemble the following:

      JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "sts:AssumeRoleWithWebIdentity",
            "Principal": {
                "Federated": "arn:aws-eusc:iam::111122223333:oidc-provider/
appleid.apple.com"
            },
            "Condition": {
                "StringEquals": {
                    "appleid.apple.com:aud": [
                        "com.example.buckets"
                    ]
                }
            }
        }
    ]
}
```

      This indicates that the Apple ID identity provider used to process Sign In With Apple requests should be allowed to use the role if the web token's aud property (the audience) matches the application's bundle ID.

Adding Sign In With Apple support to your application

Once both Apple and AWS know about your application and that Sign In With Apple should be allowed to authenticate the user for the desired services, the next step is to add a Sign In With Apple button and its supporting code to the application.

This section explains how an application can provide a Sign In With Apple option for authentication, using Apple's SwiftUI and Authentication Services libraries.

Adding a "Sign In With Apple" button

To include a Sign In With Apple button, the `AuthenticationServices` module needs to be imported along with SwiftUI:

```
import SwiftUI
import AuthenticationServices
```

In your SwiftUI sign-in view, embed a Sign In With Apple button. The `SignInWithAppleButton` is used to trigger and manage the Sign In With Apple process. This button calls a completion handler with a `Result` object, which indicates whether or not a valid Apple ID was authenticated:

```
                    // Show the "Sign In With Apple" button, using the
                    // `.continue` mode, which allows the user to create
                    // a new ID if they don't already have one. When SIWA
                    // is complete, the view model's `handleSignInResult()`
                    // function is called to turn the JWT token into AWS
                    // credentials.
                    SignInWithAppleButton(.continue) { request in
                        request.requestedScopes = [.email, .fullName ]
                    } onCompletion: { result in
                        Task {
                            do {
                                try await viewModel.handleSignInResult(result)
                            } catch BucketsAppError.signInWithAppleCanceled {
                                // The "error" is actually Sign In With Apple being
                                // canceled by the user, so end the sign in
                                // attempt.
                                return
                            } catch let error as BucketsAppError {
                                // Handle AWS errors.
                                viewModel.error = error
                                return
```

```
                            }
                        }
                    }
```

This example `SignInWithAppleButton` uses a function named `handleSignInResult()` as its completion handler. This function is passed a `Result` that contains an `ASAuthorization` object if Sign In With Apple succeeded. If sign in failed or was canceled, the `Result` contains an `Error` instead. If the error actually indicates that sign in was canceled by the user, the sign in attempt is ended. Actual errors are stored for display by a SwiftUI alert sheet.

> **Note**
>
> By default, the token returned by Sign In With Apple expires one day from its creation time.

Processing the JSON Web Token

A number of variables are used by the example application to store information related to the user and their sign-in session:

```
    /// The unique string assigned by Sign In With Apple for this login
    /// session. This ID is valid across application launches until it
    /// is signed out from Sign In With Apple.
    var userID = ""

    /// The user's email address.
    ///
    /// This is only returned by SIWA if the user has just created
    /// the app's SIWA account link. Otherwise, it's returned as `nil`
    /// by SIWA and must be retrieved from local storage if needed.
    var email = ""

    /// The user's family (last) name.
    ///
    /// This is only returned by SIWA if the user has just created
    /// the app's SIWA account link. Otherwise, it's returned as `nil`
    /// by SIWA and must be retrieved from local storage if needed.
    var familyName = ""

    /// The user's given (first) name.
    ///
    /// This is only returned by SIWA if the user has just created
```

```
    /// the app's SIWA account link. Otherwise, it's returned as `nil` by SIWA
    /// and must be retrieved from local storage if needed.
    var givenName = ""

    /// The AWS account number provided by the user.
    var awsAccountNumber = ""

    /// The AWS IAM role name given by the user.
    var awsIAMRoleName = ""

    /// The credential identity resolver created by the AWS SDK for
    /// Swift. This resolves temporary credentials using
    /// `AssumeRoleWithWebIdentity`.
    var identityResolver: STSWebIdentityAWSCredentialIdentityResolver? = nil
```

These are used to record the AWS account information (the AWS account number and the name of the IAM role to use, as entered by the user), as well as the user's information returned by Sign In With Apple. The `STSWebIdentityAWSCredentialIdentityResolver` is used to convert the JWT token into valid, temporary AWS credentials when creating a service client object.

The `handleSignInResult()` function looks like this:

```
    /// Called by the Sign In With Apple button when a JWT token has
    /// been returned by the Sign In With Apple service. This function
    /// in turn handles fetching AWS credentials using that token.
    ///
    /// - Parameters:
    ///   - result: The Swift `Result` object passed to the Sign In
    ///     With Apple button's `onCompletion` handler. If the sign
    ///     in request succeeded, this contains an `ASAuthorization`
    ///     object that contains the Apple ID sign in information.
    func handleSignInResult(_ result: Result<ASAuthorization, Error>) async throws {
        switch result {
        case .success(let auth):
            // Sign In With Apple returned a JWT identity token. Gather
            // the information it contains and prepare to convert the
            // token into AWS credentials.

            guard let credential = auth.credential as?
 ASAuthorizationAppleIDCredential,
                  let webToken = credential.identityToken,
                  let tokenString = String(data: webToken, encoding: .utf8)
            else {
```

```
                throw BucketsAppError.credentialsIncomplete
            }

            userID = credential.user

            // If the email field has a value, set the user's recorded email
            // address. Otherwise, keep the existing one.
            email = credential.email ?? self.email

            // Similarly, if the name is present in the credentials, use it.
            // Otherwise, the last known name is retained.
            if let name = credential.fullName {
                self.familyName = name.familyName ?? self.familyName
                self.givenName = name.givenName ?? self.givenName
            }

            // Use the JWT token to request a set of temporary AWS
            // credentials. Upon successful return, the
            // `credentialsProvider` can be used when configuring
            // any AWS service.

            try await authenticate(withWebIdentity: tokenString)
        case .failure(let error as ASAuthorizationError):
            if error.code == .canceled {
                throw BucketsAppError.signInWithAppleCanceled
            } else {
                throw BucketsAppError.signInWithAppleFailed
            }
        case .failure:
            throw BucketsAppError.signInWithAppleFailed
        }

        // Successfully signed in. Fetch the bucket list.
        do {
            try await self.getBucketList()
        } catch {
            throw BucketsAppError.bucketListMissing
        }
    }
```

If sign in is successful, the `Result` provides details about the authentication returned by Sign In With Apple. The authentication credential contains a JSON Web Token, as well as a user ID which is unique for the current session.

If the authentication represents a new link between Sign In With Apple and this application, it may contain the user's email address and full name. If it does, this function retrieves and stores that information. The email address and user name *will never be provided again* by Sign In With Apple.

> **Note**
>
> Sign In With Apple only includes personally identifiable information (PII) the when the user first associates their Apple ID with your application. For all subsequent connections, Sign In With Apple only provides a unique user ID. If the application needs any of this PII, it's the app's responsibility to securely save it *locally* and *securely*. The example application stores the information in the Keychain.

The JWT is converted to a string, which is passed to a function called `authenticate(withWebIdentity: region:)` to actually create the web identity resolver.

Creating the web identity credential identity resolver

The `authenticate(withWebIdentity: region:)` function performs credential identity resolution by first writing the token to a local file, then creating an object of type `STSWebIdentityAWSCredentialIdentityResolver`, specifying the stored web identity token file when doing so. This AWS Security Token Service (AWS STS) credential identity resolver is used when creating service clients.

After authenticating with Sign In With Apple, a function named `saveUserData()` is called to securely store the user's information in the Keychain. This lets the sign in screen automatically fill in the form fields, and lets the application remember the user's email address and name if available.

```
    /// Convert the given JWT identity token string into the temporary
    /// AWS credentials needed to allow this application to operate, as
    /// specified using the Apple Developer portal and the AWS Identity
    /// and Access Management (IAM) service.
    ///
    /// - Parameters:
    ///   - tokenString: The string version of the JWT identity token
    ///     returned by Sign In With Apple.
    ///   - region: An optional string specifying the AWS Region to
    ///     access. If not specified, "us-east-1" is assumed.
    func authenticate(withWebIdentity tokenString: String,
                      region: String = "us-east-1") async throws {
```

```
        // If the role is empty, pass `nil` to use the default role for
        // the user.

        let roleARN = "arn:aws:iam::\(awsAccountNumber):role/\(awsIAMRoleName)"

        // Use the AWS Security Token Service (STS) action
        // `AssumeRoleWithWebIdentity` to convert the JWT token into a
        // set of temporary AWS credentials. The first step: write the token
        // to disk so it can be used by the
        // `STSWebIdentityAWSCredentialIdentityResolver`.

        let tokenFileURL = createTokenFileURL()
        let tokenFilePath = tokenFileURL.path
        do {
            try tokenString.write(to: tokenFileURL, atomically: true, encoding: .utf8)
        } catch {
            throw BucketsAppError.tokenFileError()
        }

        // Create an identity resolver that uses the JWT token received
        // from Apple to create AWS credentials.

        do {
            identityResolver = try STSWebIdentityAWSCredentialIdentityResolver(
                region: region,
                roleArn: roleARN,
                roleSessionName: "BucketsExample",
                tokenFilePath: tokenFilePath
            )
        } catch {
            throw BucketsAppError.assumeRoleFailed
        }

        // Save the user's data securely to local storage so it's available
        // in the future.
        //
        // IMPORTANT: Any potential Personally Identifiable Information _must_
        // be saved securely, such as by using the Keychain or an appropriate
        // encrypting technique.

        saveUserData()
    }
```

> **Important**
>
> The token is written to disk since the SDK expects to load the token from a file. This token file *must remain in place* until you have finished using the service client. When you're done using the client, delete the token file.

Creating a service client using the web identity credential identity resolver

To access an AWS service, create a client configuration object that includes the `awsCredentialIdentityResolver` property. This property's value should be the web identity credential identity resolver created by the `authorize(withWebIdentity: region:)` function:

```
    /// Fetches a list of the user's Amazon S3 buckets.
    ///
    /// The bucket names are stored in the view model's `bucketList`
    /// property.
    func getBucketList() async throws {
        // If there's no identity resolver yet, return without doing anything.
        guard let identityResolver = identityResolver else {
            return
        }

        // Create an Amazon S3 client configuration that uses the
        // credential identity resolver created from the JWT token
        // returned by Sign In With Apple.
        let config = try await S3Client.S3ClientConfig(
            awsCredentialIdentityResolver: identityResolver,
            region: "us-east-1"
        )
        let s3 = S3Client(config: config)

        let output = try await s3.listBuckets(
            input: ListBucketsInput()
        )

        guard let buckets = output.buckets else {
            throw BucketsAppError.bucketListMissing
        }

        // Add the names of all the buckets to `bucketList`. Each
        // name is stored as a new `IDString` for use with the SwiftUI
```

```
        // `List`.
        for bucket in buckets {
            self.bucketList.append(IDString(bucket.name ?? "<unknown>"))
        }
    }
```

This function creates an `S3Client` configured to use our web identity credential identity resolver. That client is then used to fetch a list of Amazon S3 buckets. The buckets' names are then added to the bucket list. The SwiftUI `List` view automatically refreshes to display the newly-added bucket names.

Working with AWS services using the AWS SDK for Swift

This chapter contains information about how to work with AWS services by using the SDK for Swift.

Contents

- Working with Amazon S3 using the AWS SDK for Swift
  - Multipart Amazon S3 uploads using AWS SDK for Swift
    - Overview
    - The multipart upload process
    - Starting a multipart upload
    - Uploading the parts
    - Completing a multipart upload
    - Additional information
  - Binary streaming with AWS SDK for Swift
    - Overview
    - Streaming incoming data
    - Streaming outgoing data
  - Data integrity protection with checksums
    - Upload an object
      - Use a pre-calculated checksum value
      - Multipart uploads
    - Download an object

Working with Amazon S3 using the AWS SDK for Swift

Your main interface to the Amazon Simple Storage Service for the SDK for Swift is the `S3Client`. Use the `S3Client` like other service clients in the SDK to make requests to Amazon S3.

Resources to help you use the SDK for Swift with S3 are:

- The AWS SDK for Swift API reference for S3.
- The S3 service User Guide and service API reference.

The following topics present guided code examples for select SDK for Swift APIs that work with S3.

Topics

- Multipart Amazon S3 uploads using AWS SDK for Swift
- Binary streaming with AWS SDK for Swift
- Data integrity protection with checksums

Multipart Amazon S3 uploads using AWS SDK for Swift

Overview

Multipart upload provides a way to upload a single large object in multiple parts, which are combined by Amazon S3 when the upload or copy is complete. This can improve performance because it lets multiple parts be uploaded in parallel, allows pause and resume of uploads, and provides better performance when handling errors since only the part or parts that failed need to be re-uploaded. See Uploading and copying objects using multipart upload in *Amazon Simple Storage Service User Guide* for more information.

The multipart upload process

> **Note**
>
> This section is only a brief review of the multipart upload process. For a more detailed look at the process, including more information about limitations and additional capabilities, see Multipart upload process in *Amazon Simple Storage Service User Guide*.

Fundamentally, multipart upload involves three steps: initiating the upload, uploading the object's parts, and completing the upload once all of the parts have been uploaded. Parts can be uploaded in any order. When uploading each part, you assign it a part number corresponding to the order in which it belongs in the fully assembled file. When received by Amazon S3, each part is assigned an ETag, which is returned to you. The client and the server use the ETag and part number to match the data to the position it occupies in the object during the process of reassembling the content into the complete uploaded object.

This example can be found in its entirety on GitHub.

Starting a multipart upload

The first step is to call the SDK for Swift function `S3Client.createMultipartUpload(input:)` with the name of the bucket to store the object into and the key (name) for the new object.

```
    /// Start a multi-part upload to Amazon S3.
    /// - Parameters:
    ///   - bucket: The name of the bucket to upload into.
    ///   - key: The name of the object to store in the bucket.
    ///
    /// - Returns: A string containing the `uploadId` of the multi-part
    ///   upload job.
    ///
    /// - Throws:
    func startMultipartUpload(client: S3Client, bucket: String, key: String) async
throws -> String {
        let multiPartUploadOutput: CreateMultipartUploadOutput

        // First, create the multi-part upload.

        do {
            multiPartUploadOutput = try await client.createMultipartUpload(
                input: CreateMultipartUploadInput(
                    bucket: bucket,
                    key: key
                )
            )
        } catch {
            throw TransferError.multipartStartError
        }

        // Get the upload ID. This needs to be included with each part sent.

        guard let uploadID = multiPartUploadOutput.uploadId else {
            throw TransferError.uploadError("Unable to get the upload ID")
        }

        return uploadID
    }
```

The output from `createMultipartUpload(input:)` is a string which uniquely identifies the upload that has been started. This ID is used for each call to `uploadPart(input:)` to match the

uploaded part to a particular upload. Since each upload has a unique ID, you can have multiple multipart uploads in progress at the same time.

Uploading the parts

After creating the multipart upload, upload the numbered parts by calling `S3Client.uploadPart(input:)` once for each part. Each part (except the last) must be at least 5 MB in size.

```
            for partNumber in 1...blockCount {
                let data: Data
                let startIndex = UInt64(partNumber - 1) * UInt64(blockSize)

                // Read the block from the file.

                data = try readFileBlock(file: fileHandle, startIndex: startIndex,
 size: blockSize)

                // Upload the part to Amazon S3 and append the `CompletedPart` to
                // the array `completedParts` for use after all parts are uploaded.

                let completedPart = try await uploadPart(
                    client: s3Client, uploadID: uploadID,
                    bucket: bucket, key: fileName,
                    partNumber: partNumber, data: data
                )
                completedParts.append(completedPart)

                let percent = Double(partNumber) / Double(blockCount) * 100
                print(String(format: " %.1f%%", percent))
            }

            .
            .
            .

    /// Upload the specified data as part of an Amazon S3 multi-part upload.
    ///
    /// - Parameters:
    ///   - client: The S3Client to use to upload the part.
    ///   - uploadID: The upload ID of the multi-part upload to add the part to.
    ///   - bucket: The name of the bucket the data is being written to.
```

```
    ///   - key: A string giving the key which names the Amazon S3 object the file is
being added to.
    ///   - partNumber: The part number within the file that the specified data
represents.
    ///   - data: The data to send as the specified object part number in the object.
    ///
    /// - Throws: `TransferError.uploadError`
    ///
    /// - Returns: A `CompletedPart` object describing the part that was uploaded.
    ///   contains the part number as well as the ETag returned by Amazon S3.
    func uploadPart(client: S3Client, uploadID: String, bucket: String,
                    key: String, partNumber: Int, data: Data)
                    async throws -> S3ClientTypes.CompletedPart {
        let uploadPartInput = UploadPartInput(
            body: ByteStream.data(data),
            bucket: bucket,
            key: key,
            partNumber: partNumber,
            uploadId: uploadID
        )

        // Upload the part.
        do {
            let uploadPartOutput = try await client.uploadPart(input: uploadPartInput)

            guard let eTag = uploadPartOutput.eTag else {
                throw TransferError.uploadError("Missing eTag")
            }

            return S3ClientTypes.CompletedPart(
                eTag: eTag,
                partNumber: partNumber
            )
        } catch {
            throw TransferError.uploadError(error.localizedDescription)
        }
    }
```

This example iterates over chunks of the file, reading the data from the source file then uploading it with the corresponding part number. The response includes the ETag assigned to the newly uploaded part. This and the part number are used to create an `S3ClientTypes.CompletedPart` record matching the ETag to the part number, and this record is added to the array of completed

part records. This will be needed when all the parts have been uploaded and it's time to call `S3Client.completeMultipartUpload(input:)`.

Completing a multipart upload

Once all parts have been uploaded, call the `S3Client` function `completeMultipartUpload(input:)`. This takes as input the names of the bucket and key for the object, the upload ID string, and the array that matches each of the object's part numbers to the corresponding ETags, as created in the section above.

```
    /// Complete a multi-part upload using an array of `CompletedMultipartUpload`
    /// objects describing the completed parts.
    ///
    /// - Parameters:
    ///   - client: The S3Client to finish uploading with.
    ///   - uploadId: The multi-part upload's ID string.
    ///   - bucket: The name of the bucket the upload is targeting.
    ///   - key: The name of the object being written to the bucket.
    ///   - parts: An array of `CompletedPart` objects describing each part
    ///     of the upload.
    ///
    /// - Throws: `TransferError.multipartFinishError`
    func finishMultipartUpload(client: S3Client, uploadId: String, bucket: String, key:
 String,
                               parts: [S3ClientTypes.CompletedPart]) async throws {
        do {
            let partInfo = S3ClientTypes.CompletedMultipartUpload(parts: parts)
            let multiPartCompleteInput = CompleteMultipartUploadInput(
                bucket: bucket,
                key: key,
                multipartUpload: partInfo,
                uploadId: uploadId
            )
            _ = try await client.completeMultipartUpload(input: multiPartCompleteInput)
        } catch {
            dump(error)
            throw TransferError.multipartFinishError(error.localizedDescription)
        }
    }
```

Additional information

- Uploading and copying objects using multipart upload

Binary streaming with AWS SDK for Swift

Overview

In the AWS SDK for Swift, binary data can be represented as a stream of data directly from a file or other resource. This is done using the Smithy `ByteStream` type, which represents an abstract stream of bytes.

Streams are handled automatically by the SDK. Incoming data from downloads is accepted using a stream so you don't have to manage receiving and combining multiple inbound chunks of an object yourself. Similarly, you can choose to upload a file using a stream as the data source, which lets you avoid needing to write multipart upload code manually.

The example code in this section can be found in its entirety on GitHub.

Streaming incoming data

The `*Output` structs returned by functions that receive incoming data contain a `body` property whose value is a `ByteStream`. The `ByteStream` returns its data either as a Swift `Data` object or as a Smithy `ReadableStream`, from which you can read the data.

```
    /// Download a file from the specified bucket.
    ///
    /// - Parameters:
    ///   - bucket: The Amazon S3 bucket name to get the file from.
    ///   - key: The name (or path) of the file to download from the bucket.
    ///   - destPath: The pathname on the local filesystem at which to store
    ///     the downloaded file.
    func downloadFile(bucket: String, key: String, destPath: String?) async throws {
        let fileURL: URL

        // If no destination path was provided, use the key as the name to use
        // for the file in the downloads folder.

        if destPath == nil {
            do {
                try fileURL = FileManager.default.url(
```

```
                    for: .downloadsDirectory,
                    in: .userDomainMask,
                    appropriateFor: URL(string: key),
                    create: true
                ).appendingPathComponent(key)
            } catch {
                throw TransferError.directoryError
            }
        } else {
            fileURL = URL(fileURLWithPath: destPath!)
        }

        let config = try await S3Client.S3ClientConfig(region: region)
        let s3Client = S3Client(config: config)

        // Create a `FileHandle` referencing the local destination. Then
        // create a `ByteStream` from that.

        FileManager.default.createFile(atPath: fileURL.path, contents: nil, attributes:
 nil)
        let fileHandle = try FileHandle(forWritingTo: fileURL)

        // Download the file using `GetObject`.

        let getInput = GetObjectInput(
            bucket: bucket,
            key: key
        )

        do {
            let getOutput = try await s3Client.getObject(input: getInput)

            guard let body = getOutput.body else {
                throw TransferError.downloadError("Error: No data returned for
 download")
            }

            // If the body is returned as a `Data` object, write that to the
            // file. If it's a stream, read the stream chunk by chunk,
            // appending each chunk to the destination file.

            switch body {
            case .data:
                guard let data = try await body.readData() else {
```

```
                    throw TransferError.downloadError("Download error")
                }

                // Write the `Data` to the file.

                do {
                    try data.write(to: fileURL)
                } catch {
                    throw TransferError.writeError
                }
                break

            case .stream(let stream as ReadableStream):
                while (true) {
                    let chunk = try await stream.readAsync(upToCount: 5 * 1024 * 1024)
                    guard let chunk = chunk else {
                        break
                    }

                    // Write the chunk to the destination file.

                    do {
                        try fileHandle.write(contentsOf: chunk)
                    } catch {
                        throw TransferError.writeError
                    }
                }

                break
            default:
                throw TransferError.downloadError("Received data is unknown object
type")
            }
        } catch {
            throw TransferError.downloadError("Error downloading the file: \(error)")
        }

        print("File downloaded to \(fileURL.path).")
    }
```

This function builds a Swift URL representing the destination path for the downloaded file, using either the path specified by `destPath`, or by creating a file in the user's `Downloads` directory with the same name as the object being downloaded. Then an Amazon S3 client is created, the file

is created using the Foundation `FileManager` class, and the download is started by calling the `S3Client.getObject(input:)` function.

If the returned `GetObjectOutput` object's body property matches `.data`, the stream's contents are retrieved using the `ByteStream` function `readData()`, and the data is written to the file. If body matches `.stream`, the associated `ReadableStream` is read by repeatedly calling the stream's `readAsync(upToCount:)` function to get chunks of the file, writing each chunk to the file until no chunk is received, at which point the download ends. Any other type of body is unknown, causing an error to throw.

Streaming outgoing data

When sending data to an AWS service, you can use a `ByteStream` to send data too large to reasonably store in memory in its entirety, or data that is being generated in real time.

```
    /// Upload a file to the specified bucket.
    ///
    /// - Parameters:
    ///   - bucket: The Amazon S3 bucket name to store the file into.
    ///   - key: The name (or path) of the file to upload to in the `bucket`.
    ///   - sourcePath: The pathname on the local filesystem of the file to
    ///     upload.
    func uploadFile(sourcePath: String, bucket: String, key: String?) async throws {
        let fileURL: URL = URL(fileURLWithPath: sourcePath)
        let fileName: String

        // If no key was provided, use the last component of the filename.

        if key == nil {
            fileName = fileURL.lastPathComponent
        } else {
            fileName = key!
        }

        let s3Client = try await S3Client()

        // Create a FileHandle for the source file.

        let fileHandle = FileHandle(forReadingAtPath: sourcePath)
        guard let fileHandle = fileHandle else {
            throw TransferError.readError
        }
```

```
        // Create a byte stream to retrieve the file's contents. This uses the
        // Smithy FileStream and ByteStream types.

        let stream = FileStream(fileHandle: fileHandle)
        let body = ByteStream.stream(stream)

        // Create a `PutObjectInput` with the ByteStream as the body of the
        // request's data. The AWS SDK for Swift will handle sending the
        // entire file in chunks, regardless of its size.

        let putInput = PutObjectInput(
            body: body,
            bucket: bucket,
            key: fileName
        )

        do {
            _ = try await s3Client.putObject(input: putInput)
        } catch {
            throw TransferError.uploadError("Error uploading the file: \(error)")
        }

        print("File uploaded to \(fileURL.path).")
    }
```

This function opens the source file for reading using the Foundation `FileHandle` class. The `FileHandle` is then used to create a Smithy `ByteStream` object. The stream is specified as the `body` when setting up the `PutObjectInput`, which is in turn used to upload the file. Since a stream was specified as the `body`, the SDK automatically continues to send the stream's data until the end of the file is reached.

Data integrity protection with checksums

Amazon Simple Storage Service (Amazon S3) provides the ability to specify a checksum when you upload an object. When you specify a checksum, it is stored with the object and can be validated when the object is downloaded.

Checksums provide an additional layer of data integrity when you transfer files. With checksums, you can verify data consistency by confirming that the received file matches the original file. For more information about checksums with Amazon S3, see the Amazon Simple Storage Service User Guide including the supported algorithms.

You have the flexibility to choose the algorithm that best fits your needs and let the SDK calculate the checksum. Alternatively, you can provide a pre-computed checksum value by using one of the supported algorithms.

> **Note**
>
> Beginning with version 1.1.0 of the AWS SDK for Swift, the SDK provides default integrity protections by automatically calculating a `CRC32` checksum for uploads. The SDK calculates this checksum if you don't provide a precalculated checksum value or if you don't specify an algorithm that the SDK should use to calculate a checksum.
> The SDK also provides global settings for data integrity protections that you can set externally, which you can read about in the AWS SDKs and Tools Reference Guide.

We discuss checksums in two request phases: uploading an object and downloading an object.

Upload an object

You upload objects to Amazon S3 with the SDK for Swift by using the `putObject(input:)` function, setting the `checksumAlgorithm` property in the `PutObjectInput` struct to the desired checksum algorithm.

The following code snippet shows a request to upload an object with a `SHA256` checksum. When the SDK sends the request, it calculates the `SHA256` checksum and uploads the object. Amazon S3 stores the checksum with the object.

```
let output = try await s3Client.putObject(
    input: PutObjectInput(
        body: dataStream,
        bucket: "amzn-s3-demo-bucket",
        checksumAlgorithm: .sha256,
        key: "key"
    )
)
```

If you don't provide a checksum algorithm with the request, the checksum behavior varies depending on the version of the SDK that you use as shown in the following table.

Checksum behavior when no checksum algorithm is provided

Swift SDK version	Checksum behavior
earlier than 1.1.0	The SDK doesn't automatically calculate a CRC-based checksum and provide it in the request.
1.1.0 or later	The SDK uses the CRC32 algorithm to calculate the checksum and provides it in the request. Amazon S3 validates the integrity of the transfer by computing its own CRC32 checksum and compares it to the checksum provided by the SDK. If the checksums match, the checksum is saved with the object.

Use a pre-calculated checksum value

A pre-calculated checksum value provided with the request disables automatic computation by the SDK and uses the provided value instead.

The following example shows a request with a pre-calculated SHA256 checksum.

```
let output = try await s3Client.putObject(
    input: PutObjectInput(
        body: dataStream,
        bucket: "amzn-s3-demo-bucket",
        checksumAlgorithm: .sha256,
        checksumSHA256 =
 "cfb6d06da6e6f51c22ae3e549e33959dbb754db75a93665b8b579605464ce299",
        key: "key"
    )
)
```

If Amazon S3 determines the checksum value is incorrect for the specified algorithm, the service returns an error response.

Multipart uploads

You can also use checksums with multipart uploads.

You must specify the checksum algorithm when calling `createMultipartUpload(input:)`, as well as in each call to `uploadPart(input:)`. Also, each part's returned checksum must be included in the list of completed parts passed into `completeMultipartUpload(input:)`.

```
    /// Upload a file to Amazon S3.
    ///
    /// - Parameters:
    ///   - file: The path of the local file to upload to Amazon S3.
    ///   - bucket: The name of the bucket to upload the file into.
    ///   - key: The key (name) to give the object on Amazon S3.
    ///
    /// - Throws: Errors from `TransferError`
    func uploadFile(file: String, bucket: String, key: String?) async throws {
        let fileURL = URL(fileURLWithPath: file)
        let fileName: String

        // If no key was provided, use the last component of the filename.

        if key == nil {
            fileName = fileURL.lastPathComponent
        } else {
            fileName = key!
        }

        // Create an Amazon S3 client in the desired Region.

        let config = try await S3Client.S3ClientConfig(region: region)
        let s3Client = S3Client(config: config)

        print("Uploading file from \(fileURL.path) to \(bucket)/\(fileName).")

        let multiPartUploadOutput: CreateMultipartUploadOutput

        // First, create the multi-part upload, using SHA256 checksums.

        do {
            multiPartUploadOutput = try await s3Client.createMultipartUpload(
                input: CreateMultipartUploadInput(
                    bucket: bucket,
                    checksumAlgorithm: .sha256,
                    key: key
                )
            )
```

```
        } catch {
            throw TransferError.multipartStartError
        }

        // Get the upload ID. This needs to be included with each part sent.

        guard let uploadID = multiPartUploadOutput.uploadId else {
            throw TransferError.uploadError("Unable to get the upload ID")
        }

        // Open a file handle and prepare to send the file in chunks. Each chunk
        // is 5 MB, which is the minimum size allowed by Amazon S3.

        do {
            let blockSize = Int(5 * 1024 * 1024)
            let fileHandle = try FileHandle(forReadingFrom: fileURL)
            let fileSize = try getFileSize(file: fileHandle)
            let blockCount = Int(ceil(Double(fileSize) / Double(blockSize)))
            var completedParts: [S3ClientTypes.CompletedPart] = []

            // Upload the blocks one at as Amazon S3 object parts.

            print("Uploading...")

            for partNumber in 1...blockCount {
                let data: Data
                let startIndex = UInt64(partNumber - 1) * UInt64(blockSize)

                // Read the block from the file.

                data = try readFileBlock(file: fileHandle, startIndex: startIndex,
 size: blockSize)

                let uploadPartInput = UploadPartInput(
                    body: ByteStream.data(data),
                    bucket: bucket,
                    checksumAlgorithm: .sha256,
                    key: key,
                    partNumber: partNumber,
                    uploadId: uploadID
                )

                // Upload the part with a SHA256 checksum.
```

```
                do {
                    let uploadPartOutput = try await s3Client.uploadPart(input:
uploadPartInput)

                    guard let eTag = uploadPartOutput.eTag else {
                        throw TransferError.uploadError("Missing eTag")
                    }
                    guard let checksum = uploadPartOutput.checksumSHA256 else {
                        throw TransferError.checksumError
                    }
                    print("Part \(partNumber) checksum: \(checksum)")

                    // Append the completed part description (including its
                    // checksum, ETag, and part number) to the
                    // `completedParts` array.

                    completedParts.append(
                        S3ClientTypes.CompletedPart(
                            checksumSHA256: checksum,
                            eTag: eTag,
                            partNumber: partNumber
                        )
                    )
                } catch {
                    throw TransferError.uploadError(error.localizedDescription)
                }
            }

            // Tell Amazon S3 that all parts have been uploaded.

            do {
                let partInfo = S3ClientTypes.CompletedMultipartUpload(parts:
completedParts)
                let multiPartCompleteInput = CompleteMultipartUploadInput(
                    bucket: bucket,
                    key: key,
                    multipartUpload: partInfo,
                    uploadId: uploadID
                )
                _ = try await s3Client.completeMultipartUpload(input:
multiPartCompleteInput)
            } catch {
                throw TransferError.multipartFinishError(error.localizedDescription)
            }
```

```
        } catch {
            throw TransferError.uploadError("Error uploading the file: \(error)")
        }

        print("Done. Uploaded as \(fileName) in bucket \(bucket).")
    }
```

Download an object

When you use the `getObject(input:)` method to download an object, the SDK automatically validates the checksum when the `checksumMode` property of the `GetObjectInput` struct is set to `ChecksumMode.enabled`.

The request in the following snippet directs the SDK to validate the checksum in the response by calculating the checksum and comparing the values.

```
let output = GetObject(
    input: GetObjectInput(
        bucket: "amzn-s3-demo-bucket",
        checksumMode: .enabled
        key: "key"
    )
)
```

> **Note**
>
> If the object wasn't uploaded with a checksum, no validation takes place.

SDK for Swift code examples

The code examples in this topic show you how to use the AWS SDK for Swift with AWS.

Basics are code examples that show you how to perform the essential operations within a service.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Scenarios are code examples that show you how to accomplish specific tasks by calling multiple functions within a service or combined with other AWS services.

Some services contain additional example categories that show how to leverage libraries or functions specific to the service.

Services

- Amazon Bedrock examples using SDK for Swift
- Amazon Bedrock Runtime examples using SDK for Swift
- Amazon Cognito Identity examples using SDK for Swift
- Amazon Cognito Identity Provider examples using SDK for Swift
- DynamoDB examples using SDK for Swift
- Amazon EC2 examples using SDK for Swift
- AWS Glue examples using SDK for Swift
- IAM examples using SDK for Swift
- Lambda examples using SDK for Swift
- Amazon RDS examples using SDK for Swift
- Amazon S3 examples using SDK for Swift
- Amazon SNS examples using SDK for Swift
- Amazon SQS examples using SDK for Swift
- AWS STS examples using SDK for Swift
- Amazon Transcribe Streaming examples using SDK for Swift

Amazon Bedrock examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with Amazon Bedrock.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Get started
- Actions

Get started

Hello Amazon Bedrock

The following code example shows how to get started using Amazon Bedrock.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import ArgumentParser
import AWSClientRuntime
import Foundation

import AWSBedrock

struct ExampleCommand: ParsableCommand {
    static var configuration = CommandConfiguration(
        commandName: "ListFoundationModels",
        abstract: """
        This example demonstrates how to retrieve a list of the available
```

```
        foundation models from Amazon Bedrock.
        """,
        discussion: """
        """
    )

    /// Construct a string listing the specified modalities.
    ///
    /// - Parameter modalities: An array of the modalities to list.
    ///
    /// - Returns: A string with a human-readable list of modalities.
    func buildModalityList(modalities: [BedrockClientTypes.ModelModality]?) ->
String {
        var first = true
        var str = ""

        if modalities == nil {
            return "<none>"
        }

        for modality in modalities! {
            if !first {
                str += ", "
            }
            first = false
            str += modality.rawValue
        }

        return str
    }

    /// Construct a string listing the specified customizations.
    ///
    /// - Parameter customizations: An array of the customizations to list.
    ///
    /// - Returns: A string listing the customizations.
    func buildCustomizationList(customizations:
[BedrockClientTypes.ModelCustomization]?) -> String {
        var first = true
        var str = ""

        if customizations == nil {
            return "<none>"
        }
```

```
        for customization in customizations! {
            if !first {
                str += ", "
            }
            first = false
            str += customization.rawValue
        }

        return str
    }

    /// Construct a string listing the specified inferences.
    ///
    /// - Parameter inferences: An array of inferences to list.
    ///
    /// - Returns: A string listing the specified inferences.
    func buildInferenceList(inferences: [BedrockClientTypes.InferenceType]?) ->
 String {
        var first = true
        var str = ""

        if inferences == nil {
            return "<none>"
        }

        for inference in inferences! {
            if !first {
                str += ", "
            }
            first = false
            str += inference.rawValue
        }

        return str
    }

    /// Called by ``main()`` to run the bulk of the example.
    func runAsync() async throws {
        // Always use the Region "us-east-1" to have access to the most models.
        let config = try await BedrockClient.BedrockClientConfig(region: "us-
east-1")
        let bedrockClient = BedrockClient(config: config)
```

```
        let output = try await bedrockClient.listFoundationModels(
            input: ListFoundationModelsInput()
        )

        guard let summaries = output.modelSummaries else {
            print("No models returned.")
            return
        }

        // Output a list of the models with their details.
        for summary in summaries {
            print("==========================================")
            print(" Model ID: \(summary.modelId ?? "<unknown>")")
            print("------------------------------------------")
            print(" Name: \(summary.modelName ?? "<unknown>")")
            print(" Provider: \(summary.providerName ?? "<unknown>")")
            print(" Input modalities: \(buildModalityList(modalities:
 summary.inputModalities))")
            print(" Output modalities: \(buildModalityList(modalities:
 summary.outputModalities))")
            print(" Supported customizations:
 \(buildCustomizationList(customizations: summary.customizationsSupported ))")
            print(" Supported inference types: \(buildInferenceList(inferences:
 summary.inferenceTypesSupported))")
            print("------------------------------------------\n")
        }

        print("\(summaries.count) models available.")
    }
}

/// The program's asynchronous entry point.
@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.runAsync()
        } catch {
            ExampleCommand.exit(withError: error)
        }
    }
```

```
}
```

- For API details, see ListFoundationModels in *AWS SDK for Swift API reference*.

Actions

`ListFoundationModels`

The following code example shows how to use `ListFoundationModels`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSBedrock

        // Always use the Region "us-east-1" to have access to the most models.
        let config = try await BedrockClient.BedrockClientConfig(region: "us-
east-1")
        let bedrockClient = BedrockClient(config: config)

        let output = try await bedrockClient.listFoundationModels(
            input: ListFoundationModelsInput()
        )

        guard let summaries = output.modelSummaries else {
            print("No models returned.")
            return
        }

        // Output a list of the models with their details.
        for summary in summaries {
            print("==========================================")
            print(" Model ID: \(summary.modelId ?? "<unknown>")")
            print("------------------------------------------")
            print(" Name: \(summary.modelName ?? "<unknown>")")
```

```
            print(" Provider: \(summary.providerName ?? "<unknown>")")
            print(" Input modalities: \(buildModalityList(modalities:
summary.inputModalities))")
            print(" Output modalities: \(buildModalityList(modalities:
summary.outputModalities))")
            print(" Supported customizations:
\(buildCustomizationList(customizations: summary.customizationsSupported ))")
            print(" Supported inference types: \(buildInferenceList(inferences:
summary.inferenceTypesSupported))")
            print("-----------------------------------------\n")
        }
```

- For API details, see ListFoundationModels in *AWS SDK for Swift API reference*.

Amazon Bedrock Runtime examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with Amazon Bedrock Runtime.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Amazon Nova
- Amazon Nova Canvas
- Amazon Nova Reel
- Anthropic Claude
- Meta Llama

Amazon Nova

Converse

The following code example shows how to send a text message to Amazon Nova, using Bedrock's Converse API.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

Send a text message to Amazon Nova, using Bedrock's Converse API.

```
// An example demonstrating how to use the Conversation API to send
// a text message to Amazon Nova.

import AWSBedrockRuntime

func converse(_ textPrompt: String) async throws -> String {

    // Create a Bedrock Runtime client in the AWS Region you want to use.
    let config =
        try await BedrockRuntimeClient.BedrockRuntimeClientConfig()
    let client = BedrockRuntimeClient(config: config)

    // Set the model ID.
    let modelId = "amazon.nova-micro-v1:0"

    // Start a conversation with the user message.
    let message = BedrockRuntimeClientTypes.Message(
        content: [.text(textPrompt)],
        role: .user
    )

    // Optionally use inference parameters
    let inferenceConfig =
        BedrockRuntimeClientTypes.InferenceConfiguration(
            maxTokens: 512,
            stopSequences: ["END"],
            temperature: 0.5,
            topp: 0.9
        )

    // Create the ConverseInput to send to the model
    let input = ConverseInput(
        inferenceConfig: inferenceConfig, messages: [message], modelId: modelId)
```

```
    // Send the ConverseInput to the model
    let response = try await client.converse(input: input)

    // Extract and return the response text.
    if case let .message(msg) = response.output {
        if case let .text(textResponse) = msg.content![0] {
            return textResponse
        } else {
            return "No text response found in message content"
        }
    } else {
        return "No message found in converse output"
    }
}
```

- For API details, see Converse in *AWS SDK for Swift API reference*.

ConverseStream

The following code example shows how to send a text message to Amazon Nova, using Bedrock's Converse API and process the response stream in real-time.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

Send a text message to Amazon Nova, using Bedrock's Converse API and process the response stream in real-time.

```
// An example demonstrating how to use the Conversation API to send a text message
// to Amazon Nova and print the response stream

import AWSBedrockRuntime

func printConverseStream(_ textPrompt: String) async throws {
```

```
    // Create a Bedrock Runtime client in the AWS Region you want to use.
    let config =
        try await BedrockRuntimeClient.BedrockRuntimeClientConfig()
    let client = BedrockRuntimeClient(config: config)

    // Set the model ID.
    let modelId = "amazon.nova-lite-v1:0"

    // Start a conversation with the user message.
    let message = BedrockRuntimeClientTypes.Message(
        content: [.text(textPrompt)],
        role: .user
    )

    // Optionally use inference parameters.
    let inferenceConfig =
        BedrockRuntimeClientTypes.InferenceConfiguration(
            maxTokens: 512,
            stopSequences: ["END"],
            temperature: 0.5,
            topp: 0.9
        )

    // Create the ConverseStreamInput to send to the model.
    let input = ConverseStreamInput(
        inferenceConfig: inferenceConfig, messages: [message], modelId: modelId)

    // Send the ConverseStreamInput to the model.
    let response = try await client.converseStream(input: input)

    // Extract the streaming response.
    guard let stream = response.stream else {
        print("No stream available")
        return
    }

    // Extract and print the streamed response text in real-time.
    for try await event in stream {
        switch event {
        case .messagestart(_):
            print("\nNova Lite:")

        case .contentblockdelta(let deltaEvent):
```

```
            if case .text(let text) = deltaEvent.delta {
                print(text, terminator: "")
            }

        default:
            break
        }
    }
}
```

- For API details, see ConverseStream in *AWS SDK for Swift API reference*.

Amazon Nova Canvas

InvokeModel

The following code example shows how to invoke Amazon Nova Canvas on Amazon Bedrock to generate an image.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

Create an image with Amazon Nova Canvas.

```
// Use the native inference API to create an image with Amazon Nova Canvas

import AWSBedrockRuntime
import AWSSDKIdentity
import Foundation

struct NovaImageOutput: Decodable {
    let images: [Data]
}

func generateImage(_ textPrompt: String) async throws {
```

```
    // Create a Bedrock Runtime client in the AWS Region you want to use.
    var config =
        try await BedrockRuntimeClient.BedrockRuntimeClientConfig()
    config.awsCredentialIdentityResolver = try SSOAWSCredentialIdentityResolver()

    let client = BedrockRuntimeClient(config: config)

    // Set the model ID.
    let modelId = "amazon.nova-canvas-v1:0"

    //  Format the request payload using the model's native structure.
    let input = InvokeModelInput(
        accept: "application/json",
        body: """
            {
                "textToImageParams": {
                    "text": "\(textPrompt)"
                },
                "taskType": "TEXT_IMAGE",
                "imageGenerationConfig": {
                    "seed": 42,
                    "quality": "standard",
                    "width": 512,
                    "height": 512,
                    "numberOfImages": 1
                }
            }
            """.data(using: .utf8),
        modelId: modelId
    )

    // Invoke the model with the request.
    let response = try await client.invokeModel(input: input)

    // Decode the response body.
    let output = try JSONDecoder().decode(NovaImageOutput.self, from:
response.body!)

    // Extract the image data.
    guard let data = output.images.first else {
        print("No image data found")
        return
    }
```

```
    // Save the generated image to a local folder.
    let fileURL = URL.documentsDirectory.appending(path: "nova_canvas.png")
    print(fileURL)
    try data.write(to: fileURL)
    print("Image is saved at \(fileURL)")
}
```

- For API details, see InvokeModel in *AWS SDK for Swift API reference*.

Amazon Nova Reel

Text-to-video

The following code example shows how to use Amazon Nova Reel to generate a video from a text prompt.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

Use Amazon Nova Reel to generate a video from a text prompt.

```
// This example demonstrates how to use Amazon Nova Reel to generate a video from a
 text prompt.
// It shows how to:
// - Set up the Amazon Bedrock runtime client
// - Configure a text-to-video request
// - Submit an asynchronous job for video generation
// - Poll for job completion status
// - Access the generated video from S3

import AWSBedrockRuntime
import Foundation
import Smithy

func startTextToVideoGenerationJob(
```

```
    bedrockRuntimeClient: BedrockRuntimeClient, prompt: String, outputS3Uri: String
) async throws -> String? {
    // Specify the model ID for text-to-video generation
    let modelId = "amazon.nova-reel-v1:0"

    // Configure the video generation request with additional parameters
    let modelInputSource: [String: Any] = [
        "taskType": "TEXT_VIDEO",
        "textToVideoParams": [
            "text": "\(prompt)"
        ],
        "videoGenerationConfig": [
            "durationSeconds": 6,
            "fps": 24,
            "dimension": "1280x720",
        ],
    ]

    let modelInput = try Document.make(from: modelInputSource)

    let input = StartAsyncInvokeInput(
        modelId: modelId,
        modelInput: modelInput,
        outputDataConfig: .s3outputdataconfig(
            BedrockRuntimeClientTypes.AsyncInvokeS3OutputDataConfig(
                s3Uri: outputS3Uri
            )
        )
    )

    // Invoke the model asynchronously
    let output = try await bedrockRuntimeClient.startAsyncInvoke(input: input)
    return output.invocationArn
}

func queryJobStatus(
    bedrockRuntimeClient: BedrockRuntimeClient,
    invocationArn: String?
) async throws -> GetAsyncInvokeOutput {
    try await bedrockRuntimeClient.getAsyncInvoke(
        input: GetAsyncInvokeInput(invocationArn: invocationArn))
}

func main() async throws {
```

```
    // Create a Bedrock Runtime client
    let config =
        try await BedrockRuntimeClient.BedrockRuntimeClientConfig()
    let client = BedrockRuntimeClient(config: config)

    // Specify the S3 location for the output video
    let bucket = "s3://REPLACE-WITH-YOUR-S3-BUCKET-NAM"

    print("Submitting video generation job...")
    let invocationArn = try await startTextToVideoGenerationJob(
        bedrockRuntimeClient: client,
        prompt: "A pomegranate juice in a railway station",
        outputS3Uri: bucket
    )
    print("Job started with invocation ARN: \(String(describing:invocationArn))")

    // Poll for job completion
    var status: BedrockRuntimeClientTypes.AsyncInvokeStatus?
    var isReady = false
    var hasFailed = false

    while !isReady && !hasFailed {
        print("\nPolling job status...")
        status = try await queryJobStatus(
            bedrockRuntimeClient: client, invocationArn: invocationArn
        ).status
        switch status {
        case .completed:
            isReady = true
            print("Video is ready\nCheck S3 bucket: \(bucket)")
        case .failed:
            hasFailed = true
            print("Something went wrong")
        case .inProgress:
            print("Job is in progress...")
            try await Task.sleep(nanoseconds: 15 * 1_000_000_000)  // 15 seconds
        default:
            isReady = true
        }
    }
}

do {
    try await main()
```

```
} catch {
    print("An error occurred: \(error)")
}
```

- For API details, see the following topics in *AWS SDK for Swift API reference*.
  - GetAsyncInvoke
  - StartAsyncInvoke

Anthropic Claude

Converse

The following code example shows how to send a text message to Anthropic Claude, using Bedrock's Converse API.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

Send a text message to Anthropic Claude, using Bedrock's Converse API.

```
// An example demonstrating how to use the Conversation API to send
// a text message to Anthropic Claude.

import AWSBedrockRuntime

func converse(_ textPrompt: String) async throws -> String {

    // Create a Bedrock Runtime client in the AWS Region you want to use.
    let config =
        try await BedrockRuntimeClient.BedrockRuntimeClientConfig()
    let client = BedrockRuntimeClient(config: config)

    // Set the model ID.
    // Claude Haiku 5.5 is available through an inference profile, so this uses
```

```
    // the "us." cross-Region profile ID that matches the default Region.
    let modelId = "us.anthropic.claude-haiku-5-5"

    // Start a conversation with the user message.
    let message = BedrockRuntimeClientTypes.Message(
        content: [.text(textPrompt)],
        role: .user
    )

    // Optionally use inference parameters
    let inferenceConfig =
        BedrockRuntimeClientTypes.InferenceConfiguration(
            maxTokens: 512,
            stopSequences: ["END"]
        )

    // Create the ConverseInput to send to the model
    let input = ConverseInput(
        inferenceConfig: inferenceConfig, messages: [message], modelId: modelId)

    // Send the ConverseInput to the model
    let response = try await client.converse(input: input)

    // Extract and return the response text.
    if case let .message(msg) = response.output {
        if case let .text(textResponse) = msg.content![0] {
            return textResponse
        } else {
            return "No text response found in message content"
        }
    } else {
        return "No message found in converse output"
    }
}
```

- For API details, see Converse in *AWS SDK for Swift API reference*.

ConverseStream

The following code example shows how to send a text message to Anthropic Claude, using Bedrock's Converse API and process the response stream in real-time.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

Send a text message to Anthropic Claude, using Bedrock's Converse API and process the response stream in real-time.

```
// An example demonstrating how to use the Conversation API to send a text message
// to Anthropic Claude and print the response stream

import AWSBedrockRuntime

func printConverseStream(_ textPrompt: String) async throws {

    // Create a Bedrock Runtime client in the AWS Region you want to use.
    let config =
        try await BedrockRuntimeClient.BedrockRuntimeClientConfig()
    let client = BedrockRuntimeClient(config: config)

    // Set the model ID.
    // Claude Haiku 5.5 is available through an inference profile, so this uses
    // the "us." cross-Region profile ID that matches the default Region.
    let modelId = "us.anthropic.claude-haiku-5-5"

    // Start a conversation with the user message.
    let message = BedrockRuntimeClientTypes.Message(
        content: [.text(textPrompt)],
        role: .user
    )

    // Optionally use inference parameters.
    let inferenceConfig =
        BedrockRuntimeClientTypes.InferenceConfiguration(
            maxTokens: 512,
            stopSequences: ["END"]
        )

    // Create the ConverseStreamInput to send to the model.
```

```
    let input = ConverseStreamInput(
        inferenceConfig: inferenceConfig, messages: [message], modelId: modelId)

    // Send the ConverseStreamInput to the model.
    let response = try await client.converseStream(input: input)

    // Extract the streaming response.
    guard let stream = response.stream else {
        print("No stream available")
        return
    }

    // Extract and print the streamed response text in real-time.
    for try await event in stream {
        switch event {
        case .messagestart(_):
            print("\nAnthropic Claude:")

        case .contentblockdelta(let deltaEvent):
            if case .text(let text) = deltaEvent.delta {
                print(text, terminator: "")
            }

        default:
            break
        }
    }
}
```

- For API details, see ConverseStream in *AWS SDK for Swift API reference*.

Meta Llama

Converse

The following code example shows how to send a text message to Meta Llama, using Bedrock's Converse API.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

Send a text message to Meta Llama, using Bedrock's Converse API.

```
// An example demonstrating how to use the Conversation API to send
// a text message to Meta Llama.

import AWSBedrockRuntime

func converse(_ textPrompt: String) async throws -> String {

    // Create a Bedrock Runtime client in the AWS Region you want to use.
    let config =
        try await BedrockRuntimeClient.BedrockRuntimeClientConfig()
    let client = BedrockRuntimeClient(config: config)

    // Set the model ID.
    let modelId = "meta.llama3-8b-instruct-v1:0"

    // Start a conversation with the user message.
    let message = BedrockRuntimeClientTypes.Message(
        content: [.text(textPrompt)],
        role: .user
    )

    // Optionally use inference parameters
    let inferenceConfig =
        BedrockRuntimeClientTypes.InferenceConfiguration(
            maxTokens: 512,
            stopSequences: ["END"],
            temperature: 0.5,
            topp: 0.9
        )

    // Create the ConverseInput to send to the model
    let input = ConverseInput(
        inferenceConfig: inferenceConfig, messages: [message], modelId: modelId)
```

```
    // Send the ConverseInput to the model
    let response = try await client.converse(input: input)

    // Extract and return the response text.
    if case let .message(msg) = response.output {
        if case let .text(textResponse) = msg.content![0] {
            return textResponse
        } else {
            return "No text response found in message content"
        }
    } else {
        return "No message found in converse output"
    }
}
```

- For API details, see Converse in *AWS SDK for Swift API reference*.

ConverseStream

The following code example shows how to send a text message to Meta Llama, using Bedrock's Converse API and process the response stream in real-time.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

Send a text message to Meta Llama, using Bedrock's Converse API and process the response stream in real-time.

```
// An example demonstrating how to use the Conversation API to send a text message
// to Meta Llama and print the response stream.

import AWSBedrockRuntime

func printConverseStream(_ textPrompt: String) async throws {
```

```
    // Create a Bedrock Runtime client in the AWS Region you want to use.
    let config =
        try await BedrockRuntimeClient.BedrockRuntimeClientConfig()
    let client = BedrockRuntimeClient(config: config)

    // Set the model ID.
    let modelId = "meta.llama3-8b-instruct-v1:0"

    // Start a conversation with the user message.
    let message = BedrockRuntimeClientTypes.Message(
        content: [.text(textPrompt)],
        role: .user
    )

    // Optionally use inference parameters.
    let inferenceConfig =
        BedrockRuntimeClientTypes.InferenceConfiguration(
            maxTokens: 512,
            stopSequences: ["END"],
            temperature: 0.5,
            topp: 0.9
        )

    // Create the ConverseStreamInput to send to the model.
    let input = ConverseStreamInput(
        inferenceConfig: inferenceConfig, messages: [message], modelId: modelId)

    // Send the ConverseStreamInput to the model.
    let response = try await client.converseStream(input: input)

    // Extract the streaming response.
    guard let stream = response.stream else {
        print("No stream available")
        return
    }

    // Extract and print the streamed response text in real-time.
    for try await event in stream {
        switch event {
        case .messagestart(_):
            print("\nMeta Llama:")

        case .contentblockdelta(let deltaEvent):
```

```
            if case .text(let text) = deltaEvent.delta {
                print(text, terminator: "")
            }

        default:
            break
        }
    }
}
```

- For API details, see ConverseStream in *AWS SDK for Swift API reference*.

Amazon Cognito Identity examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with Amazon Cognito Identity.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Actions

Actions

`CreateIdentityPool`

The following code example shows how to use `CreateIdentityPool`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSCognitoIdentity


    /// Create a new identity pool and return its ID.
    ///
    /// - Parameters:
    ///     - name: The name to give the new identity pool.
    ///
    /// - Returns: A string containing the newly created pool's ID, or `nil`
    ///   if an error occurred.
    ///
    func createIdentityPool(name: String) async throws -> String? {
        do {
            let cognitoInputCall = CreateIdentityPoolInput(developerProviderName:
 "com.exampleco.CognitoIdentityDemo",
                                                           identityPoolName: name)

            let result = try await cognitoIdentityClient.createIdentityPool(input:
 cognitoInputCall)
            guard let poolId = result.identityPoolId else {
                return nil
            }

            return poolId
        } catch {
            print("ERROR: createIdentityPool:", dump(error))
            throw error
        }
    }
```

- For more information, see AWS SDK for Swift developer guide.
- For API details, see CreateIdentityPool in *AWS SDK for Swift API reference*.

`DeleteIdentityPool`

The following code example shows how to use `DeleteIdentityPool`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSCognitoIdentity


    /// Delete the specified identity pool.
    ///
    /// - Parameters:
    ///   - id: The ID of the identity pool to delete.
    ///
    func deleteIdentityPool(id: String) async throws {
        do {
            let input = DeleteIdentityPoolInput(
                identityPoolId: id
            )

            _ = try await cognitoIdentityClient.deleteIdentityPool(input: input)
        } catch {
            print("ERROR: deleteIdentityPool:", dump(error))
            throw error
        }
    }
```

- For more information, see AWS SDK for Swift developer guide.
- For API details, see DeleteIdentityPool in *AWS SDK for Swift API reference*.

`ListIdentityPools`

The following code example shows how to use `ListIdentityPools`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSCognitoIdentity


    /// Return the ID of the identity pool with the specified name.
    ///
    /// - Parameters:
    ///   - name: The name of the identity pool whose ID should be returned.
    ///
    /// - Returns: A string containing the ID of the specified identity pool
    ///   or `nil` on error or if not found.
    ///
    func getIdentityPoolID(name: String) async throws -> String? {
        let listPoolsInput = ListIdentityPoolsInput(maxResults: 25)
        // Use "Paginated" to get all the objects.
        // This lets the SDK handle the 'nextToken' field in
 "ListIdentityPoolsOutput".
        let pages = cognitoIdentityClient.listIdentityPoolsPaginated(input:
 listPoolsInput)

        do {
            for try await page in pages {
                guard let identityPools = page.identityPools else {
                    print("ERROR: listIdentityPoolsPaginated returned nil
 contents.")
                    continue
                }

                /// Read pages of identity pools from Cognito until one is found
                /// whose name matches the one specified in the `name` parameter.
                /// Return the matching pool's ID.

                for pool in identityPools {
                    if pool.identityPoolName == name {
                        return pool.identityPoolId!
```

```
                    }
                }
            }
        } catch {
            print("ERROR: getIdentityPoolID:", dump(error))
            throw error
        }

        return nil
    }
```

Get the ID of an existing identity pool or create it if it doesn't already exist.

```
import AWSCognitoIdentity


    /// Return the ID of the identity pool with the specified name.
    ///
    /// - Parameters:
    ///   - name: The name of the identity pool whose ID should be returned
    ///
    /// - Returns: A string containing the ID of the specified identity pool.
    ///   Returns `nil` if there's an error or if the pool isn't found.
    ///
    public func getOrCreateIdentityPoolID(name: String) async throws -> String? {
        // See if the pool already exists. If it doesn't, create it.

        do {
            guard let poolId = try await getIdentityPoolID(name: name) else {
                return try await createIdentityPool(name: name)
            }

            return poolId
        } catch {
            print("ERROR: getOrCreateIdentityPoolID:", dump(error))
            throw error
        }
    }
```

- For more information, see AWS SDK for Swift developer guide.
- For API details, see ListIdentityPools in *AWS SDK for Swift API reference*.

Amazon Cognito Identity Provider examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with Amazon Cognito Identity Provider.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Scenarios are code examples that show you how to accomplish specific tasks by calling multiple functions within a service or combined with other AWS services.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Actions
- Scenarios

Actions

`AdminGetUser`

The following code example shows how to use `AdminGetUser`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
```

```
import AWSCognitoIdentityProvider

    /// Get information about a specific user in a user pool.
    ///
    /// - Parameters:
    ///   - cipClient: The Amazon Cognito Identity Provider client to use.
    ///   - userName: The user to retrieve information about.
    ///   - userPoolId: The user pool to search for the specified user.
    ///
    /// - Returns: `true` if the user's information was successfully
    ///   retrieved. Otherwise returns `false`.
    func adminGetUser(cipClient: CognitoIdentityProviderClient, userName: String,
                      userPoolId: String) async -> Bool {
        do {
            let output = try await cipClient.adminGetUser(
                input: AdminGetUserInput(
                    userPoolId: userPoolId,
                    username: userName
                )
            )

            guard let userStatus = output.userStatus else {
                print("*** Unable to get the user's status.")
                return false
            }

            print("User status: \(userStatus)")
            return true
        } catch {
            return false
        }
    }
```

- For API details, see AdminGetUser in *AWS SDK for Swift API reference*.

AdminInitiateAuth

The following code example shows how to use `AdminInitiateAuth`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSCognitoIdentityProvider

    /// Begin an authentication session.
    ///
    /// - Parameters:
    ///   - cipClient: The `CongitoIdentityProviderClient` to use.
    ///   - clientId: The app client ID to use.
    ///   - userName: The username to check.
    ///   - password: The user's password.
    ///   - userPoolId: The user pool to use.
    ///
    /// - Returns: The session token associated with this authentication
    ///   session.
    func initiateAuth(cipClient: CognitoIdentityProviderClient, clientId: String,
                      userName: String, password: String,
                      userPoolId: String) async -> String? {
        var authParams: [String: String] = [:]

        authParams["USERNAME"] = userName
        authParams["PASSWORD"] = password

        do {
            let output = try await cipClient.adminInitiateAuth(
                input: AdminInitiateAuthInput(
                    authFlow:
 CognitoIdentityProviderClientTypes.AuthFlowType.adminUserPasswordAuth,
                    authParameters: authParams,
                    clientId: clientId,
                    userPoolId: userPoolId
                )
            )

            guard let challengeName = output.challengeName else {
```

```
                print("*** Invalid response from the auth service.")
                return nil
            }

            print("=====> Response challenge is \(challengeName)")

            return output.session
        } catch _ as UserNotFoundException {
            print("*** The specified username, \(userName), doesn't exist.")
            return nil
        } catch _ as UserNotConfirmedException {
            print("*** The user \(userName) has not been confirmed.")
            return nil
        } catch {
            print("*** An unexpected error occurred.")
            return nil
        }
    }
```

- For API details, see AdminInitiateAuth in *AWS SDK for Swift API reference*.

`AdminRespondToAuthChallenge`

The following code example shows how to use `AdminRespondToAuthChallenge`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSCognitoIdentityProvider

    /// Respond to the authentication challenge received from Cognito after
    /// initiating an authentication session. This involves sending a current
    /// MFA code to the service.
    ///
    /// - Parameters:
```

```
   ///   - cipClient: The `CognitoIdentityProviderClient` to use.
   ///   - userName: The user's username.
   ///   - clientId: The app client ID.
   ///   - userPoolId: The user pool to sign into.
   ///   - mfaCode: The 6-digit MFA code currently displayed by the user's
   ///     authenticator.
   ///   - session: The authentication session to continue processing.
   func adminRespondToAuthChallenge(cipClient: CognitoIdentityProviderClient,
userName: String,
                                    clientId: String, userPoolId: String, mfaCode:
String,
                                    session: String) async {
       print("=====> SOFTWARE_TOKEN_MFA challenge is generated...")

       var challengeResponsesOb: [String: String] = [:]
       challengeResponsesOb["USERNAME"] = userName
       challengeResponsesOb["SOFTWARE_TOKEN_MFA_CODE"] = mfaCode

       do {
           let output = try await cipClient.adminRespondToAuthChallenge(
               input: AdminRespondToAuthChallengeInput(
                   challengeName:
CognitoIdentityProviderClientTypes.ChallengeNameType.softwareTokenMfa,
                   challengeResponses: challengeResponsesOb,
                   clientId: clientId,
                   session: session,
                   userPoolId: userPoolId
               )
           )

           guard let authenticationResult = output.authenticationResult else {
               print("*** Unable to get authentication result.")
               return
           }

           print("=====> Authentication result (JWTs are redacted):")
           print(authenticationResult)
       } catch _ as SoftwareTokenMFANotFoundException {
           print("*** The specified user pool isn't configured for MFA.")
           return
       } catch _ as CodeMismatchException {
           print("*** The specified MFA code doesn't match the expected value.")
           return
       } catch _ as UserNotFoundException {
```

```
            print("*** The specified username, \(userName), doesn't exist.")
            return
        } catch _ as UserNotConfirmedException {
            print("*** The user \(userName) has not been confirmed.")
            return
        } catch let error as NotAuthorizedException {
            print("*** Unauthorized access. Reason: \(error.properties.message ??
 "<unknown>")")
        } catch {
            print("*** Error responding to the MFA challenge.")
            return
        }
    }
```

- For API details, see AdminRespondToAuthChallenge in *AWS SDK for Swift API reference*.

`AssociateSoftwareToken`

The following code example shows how to use `AssociateSoftwareToken`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSCognitoIdentityProvider

    /// Request and display an MFA secret token that the user should enter
    /// into their authenticator to set it up for the user account.
    ///
    /// - Parameters:
    ///   - cipClient: The `CognitoIdentityProviderClient` to use.
    ///   - authSession: The authentication session to request an MFA secret
    ///     for.
    ///
    /// - Returns: A string containing the MFA secret token that should be
```

```
    ///   entered into the authenticator software.
    func getSecretForAppMFA(cipClient: CognitoIdentityProviderClient, authSession:
String?) async -> String? {
        do {
            let output = try await cipClient.associateSoftwareToken(
                input: AssociateSoftwareTokenInput(
                    session: authSession
                )
            )

            guard let secretCode = output.secretCode else {
                print("*** Unable to get the secret code")
                return nil
            }

            print("=====> Enter this token into Google Authenticator:
\(secretCode)")
            return output.session
        } catch _ as SoftwareTokenMFANotFoundException {
            print("*** The specified user pool isn't configured for MFA.")
            return nil
        } catch {
            print("*** An unexpected error occurred getting the secret for the app's
MFA.")
            return nil
        }
    }
```

- For API details, see AssociateSoftwareToken in *AWS SDK for Swift API reference*.

`ConfirmSignUp`

The following code example shows how to use `ConfirmSignUp`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSCognitoIdentityProvider

    /// Submit a confirmation code for the specified user. This is the code as
    /// entered by the user after they've received it by email or text
    /// message.
    ///
    /// - Parameters:
    ///   - cipClient: The `CognitoIdentityProviderClient` to use.
    ///   - clientId: The app client ID the user is signing up for.
    ///   - userName: The username of the user whose code is being sent.
    ///   - code: The user's confirmation code.
    ///
    /// - Returns: `true` if the code was successfully confirmed; otherwise `false`.
    func confirmSignUp(cipClient: CognitoIdentityProviderClient, clientId: String,
                       userName: String, code: String) async -> Bool {
        do {
            _ = try await cipClient.confirmSignUp(
                input: ConfirmSignUpInput(
                    clientId: clientId,
                    confirmationCode: code,
                    username: userName
                )
            )

            print("=====> \(userName) has been confirmed.")
            return true
        } catch {
            print("=====> \(userName)'s code was entered incorrectly.")
            return false
        }
    }
```

- For API details, see ConfirmSignUp in *AWS SDK for Swift API reference*.

`ListUsers`

The following code example shows how to use `ListUsers`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
        do {
            let output = try await cognitoClient.listUsers(
                input: ListUsersInput(
                    userPoolId: poolId
                )
            )

            guard let users = output.users else {
                print("No users found.")
                return
            }

            print("\(users.count) user(s) found.")
            for user in users {
                print("  \(user.username ?? "<unknown>")")
            }
        } catch _ as NotAuthorizedException {
            print("*** Please authenticate with AWS before using this command.")
            return
        } catch _ as ResourceNotFoundException {
            print("*** The specified User Pool was not found.")
            return
        } catch {
            print("*** An unexpected type of error occurred.")
            return
        }
```

- For API details, see ListUsers in *AWS SDK for Swift API reference*.

`ResendConfirmationCode`

The following code example shows how to use `ResendConfirmationCode`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSCognitoIdentityProvider

    /// Requests a new confirmation code be sent to the given user's contact
    /// method.
    ///
    /// - Parameters:
    ///   - cipClient: The `CognitoIdentityProviderClient` to use.
    ///   - clientId: The application client ID.
    ///   - userName: The user to resend a code for.
    ///
    /// - Returns: `true` if a new code was sent successfully, otherwise
    ///   `false`.
    func resendConfirmationCode(cipClient: CognitoIdentityProviderClient, clientId:
 String,
                                userName: String) async -> Bool {
        do {
            let output = try await cipClient.resendConfirmationCode(
                input: ResendConfirmationCodeInput(
                    clientId: clientId,
                    username: userName
                )
            )

            guard let deliveryMedium = output.codeDeliveryDetails?.deliveryMedium
 else {
                print("*** Unable to get the delivery method for the resent code.")
                return false
            }

            print("=====> A new code has been sent by \(deliveryMedium)")
            return true
        } catch {
            print("*** Unable to resend the confirmation code to user \(userName).")
```

```
            return false
        }
    }
```

- For API details, see ResendConfirmationCode in *AWS SDK for Swift API reference*.

`SignUp`

The following code example shows how to use `SignUp`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSCognitoIdentityProvider

    /// Create a new user in a user pool.
    ///
    /// - Parameters:
    ///   - cipClient: The `CognitoIdentityProviderClient` to use.
    ///   - clientId: The ID of the app client to create a user for.
    ///   - userName: The username for the new user.
    ///   - password: The new user's password.
    ///   - email: The new user's email address.
    ///
    /// - Returns: `true` if successful; otherwise `false`.
    func signUp(cipClient: CognitoIdentityProviderClient, clientId: String,
 userName: String, password: String, email: String) async -> Bool {
        let emailAttr = CognitoIdentityProviderClientTypes.AttributeType(
            name: "email",
            value: email
        )

        let userAttrsList = [emailAttr]
```

```
        do {
            _ = try await cipClient.signUp(
                input: SignUpInput(
                    clientId: clientId,
                    password: password,
                    userAttributes: userAttrsList,
                    username: userName
                )

            )

            print("=====> User \(userName) signed up.")
        } catch _ as AWSCognitoIdentityProvider.UsernameExistsException {
            print("*** The username \(userName) already exists. Please use a
different one.")
            return false
        } catch let error as AWSCognitoIdentityProvider.InvalidPasswordException {
            print("*** Error: The specified password is invalid. Reason:
\(error.properties.message ?? "<none available>").")
            return false
        } catch _ as AWSCognitoIdentityProvider.ResourceNotFoundException {
            print("*** Error: The specified client ID (\(clientId)) doesn't exist.")
            return false
        } catch {
            print("*** Unexpected error: \(error)")
            return false
        }

        return true
    }
```

- For API details, see SignUp in *AWS SDK for Swift API reference*.

`VerifySoftwareToken`

The following code example shows how to use `VerifySoftwareToken`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSCognitoIdentityProvider

    /// Confirm that the user's TOTP authenticator is configured correctly by
    /// sending a code to it to check that it matches successfully.
    ///
    /// - Parameters:
    ///   - cipClient: The `CongnitoIdentityProviderClient` to use.
    ///   - session: An authentication session previously returned by an
    ///     `associateSoftwareToken()` call.
    ///   - mfaCode: The 6-digit code currently displayed by the user's
    ///     authenticator, as provided by the user.
    func verifyTOTP(cipClient: CognitoIdentityProviderClient, session: String?,
 mfaCode: String?) async {
        do {
            let output = try await cipClient.verifySoftwareToken(
                input: VerifySoftwareTokenInput(
                    session: session,
                    userCode: mfaCode
                )
            )

            guard let tokenStatus = output.status else {
                print("*** Unable to get the token's status.")
                return
            }
            print("=====> The token's status is: \(tokenStatus)")
        } catch _ as SoftwareTokenMFANotFoundException {
            print("*** The specified user pool isn't configured for MFA.")
            return
        } catch _ as CodeMismatchException {
            print("*** The specified MFA code doesn't match the expected value.")
            return
        } catch _ as UserNotFoundException {
```

```
            print("*** The specified username doesn't exist.")
            return
        } catch _ as UserNotConfirmedException {
            print("*** The user has not been confirmed.")
            return
        } catch {
            print("*** Error verifying the MFA token!")
            return
        }
    }
```

- For API details, see VerifySoftwareToken in *AWS SDK for Swift API reference*.

Scenarios

Sign up a user with a user pool that requires MFA

The following code example shows how to:

- Sign up and confirm a user with a username, password, and email address.
- Set up multi-factor authentication by associating an MFA application with the user.
- Sign in by using a password and an MFA code.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

The `Package.swift` file.

```
// swift-tools-version: 5.9
//
// The swift-tools-version declares the minimum version of Swift required to
// build this package.

import PackageDescription
```

```
let package = Package(
    name: "cognito-scenario",
    // Let Xcode know the minimum Apple platforms supported.
    platforms: [
        .macOS(.v13),
        .iOS(.v15)
    ],
    dependencies: [
        // Dependencies declare other packages that this package depends on.
        .package(
            url: "https://github.com/awslabs/aws-sdk-swift",
            from: "1.0.0"),
        .package(
            url: "https://github.com/apple/swift-argument-parser.git",
            branch: "main"
        )
    ],
    targets: [
        // Targets are the basic building blocks of a package, defining a module or
 a test suite.
        // Targets can depend on other targets in this package and products
        // from dependencies.
        .executableTarget(
            name: "cognito-scenario",
            dependencies: [
                .product(name: "AWSCognitoIdentityProvider", package: "aws-sdk-
swift"),
                .product(name: "ArgumentParser", package: "swift-argument-parser")
            ],
            path: "Sources")

    ]
)
```

The Swift code file.

```
// An example demonstrating various features of Amazon Cognito. Before running
// this Swift code example, set up your development environment, including
// your credentials.
//
// For more information, see the following documentation:
```

```
// https://docs.aws.amazon.com/sdk-for-kotlin/latest/developer-guide/setup.html
//
// TIP: To set up the required user pool, run the AWS Cloud Development Kit
// (AWS CDK) script provided in this GitHub repo at
// resources/cdk/cognito_scenario_user_pool_with_mfa.
//
// This example performs the following functions:
//
// 1. Invokes the signUp method to sign up a user.
// 2. Invokes the adminGetUser method to get the user's confirmation status.
// 3. Invokes the ResendConfirmationCode method if the user requested another
//    code.
// 4. Invokes the confirmSignUp method.
// 5. Invokes the initiateAuth to sign in. This results in being prompted to
//    set up TOTP (time-based one-time password). (The response is
//    “ChallengeName”: “MFA_SETUP”).
// 6. Invokes the AssociateSoftwareToken method to generate a TOTP MFA private
//    key. This can be used with Google Authenticator.
// 7. Invokes the VerifySoftwareToken method to verify the TOTP and register
//    for MFA.
// 8. Invokes the AdminInitiateAuth to sign in again. This results in being
//    prompted to submit a TOTP (Response: “ChallengeName”:
//    “SOFTWARE_TOKEN_MFA”).
// 9. Invokes the AdminRespondToAuthChallenge to get back a token.

import ArgumentParser
import Foundation

import AWSClientRuntime
import AWSCognitoIdentityProvider

struct ExampleCommand: ParsableCommand {
    @Argument(help: "The application clientId.")
    var clientId: String
    @Argument(help: "The user pool ID to use.")
    var poolId: String
    @Option(help: "Name of the Amazon Region to use")
    var region = "us-east-1"

    static var configuration = CommandConfiguration(
        commandName: "cognito-scenario",
        abstract: """
        Demonstrates various features of Amazon Cognito.
        """,
```

```
        discussion: """
        """
    )

    /// Prompt for an input string of at least a minimum length.
    ///
    /// - Parameters:
    ///   - prompt: The prompt string to display.
    ///   - minLength: The minimum number of characters to allow in the
    ///     response. Default value is 0.
    ///
    /// - Returns: The entered string.
    func stringRequest(_ prompt: String, minLength: Int = 1) -> String {
        while true {
            print(prompt, terminator: "")
            let str = readLine()

            guard let str else {
                continue
            }
            if str.count >= minLength {
                return str
            } else {
                print("*** Response must be at least \(minLength) character(s)
long.")
            }
        }
    }

    /// Ask a yes/no question.
    ///
    /// - Parameter prompt: A prompt string to print.
    ///
    /// - Returns: `true` if the user answered "Y", otherwise `false`.
    func yesNoRequest(_ prompt: String) -> Bool {
        while true {
            let answer = stringRequest(prompt).lowercased()
            if answer == "y" || answer == "n" {
                return answer == "y"
            }
        }
    }

    /// Get information about a specific user in a user pool.
```

```
    ///
    /// - Parameters:
    ///   - cipClient: The Amazon Cognito Identity Provider client to use.
    ///   - userName: The user to retrieve information about.
    ///   - userPoolId: The user pool to search for the specified user.
    ///
    /// - Returns: `true` if the user's information was successfully
    ///   retrieved. Otherwise returns `false`.
    func adminGetUser(cipClient: CognitoIdentityProviderClient, userName: String,
                      userPoolId: String) async -> Bool {
        do {
            let output = try await cipClient.adminGetUser(
                input: AdminGetUserInput(
                    userPoolId: userPoolId,
                    username: userName
                )
            )

            guard let userStatus = output.userStatus else {
                print("*** Unable to get the user's status.")
                return false
            }

            print("User status: \(userStatus)")
            return true
        } catch {
            return false
        }
    }

    /// Create a new user in a user pool.
    ///
    /// - Parameters:
    ///   - cipClient: The `CognitoIdentityProviderClient` to use.
    ///   - clientId: The ID of the app client to create a user for.
    ///   - userName: The username for the new user.
    ///   - password: The new user's password.
    ///   - email: The new user's email address.
    ///
    /// - Returns: `true` if successful; otherwise `false`.
    func signUp(cipClient: CognitoIdentityProviderClient, clientId: String,
userName: String, password: String, email: String) async -> Bool {
        let emailAttr = CognitoIdentityProviderClientTypes.AttributeType(
            name: "email",
```

```
            value: email
        )

        let userAttrsList = [emailAttr]

        do {
            _ = try await cipClient.signUp(
                input: SignUpInput(
                    clientId: clientId,
                    password: password,
                    userAttributes: userAttrsList,
                    username: userName
                )

            )

            print("=====> User \(userName) signed up.")
        } catch _ as AWSCognitoIdentityProvider.UsernameExistsException {
            print("*** The username \(userName) already exists. Please use a
different one.")
            return false
        } catch let error as AWSCognitoIdentityProvider.InvalidPasswordException {
            print("*** Error: The specified password is invalid. Reason:
\(error.properties.message ?? "<none available>").")
            return false
        } catch _ as AWSCognitoIdentityProvider.ResourceNotFoundException {
            print("*** Error: The specified client ID (\(clientId)) doesn't exist.")
            return false
        } catch {
            print("*** Unexpected error: \(error)")
            return false
        }

        return true
    }

    /// Requests a new confirmation code be sent to the given user's contact
    /// method.
    ///
    /// - Parameters:
    ///   - cipClient: The `CognitoIdentityProviderClient` to use.
    ///   - clientId: The application client ID.
    ///   - userName: The user to resend a code for.
    ///
```

```
    /// - Returns: `true` if a new code was sent successfully, otherwise
    ///   `false`.
    func resendConfirmationCode(cipClient: CognitoIdentityProviderClient, clientId:
String,
                                userName: String) async -> Bool {
        do {
            let output = try await cipClient.resendConfirmationCode(
                input: ResendConfirmationCodeInput(
                    clientId: clientId,
                    username: userName
                )
            )

            guard let deliveryMedium = output.codeDeliveryDetails?.deliveryMedium
else {
                print("*** Unable to get the delivery method for the resent code.")
                return false
            }

            print("=====> A new code has been sent by \(deliveryMedium)")
            return true
        } catch {
            print("*** Unable to resend the confirmation code to user \(userName).")
            return false
        }
    }

    /// Submit a confirmation code for the specified user. This is the code as
    /// entered by the user after they've received it by email or text
    /// message.
    ///
    /// - Parameters:
    ///   - cipClient: The `CognitoIdentityProviderClient` to use.
    ///   - clientId: The app client ID the user is signing up for.
    ///   - userName: The username of the user whose code is being sent.
    ///   - code: The user's confirmation code.
    ///
    /// - Returns: `true` if the code was successfully confirmed; otherwise `false`.
    func confirmSignUp(cipClient: CognitoIdentityProviderClient, clientId: String,
                       userName: String, code: String) async -> Bool {
        do {
            _ = try await cipClient.confirmSignUp(
                input: ConfirmSignUpInput(
                    clientId: clientId,
```

```
                    confirmationCode: code,
                    username: userName
                )
            )

            print("=====> \(userName) has been confirmed.")
            return true
        } catch {
            print("=====> \(userName)'s code was entered incorrectly.")
            return false
        }
    }

    /// Begin an authentication session.
    ///
    /// - Parameters:
    ///   - cipClient: The `CongitoIdentityProviderClient` to use.
    ///   - clientId: The app client ID to use.
    ///   - userName: The username to check.
    ///   - password: The user's password.
    ///   - userPoolId: The user pool to use.
    ///
    /// - Returns: The session token associated with this authentication
    ///   session.
    func initiateAuth(cipClient: CognitoIdentityProviderClient, clientId: String,
                      userName: String, password: String,
                      userPoolId: String) async -> String? {
        var authParams: [String: String] = [:]

        authParams["USERNAME"] = userName
        authParams["PASSWORD"] = password

        do {
            let output = try await cipClient.adminInitiateAuth(
                input: AdminInitiateAuthInput(
                    authFlow:
CognitoIdentityProviderClientTypes.AuthFlowType.adminUserPasswordAuth,
                    authParameters: authParams,
                    clientId: clientId,
                    userPoolId: userPoolId
                )
            )

            guard let challengeName = output.challengeName else {
```

```
                print("*** Invalid response from the auth service.")
                return nil
            }

            print("=====> Response challenge is \(challengeName)")

            return output.session
        } catch _ as UserNotFoundException {
            print("*** The specified username, \(userName), doesn't exist.")
            return nil
        } catch _ as UserNotConfirmedException {
            print("*** The user \(userName) has not been confirmed.")
            return nil
        } catch {
            print("*** An unexpected error occurred.")
            return nil
        }
    }

    /// Request and display an MFA secret token that the user should enter
    /// into their authenticator to set it up for the user account.
    ///
    /// - Parameters:
    ///   - cipClient: The `CognitoIdentityProviderClient` to use.
    ///   - authSession: The authentication session to request an MFA secret
    ///     for.
    ///
    /// - Returns: A string containing the MFA secret token that should be
    ///   entered into the authenticator software.
    func getSecretForAppMFA(cipClient: CognitoIdentityProviderClient, authSession:
 String?) async -> String? {
        do {
            let output = try await cipClient.associateSoftwareToken(
                input: AssociateSoftwareTokenInput(
                    session: authSession
                )
            )

            guard let secretCode = output.secretCode else {
                print("*** Unable to get the secret code")
                return nil
            }
```

```
            print("=====> Enter this token into Google Authenticator:
\(secretCode)")
            return output.session
        } catch _ as SoftwareTokenMFANotFoundException {
            print("*** The specified user pool isn't configured for MFA.")
            return nil
        } catch {
            print("*** An unexpected error occurred getting the secret for the app's
MFA.")
            return nil
        }
    }

    /// Confirm that the user's TOTP authenticator is configured correctly by
    /// sending a code to it to check that it matches successfully.
    ///
    /// - Parameters:
    ///   - cipClient: The `CongnitoIdentityProviderClient` to use.
    ///   - session: An authentication session previously returned by an
    ///     `associateSoftwareToken()` call.
    ///   - mfaCode: The 6-digit code currently displayed by the user's
    ///     authenticator, as provided by the user.
    func verifyTOTP(cipClient: CognitoIdentityProviderClient, session: String?,
mfaCode: String?) async {
        do {
            let output = try await cipClient.verifySoftwareToken(
                input: VerifySoftwareTokenInput(
                    session: session,
                    userCode: mfaCode
                )
            )

            guard let tokenStatus = output.status else {
                print("*** Unable to get the token's status.")
                return
            }
            print("=====> The token's status is: \(tokenStatus)")
        } catch _ as SoftwareTokenMFANotFoundException {
            print("*** The specified user pool isn't configured for MFA.")
            return
        } catch _ as CodeMismatchException {
            print("*** The specified MFA code doesn't match the expected value.")
            return
        } catch _ as UserNotFoundException {
```

```
            print("*** The specified username doesn't exist.")
            return
        } catch _ as UserNotConfirmedException {
            print("*** The user has not been confirmed.")
            return
        } catch {
            print("*** Error verifying the MFA token!")
            return
        }
    }

    /// Respond to the authentication challenge received from Cognito after
    /// initiating an authentication session. This involves sending a current
    /// MFA code to the service.
    ///
    /// - Parameters:
    ///   - cipClient: The `CognitoIdentityProviderClient` to use.
    ///   - userName: The user's username.
    ///   - clientId: The app client ID.
    ///   - userPoolId: The user pool to sign into.
    ///   - mfaCode: The 6-digit MFA code currently displayed by the user's
    ///     authenticator.
    ///   - session: The authentication session to continue processing.
    func adminRespondToAuthChallenge(cipClient: CognitoIdentityProviderClient,
userName: String,
                                     clientId: String, userPoolId: String, mfaCode:
String,
                                     session: String) async {
        print("=====> SOFTWARE_TOKEN_MFA challenge is generated...")

        var challengeResponsesOb: [String: String] = [:]
        challengeResponsesOb["USERNAME"] = userName
        challengeResponsesOb["SOFTWARE_TOKEN_MFA_CODE"] = mfaCode

        do {
            let output = try await cipClient.adminRespondToAuthChallenge(
                input: AdminRespondToAuthChallengeInput(
                    challengeName:
CognitoIdentityProviderClientTypes.ChallengeNameType.softwareTokenMfa,
                    challengeResponses: challengeResponsesOb,
                    clientId: clientId,
                    session: session,
                    userPoolId: userPoolId
                )
```

```
            )

            guard let authenticationResult = output.authenticationResult else {
                print("*** Unable to get authentication result.")
                return
            }

            print("=====> Authentication result (JWTs are redacted):")
            print(authenticationResult)
        } catch _ as SoftwareTokenMFANotFoundException {
            print("*** The specified user pool isn't configured for MFA.")
            return
        } catch _ as CodeMismatchException {
            print("*** The specified MFA code doesn't match the expected value.")
            return
        } catch _ as UserNotFoundException {
            print("*** The specified username, \(userName), doesn't exist.")
            return
        } catch _ as UserNotConfirmedException {
            print("*** The user \(userName) has not been confirmed.")
            return
        } catch let error as NotAuthorizedException {
            print("*** Unauthorized access. Reason: \(error.properties.message ??
"<unknown>")")
        } catch {
            print("*** Error responding to the MFA challenge.")
            return
        }
    }

    /// Called by ``main()`` to run the bulk of the example.
    func runAsync() async throws {
        let config = try await
CognitoIdentityProviderClient.CognitoIdentityProviderClientConfig(region: region)
        let cipClient = CognitoIdentityProviderClient(config: config)

        print("""
              This example collects information about a user, then creates that user
in the
              specified user pool. Then, it enables Multi-Factor Authentication
(MFA) for that
              user by associating an authenticator application (such as Google
Authenticator
```

```
             or a password manager that supports TOTP). Then, the user uses a code
from their
             authenticator application to sign in.

             """)

       let userName = stringRequest("Please enter a new username: ")
       let password = stringRequest("Enter a password: ")
       let email = stringRequest("Enter your email address: ", minLength: 5)

       // Submit the sign-up request to AWS.

       print("==> Signing up user \(userName)...")
       if await signUp(cipClient: cipClient, clientId: clientId,
                       userName: userName, password: password,
                       email: email) == false {
           return
       }

       // Check the user's status. This time, it should come back "unconfirmed".

       print("==> Getting the status of user \(userName) from the user pool (should
be 'unconfirmed')...")
       if await adminGetUser(cipClient: cipClient, userName: userName, userPoolId:
poolId) == false {
           return
       }

       // Ask the user if they want a replacement code sent, such as if the
       // code hasn't arrived yet. If the user responds with a "yes," send a
       // new code.

       if yesNoRequest("==> A confirmation code was sent to \(userName). Would you
like to send a new code (Y/N)? ") {
           print("==> Sending a new confirmation code...")
           if await resendConfirmationCode(cipClient: cipClient, clientId:
clientId, userName: userName) == false {
               return
           }
       }

       // Ask the user to enter the confirmation code, then send it to Amazon
       // Cognito to verify it.
```

```
        let code = stringRequest("==> Enter the confirmation code sent to
\(userName): ")
        if await confirmSignUp(cipClient: cipClient, clientId: clientId, userName:
userName, code: code) == false {
            // The code didn't match. Your application may wish to offer to
            // re-send the confirmation code here and try again.
            return
        }

        // Check the user's status again. This time it should come back
        // "confirmed".

        print("==> Rechecking status of user \(userName) in the user pool (should be
'confirmed')...")
        if await adminGetUser(cipClient: cipClient, userName: userName, userPoolId:
poolId) == false {
            return
        }
        // Check the challenge mode. Here, it should be "mfaSetup", indicating
        // that the user needs to add MFA before using it. This returns a
        // session that can be used to register MFA, or nil if an error occurs.

        let authSession = await initiateAuth(cipClient: cipClient, clientId:
clientId,
                                             userName: userName, password:
password,
                                             userPoolId: poolId)
        if authSession == nil {
            return
        }

        // Ask Cognito for an MFA secret token that the user should enter into
        // their authenticator software (such as Google Authenticator) or
        // password manager to configure it for this user account. This
        // returns a new session that should be used for the new stage of the
        // authentication process.

        let newSession = await getSecretForAppMFA(cipClient: cipClient, authSession:
authSession)
        if newSession == nil {
            return
        }

        // Ask the user to enter the current 6-digit code displayed by their
```

```
        // authenticator. Then verify that it matches the value expected for
        // the session.

        let mfaCode1 = stringRequest("==> Enter the 6-digit code displayed in your
 authenticator: ",
                                     minLength: 6)
        await verifyTOTP(cipClient: cipClient, session: newSession, mfaCode:
 mfaCode1)

        // Ask the user to authenticate now that the authenticator has been
        // configured. This creates a new session using the user's username
        // and password as already entered.

        print("\nNow starting the sign-in process for user \(userName)...\n")

        let session2 = await initiateAuth(cipClient: cipClient, clientId: clientId,
                                     userName: userName, password: password,
 userPoolId: poolId)
        guard let session2 else {
            return
        }

        // Now that we have a new auth session, `session2`, ask the user for a
        // new 6-digit code from their authenticator, and send it to the auth
        // session.

        let mfaCode2 = stringRequest("==> Wait for your authenticator to show a new
 6-digit code, then enter it: ",
                                     minLength: 6)
        await adminRespondToAuthChallenge(cipClient: cipClient, userName: userName,
                                          clientId: clientId, userPoolId: poolId,
                                          mfaCode: mfaCode2, session: session2)
    }
}

/// The program's asynchronous entry point.
@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.runAsync()
```

```
        } catch {
            ExampleCommand.exit(withError: error)
        }
    }
}
```

- For API details, see the following topics in *AWS SDK for Swift API reference*.
  - AdminGetUser
  - AdminInitiateAuth
  - AdminRespondToAuthChallenge
  - AssociateSoftwareToken
  - ConfirmDevice
  - ConfirmSignUp
  - InitiateAuth
  - ListUsers
  - ResendConfirmationCode
  - RespondToAuthChallenge
  - SignUp
  - VerifySoftwareToken

DynamoDB examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with DynamoDB.

Basics are code examples that show you how to perform the essential operations within a service.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Basics

- Actions

Basics

Learn the basics

The following code example shows how to:

- Create a table that can hold movie data.
- Put, get, and update a single movie in the table.
- Write movie data to the table from a sample JSON file.
- Query for movies that were released in a given year.
- Scan for movies that were released in a range of years.
- Delete a movie from the table, then delete the table.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

A Swift class that handles DynamoDB calls to the SDK for Swift.

```
import AWSDynamoDB
import Foundation

/// An enumeration of error codes representing issues that can arise when using
/// the `MovieTable` class.
enum MoviesError: Error {
    /// The specified table wasn't found or couldn't be created.
    case TableNotFound
    /// The specified item wasn't found or couldn't be created.
    case ItemNotFound
    /// The Amazon DynamoDB client is not properly initialized.
    case UninitializedClient
    /// The table status reported by Amazon DynamoDB is not recognized.
    case StatusUnknown
```

```
    /// One or more specified attribute values are invalid or missing.
    case InvalidAttributes
}

/// A class representing an Amazon DynamoDB table containing movie
/// information.
public class MovieTable {
    var ddbClient: DynamoDBClient?
    let tableName: String

    /// Create an object representing a movie table in an Amazon DynamoDB
    /// database.
    ///
    /// - Parameters:
    ///   - region: The optional Amazon Region to create the database in.
    ///   - tableName: The name to assign to the table. If not specified, a
    ///     random table name is generated automatically.
    ///
    /// > Note: The table is not necessarily available when this function
    /// returns. Use `tableExists()` to check for its availability, or
    /// `awaitTableActive()` to wait until the table's status is reported as
    /// ready to use by Amazon DynamoDB.
    ///
    init(region: String? = nil, tableName: String) async throws {
        do {
            var config = try await DynamoDBClient.DynamoDBClientConfig()
            if let region = region {
                config.region = region
            }

            self.ddbClient = DynamoDBClient(config: config)
            self.tableName = tableName

            try await self.createTable()
        } catch {
            print("ERROR: ", dump(error, name: "Initializing Amazon DynamoDBClient
 client"))
            throw error
        }
    }

    ///
    /// Create a movie table in the Amazon DynamoDB data store.
    ///
```

```
    private func createTable() async throws {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = CreateTableInput(
                attributeDefinitions: [
                    DynamoDBClientTypes.AttributeDefinition(attributeName: "year",
attributeType: .n),
                    DynamoDBClientTypes.AttributeDefinition(attributeName: "title",
attributeType: .s)
                ],
                billingMode: DynamoDBClientTypes.BillingMode.payPerRequest,
                keySchema: [
                    DynamoDBClientTypes.KeySchemaElement(attributeName: "year",
keyType: .hash),
                    DynamoDBClientTypes.KeySchemaElement(attributeName: "title",
keyType: .range)
                ],
                tableName: self.tableName
            )
            let output = try await client.createTable(input: input)
            if output.tableDescription == nil {
                throw MoviesError.TableNotFound
            }
        } catch {
            print("ERROR: createTable:", dump(error))
            throw error
        }
    }


    /// Check to see if the table exists online yet.
    ///
    /// - Returns: `true` if the table exists, or `false` if not.
    ///
    func tableExists() async throws -> Bool {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = DescribeTableInput(
```

```
                tableName: tableName
            )
            let output = try await client.describeTable(input: input)
            guard let description = output.table else {
                throw MoviesError.TableNotFound
            }

            return description.tableName == self.tableName
        } catch {
            print("ERROR: tableExists:", dump(error))
            throw error
        }
    }


    ///
    /// Waits for the table to exist and for its status to be active.
    ///
    func awaitTableActive() async throws {
        while try (await self.tableExists() == false) {
            do {
                let duration = UInt64(0.25 * 1_000_000_000) // Convert .25 seconds
to nanoseconds.
                try await Task.sleep(nanoseconds: duration)
            } catch {
                print("Sleep error:", dump(error))
            }
        }

        while try (await self.getTableStatus() != .active) {
            do {
                let duration = UInt64(0.25 * 1_000_000_000) // Convert .25 seconds
to nanoseconds.
                try await Task.sleep(nanoseconds: duration)
            } catch {
                print("Sleep error:", dump(error))
            }
        }
    }


    ///
    /// Deletes the table from Amazon DynamoDB.
    ///
```

```
    func deleteTable() async throws {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = DeleteTableInput(
                tableName: self.tableName
            )
            _ = try await client.deleteTable(input: input)
        } catch {
            print("ERROR: deleteTable:", dump(error))
            throw error
        }
    }


    /// Get the table's status.
    ///
    /// - Returns: The table status, as defined by the
    ///   `DynamoDBClientTypes.TableStatus` enum.
    ///
    func getTableStatus() async throws -> DynamoDBClientTypes.TableStatus {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = DescribeTableInput(
                tableName: self.tableName
            )
            let output = try await client.describeTable(input: input)
            guard let description = output.table else {
                throw MoviesError.TableNotFound
            }
            guard let status = description.tableStatus else {
                throw MoviesError.StatusUnknown
            }
            return status
        } catch {
            print("ERROR: getTableStatus:", dump(error))
            throw error
        }
    }
```

```
    /// Populate the movie database from the specified JSON file.
    ///
    /// - Parameter jsonPath: Path to a JSON file containing movie data.
    ///
    func populate(jsonPath: String) async throws {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            // Create a Swift `URL` and use it to load the file into a `Data`
            // object. Then decode the JSON into an array of `Movie` objects.

            let fileUrl = URL(fileURLWithPath: jsonPath)
            let jsonData = try Data(contentsOf: fileUrl)

            var movieList = try JSONDecoder().decode([Movie].self, from: jsonData)

            // Truncate the list to the first 200 entries or so for this example.

            if movieList.count > 200 {
                movieList = Array(movieList[...199])
            }

            // Before sending records to the database, break the movie list into
            // 25-entry chunks, which is the maximum size of a batch item request.

            let count = movieList.count
            let chunks = stride(from: 0, to: count, by: 25).map {
                Array(movieList[$0 ..< Swift.min($0 + 25, count)])
            }

            // For each chunk, create a list of write request records and populate
            // them with `PutRequest` requests, each specifying one movie from the
            // chunk. Once the chunk's items are all in the `PutRequest` list,
            // send them to Amazon DynamoDB using the
            // `DynamoDBClient.batchWriteItem()` function.

            for chunk in chunks {
                var requestList: [DynamoDBClientTypes.WriteRequest] = []

                for movie in chunk {
```

```
                    let item = try await movie.getAsItem()
                    let request = DynamoDBClientTypes.WriteRequest(
                        putRequest: .init(
                            item: item
                        )
                    )
                    requestList.append(request)
                }

                let input = BatchWriteItemInput(requestItems: [tableName:
requestList])
                _ = try await client.batchWriteItem(input: input)
            }
        } catch {
            print("ERROR: populate:", dump(error))
            throw error
        }
    }


    /// Add a movie specified as a `Movie` structure to the Amazon DynamoDB
    /// table.
    ///
    /// - Parameter movie: The `Movie` to add to the table.
    ///
    func add(movie: Movie) async throws {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            // Get a DynamoDB item containing the movie data.
            let item = try await movie.getAsItem()

            // Send the `PutItem` request to Amazon DynamoDB.

            let input = PutItemInput(
                item: item,
                tableName: self.tableName
            )
            _ = try await client.putItem(input: input)
        } catch {
            print("ERROR: add movie:", dump(error))
            throw error
```

```
        }
    }


    /// Given a movie's details, add a movie to the Amazon DynamoDB table.
    ///
    /// - Parameters:
    ///   - title: The movie's title as a `String`.
    ///   - year: The release year of the movie (`Int`).
    ///   - rating: The movie's rating if available (`Double`; default is
    ///     `nil`).
    ///   - plot: A summary of the movie's plot (`String`; default is `nil`,
    ///     indicating no plot summary is available).
    ///
    func add(title: String, year: Int, rating: Double? = nil,
             plot: String? = nil) async throws
    {
        do {
            let movie = Movie(title: title, year: year, rating: rating, plot: plot)
            try await self.add(movie: movie)
        } catch {
            print("ERROR: add with fields:", dump(error))
            throw error
        }
    }


    /// Return a `Movie` record describing the specified movie from the Amazon
    /// DynamoDB table.
    ///
    /// - Parameters:
    ///   - title: The movie's title (`String`).
    ///   - year: The movie's release year (`Int`).
    ///
    /// - Throws: `MoviesError.ItemNotFound` if the movie isn't in the table.
    ///
    /// - Returns: A `Movie` record with the movie's details.
    func get(title: String, year: Int) async throws -> Movie {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = GetItemInput(
```

```
                key: [
                    "year": .n(String(year)),
                    "title": .s(title)
                ],
                tableName: self.tableName
            )
            let output = try await client.getItem(input: input)
            guard let item = output.item else {
                throw MoviesError.ItemNotFound
            }

            let movie = try Movie(withItem: item)
            return movie
        } catch {
            print("ERROR: get:", dump(error))
            throw error
        }
    }


    /// Get all the movies released in the specified year.
    ///
    /// - Parameter year: The release year of the movies to return.
    ///
    /// - Returns: An array of `Movie` objects describing each matching movie.
    ///
    func getMovies(fromYear year: Int) async throws -> [Movie] {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = QueryInput(
                expressionAttributeNames: [
                    "#y": "year"
                ],
                expressionAttributeValues: [
                    ":y": .n(String(year))
                ],
                keyConditionExpression: "#y = :y",
                tableName: self.tableName
            )
            // Use "Paginated" to get all the movies.
```

```
            // This lets the SDK handle the 'lastEvaluatedKey' property in
"QueryOutput".

            let pages = client.queryPaginated(input: input)

            var movieList: [Movie] = []
            for try await page in pages {
                guard let items = page.items else {
                    print("Error: no items returned.")
                    continue
                }

                // Convert the found movies into `Movie` objects and return an array
                // of them.

                for item in items {
                    let movie = try Movie(withItem: item)
                    movieList.append(movie)
                }
            }
            return movieList
        } catch {
            print("ERROR: getMovies:", dump(error))
            throw error
        }
    }


    /// Return an array of `Movie` objects released in the specified range of
    /// years.
    ///
    /// - Parameters:
    ///   - firstYear: The first year of movies to return.
    ///   - lastYear: The last year of movies to return.
    ///   - startKey: A starting point to resume processing; always use `nil`.
    ///
    /// - Returns: An array of `Movie` objects describing the matching movies.
    ///
    /// > Note: The `startKey` parameter is used by this function when
    ///   recursively calling itself, and should always be `nil` when calling
    ///   directly.
    ///
    func getMovies(firstYear: Int, lastYear: Int,
```

```
                    startKey: [Swift.String: DynamoDBClientTypes.AttributeValue]? =
nil)
        async throws -> [Movie]
    {
        do {
            var movieList: [Movie] = []

            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = ScanInput(
                consistentRead: true,
                exclusiveStartKey: startKey,
                expressionAttributeNames: [
                    "#y": "year" // `year` is a reserved word, so use `#y` instead.
                ],
                expressionAttributeValues: [
                    ":y1": .n(String(firstYear)),
                    ":y2": .n(String(lastYear))
                ],
                filterExpression: "#y BETWEEN :y1 AND :y2",
                tableName: self.tableName
            )

            let pages = client.scanPaginated(input: input)

            for try await page in pages {
                guard let items = page.items else {
                    print("Error: no items returned.")
                    continue
                }

                // Build an array of `Movie` objects for the returned items.

                for item in items {
                    let movie = try Movie(withItem: item)
                    movieList.append(movie)
                }
            }
            return movieList

        } catch {
            print("ERROR: getMovies with scan:", dump(error))
```

```
            throw error
        }
    }


    /// Update the specified movie with new `rating` and `plot` information.
    ///
    /// - Parameters:
    ///   - title: The title of the movie to update.
    ///   - year: The release year of the movie to update.
    ///   - rating: The new rating for the movie.
    ///   - plot: The new plot summary string for the movie.
    ///
    /// - Returns: An array of mappings of attribute names to their new
    ///   listing each item actually changed. Items that didn't need to change
    ///   aren't included in this list. `nil` if no changes were made.
    ///
    func update(title: String, year: Int, rating: Double? = nil, plot: String? =
nil) async throws
        -> [Swift.String: DynamoDBClientTypes.AttributeValue]?
    {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            // Build the update expression and the list of expression attribute
            // values. Include only the information that's changed.

            var expressionParts: [String] = []
            var attrValues: [Swift.String: DynamoDBClientTypes.AttributeValue] = [:]

            if rating != nil {
                expressionParts.append("info.rating=:r")
                attrValues[":r"] = .n(String(rating!))
            }
            if plot != nil {
                expressionParts.append("info.plot=:p")
                attrValues[":p"] = .s(plot!)
            }
            let expression = "set \(expressionParts.joined(separator: ", "))"

            let input = UpdateItemInput(
                // Create substitution tokens for the attribute values, to ensure
```

```
                // no conflicts in expression syntax.
                expressionAttributeValues: attrValues,
                // The key identifying the movie to update consists of the release
                // year and title.
                key: [
                    "year": .n(String(year)),
                    "title": .s(title)
                ],
                returnValues: .updatedNew,
                tableName: self.tableName,
                updateExpression: expression
            )
            let output = try await client.updateItem(input: input)

            guard let attributes: [Swift.String: DynamoDBClientTypes.AttributeValue]
= output.attributes else {
                throw MoviesError.InvalidAttributes
            }
            return attributes
        } catch {
            print("ERROR: update:", dump(error))
            throw error
        }
    }


    /// Delete a movie, given its title and release year.
    ///
    /// - Parameters:
    ///   - title: The movie's title.
    ///   - year: The movie's release year.
    ///
    func delete(title: String, year: Int) async throws {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = DeleteItemInput(
                key: [
                    "year": .n(String(year)),
                    "title": .s(title)
                ],
                tableName: self.tableName
```

```
            )
            _ = try await client.deleteItem(input: input)
        } catch {
            print("ERROR: delete:", dump(error))
            throw error
        }
    }
}
```

The structures used by the MovieTable class to represent movies.

```
import Foundation
import AWSDynamoDB

/// The optional details about a movie.
public struct Details: Codable {
    /// The movie's rating, if available.
    var rating: Double?
    /// The movie's plot, if available.
    var plot: String?
}

/// A structure describing a movie. The `year` and `title` properties are
/// required and are used as the key for Amazon DynamoDB operations. The
/// `info` sub-structure's two properties, `rating` and `plot`, are optional.
public struct Movie: Codable {
    /// The year in which the movie was released.
    var year: Int
    /// The movie's title.
    var title: String
    /// A `Details` object providing the optional movie rating and plot
    /// information.
    var info: Details

    /// Create a `Movie` object representing a movie, given the movie's
    /// details.
    ///
    /// - Parameters:
    ///   - title: The movie's title (`String`).
    ///   - year: The year in which the movie was released (`Int`).
    ///   - rating: The movie's rating (optional `Double`).
```

```
    ///   - plot: The movie's plot (optional `String`)
    init(title: String, year: Int, rating: Double? = nil, plot: String? = nil) {
        self.title = title
        self.year = year

        self.info = Details(rating: rating, plot: plot)
    }

    /// Create a `Movie` object representing a movie, given the movie's
    /// details.
    ///
    /// - Parameters:
    ///   - title: The movie's title (`String`).
    ///   - year: The year in which the movie was released (`Int`).
    ///   - info: The optional rating and plot information for the movie in a
    ///     `Details` object.
    init(title: String, year: Int, info: Details?){
        self.title = title
        self.year = year

        if info != nil {
            self.info = info!
        } else {
            self.info = Details(rating: nil, plot: nil)
        }
    }

    ///
    /// Return a new `MovieTable` object, given an array mapping string to Amazon
    /// DynamoDB attribute values.
    ///
    /// - Parameter item: The item information provided to the form used by
    ///   DynamoDB. This is an array of strings mapped to
    ///   `DynamoDBClientTypes.AttributeValue` values.
    init(withItem item: [Swift.String:DynamoDBClientTypes.AttributeValue]) throws  {
        // Read the attributes.

        guard let titleAttr = item["title"],
              let yearAttr = item["year"] else {
            throw MoviesError.ItemNotFound
        }
        let infoAttr = item["info"] ?? nil

        // Extract the values of the title and year attributes.
```

```
        if case .s(let titleVal) = titleAttr {
            self.title = titleVal
        } else {
            throw MoviesError.InvalidAttributes
        }

        if case .n(let yearVal) = yearAttr {
            self.year = Int(yearVal)!
        } else {
            throw MoviesError.InvalidAttributes
        }

        // Extract the rating and/or plot from the `info` attribute, if
        // they're present.

        var rating: Double? = nil
        var plot: String? = nil

        if infoAttr != nil, case .m(let infoVal) = infoAttr {
            let ratingAttr = infoVal["rating"] ?? nil
            let plotAttr = infoVal["plot"] ?? nil

            if ratingAttr != nil, case .n(let ratingVal) = ratingAttr {
                rating = Double(ratingVal) ?? nil
            }
            if plotAttr != nil, case .s(let plotVal) = plotAttr {
                plot = plotVal
            }
        }

        self.info = Details(rating: rating, plot: plot)
    }

    ///
    /// Return an array mapping attribute names to Amazon DynamoDB attribute
    /// values, representing the contents of the `Movie` record as a DynamoDB
    /// item.
    ///
    /// - Returns: The movie item as an array of type
    ///   `[Swift.String:DynamoDBClientTypes.AttributeValue]`.
    ///
    func getAsItem() async throws ->
[Swift.String:DynamoDBClientTypes.AttributeValue]  {
```

```
        // Build the item record, starting with the year and title, which are
        // always present.

        var item: [Swift.String:DynamoDBClientTypes.AttributeValue] = [
            "year": .n(String(self.year)),
            "title": .s(self.title)
        ]

        // Add the `info` field with the rating and/or plot if they're
        // available.

        var details: [Swift.String:DynamoDBClientTypes.AttributeValue] = [:]
        if (self.info.rating != nil || self.info.plot != nil) {
            if self.info.rating != nil {
                details["rating"] = .n(String(self.info.rating!))
            }
            if self.info.plot != nil {
                details["plot"] = .s(self.info.plot!)
            }
        }
        item["info"] = .m(details)

        return item
    }
}
```

A program that uses the MovieTable class to access a DynamoDB database.

```
import ArgumentParser
import ClientRuntime
import Foundation

import AWSDynamoDB


@testable import MovieList

extension String {
    // Get the directory if the string is a file path.
    func directory() -> String {
        guard let lastIndex = lastIndex(of: "/") else {
            print("Error: String directory separator not found.")
```

```
            return ""
        }
        return String(self[...lastIndex])
    }
}

struct ExampleCommand: ParsableCommand {
    @Argument(help: "The path of the sample movie data JSON file.")
    var jsonPath: String = #file.directory() + "../../../../../resources/
sample_files/movies.json"

    @Option(help: "The AWS Region to run AWS API calls in.")
    var awsRegion: String?

    @Option(
        help: ArgumentHelp("The level of logging for the Swift SDK to perform."),
        completion: .list([
            "critical",
            "debug",
            "error",
            "info",
            "notice",
            "trace",
            "warning"
        ])
    )
    var logLevel: String = "error"

    /// Configuration details for the command.
    static var configuration = CommandConfiguration(
        commandName: "basics",
        abstract: "A basic scenario demonstrating the usage of Amazon DynamoDB.",
        discussion: """
        An example showing how to use Amazon DynamoDB to perform a series of
        common database activities on a simple movie database.
        """
    )

    /// Called by ``main()`` to asynchronously run the AWS example.
    func runAsync() async throws {
        print("Welcome to the AWS SDK for Swift basic scenario for Amazon
 DynamoDB!")

        //======================================================================
```

```
        // 1. Create the table. The Amazon DynamoDB table is represented by
        //    the `MovieTable` class.
        //===========================================================================

        let tableName = "ddb-movies-sample-\(Int.random(in: 1 ... Int.max))"

        print("Creating table \"\(tableName)\"...")

        let movieDatabase = try await MovieTable(region: awsRegion,
                                                 tableName: tableName)

        print("\nWaiting for table to be ready to use...")
        try await movieDatabase.awaitTableActive()

        //===========================================================================
        // 2. Add a movie to the table.
        //===========================================================================

        print("\nAdding a movie...")
        try await movieDatabase.add(title: "Avatar: The Way of Water", year: 2022)
        try await movieDatabase.add(title: "Not a Real Movie", year: 2023)

        //===========================================================================
        // 3. Update the plot and rating of the movie using an update
        //    expression.
        //===========================================================================

        print("\nAdding details to the added movie...")
        _ = try await movieDatabase.update(title: "Avatar: The Way of Water", year:
2022,
                                           rating: 9.2, plot: "It's a sequel.")

        //===========================================================================
        // 4. Populate the table from the JSON file.
        //===========================================================================

        print("\nPopulating the movie database from JSON...")
        try await movieDatabase.populate(jsonPath: jsonPath)

        //===========================================================================
        // 5. Get a specific movie by key. In this example, the key is a
        //    combination of `title` and `year`.
        //===========================================================================
```

```
        print("\nLooking for a movie in the table...")
        let gotMovie = try await movieDatabase.get(title: "This Is the End", year:
2013)

        print("Found the movie \"\(gotMovie.title)\", released in
\(gotMovie.year).")
        print("Rating: \(gotMovie.info.rating ?? 0.0).")
        print("Plot summary: \(gotMovie.info.plot ?? "None.")")

        //=====================================================================
        // 6. Delete a movie.
        //=====================================================================

        print("\nDeleting the added movie...")
        try await movieDatabase.delete(title: "Avatar: The Way of Water", year:
2022)

        //=====================================================================
        // 7. Use a query with a key condition expression to return all movies
        //    released in a given year.
        //=====================================================================

        print("\nGetting movies released in 1994...")
        let movieList = try await movieDatabase.getMovies(fromYear: 1994)
        for movie in movieList {
            print("    \(movie.title)")
        }

        //=====================================================================
        // 8. Use `scan()` to return movies released in a range of years.
        //=====================================================================

        print("\nGetting movies released between 1993 and 1997...")
        let scannedMovies = try await movieDatabase.getMovies(firstYear: 1993,
lastYear: 1997)
        for movie in scannedMovies {
            print("    \(movie.title) (\(movie.year))")
        }

        //=====================================================================
        // 9. Delete the table.
        //=====================================================================

        print("\nDeleting the table...")
```

```
        try await movieDatabase.deleteTable()
    }
}

@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.runAsync()
        } catch {
            ExampleCommand.exit(withError: error)
        }
    }
}
```

- For API details, see the following topics in *AWS SDK for Swift API reference*.
  - BatchWriteItem
  - CreateTable
  - DeleteItem
  - DeleteTable
  - DescribeTable
  - GetItem
  - PutItem
  - Query
  - Scan
  - UpdateItem

Actions

`BatchGetItem`

The following code example shows how to use `BatchGetItem`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSDynamoDB


    /// Gets an array of `Movie` objects describing all the movies in the
    /// specified list. Any movies that aren't found in the list have no
    /// corresponding entry in the resulting array.
    ///
    /// - Parameters
    ///     - keys: An array of tuples, each of which specifies the title and
    ///       release year of a movie to fetch from the table.
    ///
    /// - Returns:
    ///     - An array of `Movie` objects describing each match found in the
    ///     table.
    ///
    /// - Throws:
    ///     - `MovieError.ClientUninitialized` if the DynamoDB client has not
    ///     been initialized.
    ///     - DynamoDB errors are thrown without change.
    func batchGet(keys: [(title: String, year: Int)]) async throws -> [Movie] {
        do {
            guard let client = self.ddbClient else {
                throw MovieError.ClientUninitialized
            }

            var movieList: [Movie] = []
            var keyItems: [[Swift.String: DynamoDBClientTypes.AttributeValue]] = []

            // Convert the list of keys into the form used by DynamoDB.

            for key in keys {
                let item: [Swift.String: DynamoDBClientTypes.AttributeValue] = [
                    "title": .s(key.title),
                    "year": .n(String(key.year))
```

```
            ]
            keyItems.append(item)
        }

        // Create the input record for `batchGetItem()`. The list of requested
        // items is in the `requestItems` property. This array contains one
        // entry for each table from which items are to be fetched. In this
        // example, there's only one table containing the movie data.
        //
        // If we wanted this program to also support searching for matches
        // in a table of book data, we could add a second `requestItem`
        // mapping the name of the book table to the list of items we want to
        // find in it.
        let input = BatchGetItemInput(
            requestItems: [
                self.tableName: .init(
                    consistentRead: true,
                    keys: keyItems
                )
            ]
        )

        // Fetch the matching movies from the table.

        let output = try await client.batchGetItem(input: input)

        // Get the set of responses. If there aren't any, return the empty
        // movie list.

        guard let responses = output.responses else {
            return movieList
        }

        // Get the list of matching items for the table with the name
        // `tableName`.

        guard let responseList = responses[self.tableName] else {
            return movieList
        }

        // Create `Movie` items for each of the matching movies in the table
        // and add them to the `MovieList` array.

        for response in responseList {
```

```
                try movieList.append(Movie(withItem: response))
            }

            return movieList
        } catch {
            print("ERROR: batchGet", dump(error))
            throw error
        }
    }
```

- For API details, see BatchGetItem in *AWS SDK for Swift API reference*.

`BatchWriteItem`

The following code example shows how to use `BatchWriteItem`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSDynamoDB


    /// Populate the movie database from the specified JSON file.
    ///
    /// - Parameter jsonPath: Path to a JSON file containing movie data.
    ///
    func populate(jsonPath: String) async throws {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            // Create a Swift `URL` and use it to load the file into a `Data`
            // object. Then decode the JSON into an array of `Movie` objects.

            let fileUrl = URL(fileURLWithPath: jsonPath)
```

```
            let jsonData = try Data(contentsOf: fileUrl)

            var movieList = try JSONDecoder().decode([Movie].self, from: jsonData)

            // Truncate the list to the first 200 entries or so for this example.

            if movieList.count > 200 {
                movieList = Array(movieList[...199])
            }

            // Before sending records to the database, break the movie list into
            // 25-entry chunks, which is the maximum size of a batch item request.

            let count = movieList.count
            let chunks = stride(from: 0, to: count, by: 25).map {
                Array(movieList[$0 ..< Swift.min($0 + 25, count)])
            }

            // For each chunk, create a list of write request records and populate
            // them with `PutRequest` requests, each specifying one movie from the
            // chunk. Once the chunk's items are all in the `PutRequest` list,
            // send them to Amazon DynamoDB using the
            // `DynamoDBClient.batchWriteItem()` function.

            for chunk in chunks {
                var requestList: [DynamoDBClientTypes.WriteRequest] = []

                for movie in chunk {
                    let item = try await movie.getAsItem()
                    let request = DynamoDBClientTypes.WriteRequest(
                        putRequest: .init(
                            item: item
                        )
                    )
                    requestList.append(request)
                }

                let input = BatchWriteItemInput(requestItems: [tableName:
requestList])
                _ = try await client.batchWriteItem(input: input)
            }
        } catch {
            print("ERROR: populate:", dump(error))
            throw error
```

```
        }
    }
```

- For API details, see BatchWriteItem in *AWS SDK for Swift API reference*.

`CreateTable`

The following code example shows how to use `CreateTable`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSDynamoDB


    ///
    /// Create a movie table in the Amazon DynamoDB data store.
    ///
    private func createTable() async throws {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = CreateTableInput(
                attributeDefinitions: [
                    DynamoDBClientTypes.AttributeDefinition(attributeName: "year",
 attributeType: .n),
                    DynamoDBClientTypes.AttributeDefinition(attributeName: "title",
 attributeType: .s)
                ],
                billingMode: DynamoDBClientTypes.BillingMode.payPerRequest,
                keySchema: [
                    DynamoDBClientTypes.KeySchemaElement(attributeName: "year",
 keyType: .hash),
```

```
                    DynamoDBClientTypes.KeySchemaElement(attributeName: "title",
 keyType: .range)
                ],
                tableName: self.tableName
            )
            let output = try await client.createTable(input: input)
            if output.tableDescription == nil {
                throw MoviesError.TableNotFound
            }
        } catch {
            print("ERROR: createTable:", dump(error))
            throw error
        }
    }
```

- For API details, see CreateTable in *AWS SDK for Swift API reference*.

`DeleteItem`

The following code example shows how to use `DeleteItem`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSDynamoDB


    /// Delete a movie, given its title and release year.
    ///
    /// - Parameters:
    ///   - title: The movie's title.
    ///   - year: The movie's release year.
    ///
    func delete(title: String, year: Int) async throws {
```

```
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = DeleteItemInput(
                key: [
                    "year": .n(String(year)),
                    "title": .s(title)
                ],
                tableName: self.tableName
            )
            _ = try await client.deleteItem(input: input)
        } catch {
            print("ERROR: delete:", dump(error))
            throw error
        }
    }
```

- For API details, see DeleteItem in *AWS SDK for Swift API reference*.

`DeleteTable`

The following code example shows how to use `DeleteTable`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSDynamoDB


    ///
    /// Deletes the table from Amazon DynamoDB.
    ///
    func deleteTable() async throws {
```

```
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = DeleteTableInput(
                tableName: self.tableName
            )
            _ = try await client.deleteTable(input: input)
        } catch {
            print("ERROR: deleteTable:", dump(error))
            throw error
        }
    }
```

- For API details, see DeleteTable in *AWS SDK for Swift API reference*.

`GetItem`

The following code example shows how to use `GetItem`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSDynamoDB


    /// Return a `Movie` record describing the specified movie from the Amazon
    /// DynamoDB table.
    ///
    /// - Parameters:
    ///   - title: The movie's title (`String`).
    ///   - year: The movie's release year (`Int`).
    ///
    /// - Throws: `MoviesError.ItemNotFound` if the movie isn't in the table.
```

```
    ///
    /// - Returns: A `Movie` record with the movie's details.
    func get(title: String, year: Int) async throws -> Movie {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = GetItemInput(
                key: [
                    "year": .n(String(year)),
                    "title": .s(title)
                ],
                tableName: self.tableName
            )
            let output = try await client.getItem(input: input)
            guard let item = output.item else {
                throw MoviesError.ItemNotFound
            }

            let movie = try Movie(withItem: item)
            return movie
        } catch {
            print("ERROR: get:", dump(error))
            throw error
        }
    }
```

- For API details, see GetItem in *AWS SDK for Swift API reference*.

`ListTables`

The following code example shows how to use `ListTables`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSDynamoDB


    /// Get a list of the DynamoDB tables available in the specified Region.
    ///
    /// - Returns: An array of strings listing all of the tables available
    ///   in the Region specified when the session was created.
    public func getTableList() async throws -> [String] {
        let input = ListTablesInput(
        )
        return try await session.listTables(input: input)
    }
```

- For API details, see ListTables in *AWS SDK for Swift API reference*.

`PutItem`

The following code example shows how to use `PutItem`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSDynamoDB


    /// Add a movie specified as a `Movie` structure to the Amazon DynamoDB
    /// table.
    ///
    /// - Parameter movie: The `Movie` to add to the table.
    ///
    func add(movie: Movie) async throws {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }
```

```
            // Get a DynamoDB item containing the movie data.
            let item = try await movie.getAsItem()

            // Send the `PutItem` request to Amazon DynamoDB.

            let input = PutItemInput(
                item: item,
                tableName: self.tableName
            )
            _ = try await client.putItem(input: input)
        } catch {
            print("ERROR: add movie:", dump(error))
            throw error
        }
    }


    ///
    /// Return an array mapping attribute names to Amazon DynamoDB attribute
    /// values, representing the contents of the `Movie` record as a DynamoDB
    /// item.
    ///
    /// - Returns: The movie item as an array of type
    ///   `[Swift.String:DynamoDBClientTypes.AttributeValue]`.
    ///
    func getAsItem() async throws ->
[Swift.String:DynamoDBClientTypes.AttributeValue]  {
        // Build the item record, starting with the year and title, which are
        // always present.

        var item: [Swift.String:DynamoDBClientTypes.AttributeValue] = [
            "year": .n(String(self.year)),
            "title": .s(self.title)
        ]

        // Add the `info` field with the rating and/or plot if they're
        // available.

        var details: [Swift.String:DynamoDBClientTypes.AttributeValue] = [:]
        if (self.info.rating != nil || self.info.plot != nil) {
            if self.info.rating != nil {
                details["rating"] = .n(String(self.info.rating!))
            }
```

```
            if self.info.plot != nil {
                details["plot"] = .s(self.info.plot!)
            }
        }
        item["info"] = .m(details)

        return item
    }
```

- For API details, see PutItem in *AWS SDK for Swift API reference*.

`Query`

The following code example shows how to use `Query`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSDynamoDB


    /// Get all the movies released in the specified year.
    ///
    /// - Parameter year: The release year of the movies to return.
    ///
    /// - Returns: An array of `Movie` objects describing each matching movie.
    ///
    func getMovies(fromYear year: Int) async throws -> [Movie] {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = QueryInput(
                expressionAttributeNames: [
```

```
                    "#y": "year"
                ],
                expressionAttributeValues: [
                    ":y": .n(String(year))
                ],
                keyConditionExpression: "#y = :y",
                tableName: self.tableName
            )
            // Use "Paginated" to get all the movies.
            // This lets the SDK handle the 'lastEvaluatedKey' property in
"QueryOutput".

            let pages = client.queryPaginated(input: input)

            var movieList: [Movie] = []
            for try await page in pages {
                guard let items = page.items else {
                    print("Error: no items returned.")
                    continue
                }

                // Convert the found movies into `Movie` objects and return an array
                // of them.

                for item in items {
                    let movie = try Movie(withItem: item)
                    movieList.append(movie)
                }
            }
            return movieList
        } catch {
            print("ERROR: getMovies:", dump(error))
            throw error
        }
    }
```

- For API details, see Query in *AWS SDK for Swift API reference*.

Scan

The following code example shows how to use `Scan`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSDynamoDB


    /// Return an array of `Movie` objects released in the specified range of
    /// years.
    ///
    /// - Parameters:
    ///   - firstYear: The first year of movies to return.
    ///   - lastYear: The last year of movies to return.
    ///   - startKey: A starting point to resume processing; always use `nil`.
    ///
    /// - Returns: An array of `Movie` objects describing the matching movies.
    ///
    /// > Note: The `startKey` parameter is used by this function when
    ///   recursively calling itself, and should always be `nil` when calling
    ///   directly.
    ///
    func getMovies(firstYear: Int, lastYear: Int,
                   startKey: [Swift.String: DynamoDBClientTypes.AttributeValue]? =
 nil)
        async throws -> [Movie]
    {
        do {
            var movieList: [Movie] = []

            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            let input = ScanInput(
                consistentRead: true,
                exclusiveStartKey: startKey,
                expressionAttributeNames: [
                    "#y": "year" // `year` is a reserved word, so use `#y` instead.
```

```
                ],
                expressionAttributeValues: [
                    ":y1": .n(String(firstYear)),
                    ":y2": .n(String(lastYear))
                ],
                filterExpression: "#y BETWEEN :y1 AND :y2",
                tableName: self.tableName
            )

            let pages = client.scanPaginated(input: input)

            for try await page in pages {
                guard let items = page.items else {
                    print("Error: no items returned.")
                    continue
                }

                // Build an array of `Movie` objects for the returned items.

                for item in items {
                    let movie = try Movie(withItem: item)
                    movieList.append(movie)
                }
            }
            return movieList

        } catch {
            print("ERROR: getMovies with scan:", dump(error))
            throw error
        }
    }
```

- For API details, see Scan in *AWS SDK for Swift API reference*.

`UpdateItem`

The following code example shows how to use `UpdateItem`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSDynamoDB


    /// Update the specified movie with new `rating` and `plot` information.
    ///
    /// - Parameters:
    ///   - title: The title of the movie to update.
    ///   - year: The release year of the movie to update.
    ///   - rating: The new rating for the movie.
    ///   - plot: The new plot summary string for the movie.
    ///
    /// - Returns: An array of mappings of attribute names to their new
    ///   listing each item actually changed. Items that didn't need to change
    ///   aren't included in this list. `nil` if no changes were made.
    ///
    func update(title: String, year: Int, rating: Double? = nil, plot: String? =
 nil) async throws
        -> [Swift.String: DynamoDBClientTypes.AttributeValue]?
    {
        do {
            guard let client = self.ddbClient else {
                throw MoviesError.UninitializedClient
            }

            // Build the update expression and the list of expression attribute
            // values. Include only the information that's changed.

            var expressionParts: [String] = []
            var attrValues: [Swift.String: DynamoDBClientTypes.AttributeValue] = [:]

            if rating != nil {
                expressionParts.append("info.rating=:r")
                attrValues[":r"] = .n(String(rating!))
            }
```

```
            if plot != nil {
                expressionParts.append("info.plot=:p")
                attrValues[":p"] = .s(plot!)
            }
            let expression = "set \(expressionParts.joined(separator: ", "))"

            let input = UpdateItemInput(
                // Create substitution tokens for the attribute values, to ensure
                // no conflicts in expression syntax.
                expressionAttributeValues: attrValues,
                // The key identifying the movie to update consists of the release
                // year and title.
                key: [
                    "year": .n(String(year)),
                    "title": .s(title)
                ],
                returnValues: .updatedNew,
                tableName: self.tableName,
                updateExpression: expression
            )
            let output = try await client.updateItem(input: input)

            guard let attributes: [Swift.String: DynamoDBClientTypes.AttributeValue]
 = output.attributes else {
                throw MoviesError.InvalidAttributes
            }
            return attributes
        } catch {
            print("ERROR: update:", dump(error))
            throw error
        }
    }
```

- For API details, see UpdateItem in *AWS SDK for Swift API reference*.

Amazon EC2 examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with Amazon EC2.

Basics are code examples that show you how to perform the essential operations within a service.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Get started
- Basics
- Actions

Get started

Hello Amazon EC2

The following code example shows how to get started using Amazon EC2.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

The `Package.swift` file.

```
// swift-tools-version: 5.9
//
// The swift-tools-version declares the minimum version of Swift required to
// build this package.

import PackageDescription

let package = Package(
    name: "hello-ec2",
    // Let Xcode know the minimum Apple platforms supported.
    platforms: [
```

```
        .macOS(.v13),
        .iOS(.v15)
    ],
    dependencies: [
        // Dependencies declare other packages that this package depends on.
        .package(
            url: "https://github.com/awslabs/aws-sdk-swift",
            from: "1.0.0"),
        .package(
            url: "https://github.com/apple/swift-argument-parser.git",
            branch: "main"
        )
    ],
    targets: [
        // Targets are the basic building blocks of a package, defining a module or
 a test suite.
        // Targets can depend on other targets in this package and products
        // from dependencies.
        .executableTarget(
            name: "hello-ec2",
            dependencies: [
                .product(name: "AWSEC2", package: "aws-sdk-swift"),
                .product(name: "ArgumentParser", package: "swift-argument-parser")
            ],
            path: "Sources")

    ]
)
```

The `entry.swift` file.

```
// An example that shows how to use the AWS SDK for Swift to perform a simple
// operation using Amazon Elastic Compute Cloud (EC2).
//

import ArgumentParser
import Foundation

import AWSEC2

struct ExampleCommand: ParsableCommand {
    @Option(help: "The AWS Region to run AWS API calls in.")
```

```
    var awsRegion = "us-east-1"

    @Option(
        help: ArgumentHelp("The level of logging for the Swift SDK to perform."),
        completion: .list([
            "critical",
            "debug",
            "error",
            "info",
            "notice",
            "trace",
            "warning"
        ])
    )
    var logLevel: String = "error"

    static var configuration = CommandConfiguration(
        commandName: "hello-ec2",
        abstract: """
        Demonstrates a simple operation using Amazon EC2.
        """,
        discussion: """
        An example showing how to make a call to Amazon EC2 using the AWS SDK for
Swift.
        """
    )

    /// Return an array of strings giving the names of every security group
    /// the user is a member of.
    ///
    /// - Parameter ec2Client: The `EC2Client` to use when calling
    ///   `describeSecurityGroupsPaginated()`.
    ///
    /// - Returns: An array of strings giving the names of every security
    ///   group the user is a member of.
    func getSecurityGroupNames(ec2Client: EC2Client) async -> [String] {
        let pages = ec2Client.describeSecurityGroupsPaginated(
            input: DescribeSecurityGroupsInput()
        )

        var groupNames: [String] = []

        do {
            for try await page in pages {
```

```
                guard let groups = page.securityGroups else {
                    print("*** Error: No groups returned.")
                    continue
                }

                for group in groups {
                    groupNames.append(group.groupName ?? "<unknown>")
                }
            }
        } catch {
            print("*** Error: \(error.localizedDescription)")
        }

        return groupNames
    }

    /// Called by ``main()`` to run the bulk of the example.
    func runAsync() async throws {
        let ec2Config = try await EC2Client.EC2ClientConfig(region: awsRegion)
        let ec2Client = EC2Client(config: ec2Config)

        let groupNames = await getSecurityGroupNames(ec2Client: ec2Client)

        print("Found \(groupNames.count) security group(s):")

        for group in groupNames {
            print("    \(group)")
        }
    }
}

/// The program's asynchronous entry point.
@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.runAsync()
        } catch {
            ExampleCommand.exit(withError: error)
        }
    }
```

```
}
```

- For API details, see DescribeSecurityGroups in *AWS SDK for Swift API reference*.

Basics

Learn the basics

The following code example shows how to:

- Create a key pair and security group.
- Select an Amazon Machine Image (AMI) and compatible instance type, then create an instance.
- Stop and restart the instance.
- Associate an Elastic IP address with your instance.
- Connect to your instance with SSH, then clean up resources.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

The `Package.swift` file.

```
// swift-tools-version: 5.9
//
// The swift-tools-version declares the minimum version of Swift required to
// build this package.

import PackageDescription

let package = Package(
    name: "ec2-scenario",
    // Let Xcode know the minimum Apple platforms supported.
    platforms: [
        .macOS(.v13),
```

```
        .iOS(.v15)
    ],
    dependencies: [
        // Dependencies declare other packages that this package depends on.
        .package(
            url: "https://github.com/awslabs/aws-sdk-swift",
            from: "1.4.0"),
        .package(
            url: "https://github.com/apple/swift-argument-parser.git",
            branch: "main"
        )
    ],
    targets: [
        // Targets are the basic building blocks of a package, defining a module or
 a test suite.
        // Targets can depend on other targets in this package and products
        // from dependencies.
        .executableTarget(
            name: "ec2-scenario",
            dependencies: [
                .product(name: "AWSEC2", package: "aws-sdk-swift"),
                .product(name: "AWSSSM", package: "aws-sdk-swift"),
                .product(name: "ArgumentParser", package: "swift-argument-parser")
            ],
            path: "Sources")

    ]
)
```

The `entry.swift` file.

```
// An example that shows how to use the AWS SDK for Swift to perform a variety
// of operations using Amazon Elastic Compute Cloud (EC2).
//

import ArgumentParser
import Foundation
import AWSEC2

// Allow waiters to be used.

import class SmithyWaitersAPI.Waiter
```

```
import struct SmithyWaitersAPI.WaiterOptions

import AWSSSM

struct ExampleCommand: ParsableCommand {
    @Option(help: "The AWS Region to run AWS API calls in.")
    var awsRegion = "us-east-1"

    @Option(
        help: ArgumentHelp("The level of logging for the Swift SDK to perform."),
        completion: .list([
            "critical",
            "debug",
            "error",
            "info",
            "notice",
            "trace",
            "warning"
        ])
    )
    var logLevel: String = "error"

    static var configuration = CommandConfiguration(
        commandName: "ec2-scenario",
        abstract: """
        Performs various operations to demonstrate the use of Amazon EC2 using the
        AWS SDK for Swift.
        """,
        discussion: """
        """
    )

    /// Called by ``main()`` to run the bulk of the example.
    func runAsync() async throws {
        let ssmConfig = try await SSMClient.SSMClientConfig(region: awsRegion)
        let ssmClient = SSMClient(config: ssmConfig)

        let ec2Config = try await EC2Client.EC2ClientConfig(region: awsRegion)
        let ec2Client = EC2Client(config: ec2Config)

        let example = Example(ec2Client: ec2Client, ssmClient: ssmClient)

        await example.run()
    }
```

```
}

class Example {
    let ec2Client: EC2Client
    let ssmClient: SSMClient

    // Storage for AWS EC2 properties.

    var keyName: String? = nil
    var securityGroupId: String? = nil
    var instanceId: String? = nil
    var allocationId: String? = nil
    var associationId: String? = nil

    init(ec2Client: EC2Client, ssmClient: SSMClient) {
        self.ec2Client = ec2Client
        self.ssmClient = ssmClient
    }

    /// The example's main body.
    func run() async {
        //=====================================================================
        // 1. Create an RSA key pair, saving the private key as a `.pem` file.
        //    Create a `defer` block that will delete the private key when the
        //    program exits.
        //=====================================================================

        print("Creating an RSA key pair...")

        keyName = self.tempName(prefix: "ExampleKeyName")
        let keyUrl = await self.createKeyPair(name: keyName!)

        guard let keyUrl else {
            print("*** Failed to create the key pair!")
            return
        }

        print("Created the private key at: \(keyUrl.absoluteString)")

        // Schedule deleting the private key file to occur automatically when
        // the program exits, no matter how it exits.

        defer {
            do {
```

```
                try FileManager.default.removeItem(at: keyUrl)
            } catch {
                print("*** Failed to delete the private key at
\(keyUrl.absoluteString)")
            }
        }

        //=====================================================================
        // 2. List the key pairs by calling `DescribeKeyPairs`.
        //=====================================================================

        print("Describing available key pairs...")
        await self.describeKeyPairs()

        //=====================================================================
        // 3. Create a security group for the default VPC, and add an inbound
        //    rule to allow SSH from the current computer's public IPv4
        //    address.
        //=====================================================================

        print("Creating the security group...")

        let secGroupName = self.tempName(prefix: "ExampleSecurityGroup")
        let ipAddress = self.getMyIPAddress()

        guard let ipAddress else {
            print("*** Unable to get the device's IP address.")
            return
        }

        print("IP address is: \(ipAddress)")

        securityGroupId = await self.createSecurityGroup(
            name: secGroupName,
            description: "An example security group created using the AWS SDK for
Swift"
        )

        if securityGroupId == nil {
            await cleanUp()
            return
        }

        print("Created security group: \(securityGroupId ?? "<unknown>")")
```

```
        if !(await self.authorizeSecurityGroupIngress(groupId: securityGroupId!,
ipAddress: ipAddress)) {
            await cleanUp()
            return
        }

        //=====================================================================
        // 4. Display security group information for the new security group
        //    using DescribeSecurityGroups.
        //=====================================================================

        if !(await self.describeSecurityGroups(groupId: securityGroupId!)) {
            await cleanUp()
            return
        }

        //=====================================================================
        // 5. Get a list of Amazon Linux 2023 AMIs and pick one (SSM is the
        //    best practice), using path and then filter the list after the
        //    fact to include "al2023" in the Name field
        //    (ssm.GetParametersByPath). Paginate to get all images.
        //=====================================================================

        print("Searching available images for Amazon Linux 2023 images...")

        let options = await self.findAMIsMatchingFilter("al2023")

        //=====================================================================
        // 6. The information in the AMI options isn't great, so make a list
        //    of the image IDs (the "Value" field in the AMI options) and get
        //    more information about them from EC2. Display the Description
        //    field and select one of them (DescribeImages with ImageIds
        //    filter).
        //=====================================================================

        print("Images matching Amazon Linux 2023:")

        var imageIds: [String] = []
        for option in options {
            guard let id = option.value else {
                continue
            }
            imageIds.append(id)
```

```
        }

        let images = await self.describeImages(imageIds)

        // This is where you would normally let the user choose which AMI to
        // use. However, for this example, we're just going to use the first
        // one, whatever it is.

        let chosenImage = images[0]

        //=====================================================================
        // 7. Get a list of instance types that are compatible with the
        //    selected AMI's architecture (such as "x86_64") and are either
        //    small or micro. Select one (DescribeInstanceTypes).
        //=====================================================================

        print("Getting the instance types compatible with the selected image...")

        guard let arch = chosenImage.architecture else {
            print("*** The selected image doesn't have a valid architecture.")
            await cleanUp()
            return
        }

        let imageTypes = await self.getMatchingInstanceTypes(architecture: arch)

        for type in imageTypes {
            guard let instanceType = type.instanceType else {
                continue
            }
            print("    \(instanceType.rawValue)")
        }

        // This example selects the first returned instance type. A real-world
        // application would probably ask the user to select one here.

        let chosenInstanceType = imageTypes[0]

        //=====================================================================
        // 8. Create an instance with the key pair, security group, AMI, and
        //    instance type (RunInstances).
        //=====================================================================

        print("Creating an instance...")
```

```
        guard let imageId = chosenImage.imageId else {
            print("*** Cannot start image without a valid image ID.")
            await cleanUp()
            return
        }
        guard let instanceType = chosenInstanceType.instanceType else {
            print("*** Unable to start image without a valid image type.")
            await cleanUp()
            return
        }

        let instance = await self.runInstance(
            imageId: imageId,
            instanceType: instanceType,
            keyPairName: keyName!,
            securityGroups: [securityGroupId!]
        )

        guard let instance else {
            await cleanUp()
            return
        }

        instanceId = instance.instanceId
        if instanceId == nil {
            print("*** Instance is missing an ID. Canceling.")
            await cleanUp()
            return
        }

        //=====================================================================
        // 9. Wait for the instance to be ready and then display its
        //    information (DescribeInstances).
        //=====================================================================

        print("Waiting a few seconds to let the instance come up...")

        do {
            try await Task.sleep(for: .seconds(20))
        } catch {
            print("*** Error pausing the task.")
        }
        print("Success! Your new instance is ready:")
```

```
        //=====================================================================
        // 10. Display SSH connection info for the instance.
        //=====================================================================

        var runningInstance = await self.describeInstance(instanceId: instanceId!)

        if (runningInstance != nil) && (runningInstance!.publicIpAddress != nil) {
            print("\nYou can SSH to this instance using the following command:")
            print("ssh -i \(keyUrl.path) ec2-user@
\(runningInstance!.publicIpAddress!)")
        }

        //=====================================================================
        // 11. Stop the instance and wait for it to stop (StopInstances).
        //=====================================================================

        print("Stopping the instance...")

        if !(await self.stopInstance(instanceId: instanceId!, waitUntilStopped:
 true)) {
            await cleanUp()
            return
        }

        //=====================================================================
        // 12. Start the instance and wait for it to start (StartInstances).
        //=====================================================================

        print("Starting the instance again...")

        if !(await self.startInstance(instanceId: instanceId!, waitUntilStarted:
 true)) {
            await cleanUp()
            return
        }

        //=====================================================================
        // 13. Display SSH connection info for the instance. Note that it's
        //     changed.
        //=====================================================================

        runningInstance = await self.describeInstance(instanceId: instanceId!)
        if (runningInstance != nil) && (runningInstance!.publicIpAddress != nil) {
```

```
            print("\nYou can SSH to this instance using the following command.")
            print("This is probably different from when the instance was running
 before.")
            print("ssh -i \(keyUrl.path) ec2-user@
\(runningInstance!.publicIpAddress!)")
        }

        //======================================================================
        // 14. Allocate an elastic IP and associate it with the instance
        //     (AllocateAddress and AssociateAddress).
        //======================================================================

        allocationId = await self.allocateAddress()

        if allocationId == nil {
            await cleanUp()
            return
        }

        associationId = await self.associateAddress(instanceId: instanceId!,
 allocationId: allocationId)

        if associationId == nil {
            await cleanUp()
            return
        }

        //======================================================================
        // 15. Display SSH connection info for the connection. Note that the
        //     public IP is now the Elastic IP, which stays constant.
        //======================================================================

        runningInstance = await self.describeInstance(instanceId: instanceId!)
        if (runningInstance != nil) && (runningInstance!.publicIpAddress != nil) {
            print("\nYou can SSH to this instance using the following command.")
            print("This has changed again, and is now the Elastic IP.")
            print("ssh -i \(keyUrl.path) ec2-user@
\(runningInstance!.publicIpAddress!)")
        }

        //======================================================================
        // Handle all cleanup tasks
        //======================================================================
```

```
        await cleanUp()
    }

    /// Clean up by discarding and closing down all allocated EC2 items:
    ///
    /// * Elastic IP allocation and association
    /// * Terminate the instance
    /// * Delete the security group
    /// * Delete the key pair
    func cleanUp() async {
        //======================================================================
        // 16. Disassociate and delete the Elastic IP (DisassociateAddress and
        //     ReleaseAddress).
        //======================================================================

        if associationId != nil {
            await self.disassociateAddress(associationId: associationId!)
        }

        if allocationId != nil {
            await self.releaseAddress(allocationId: allocationId!)
        }

        //======================================================================
        // 17. Terminate the instance and wait for it to terminate
        //     (TerminateInstances).
        //======================================================================

        if instanceId != nil {
            print("Terminating the instance...")
            _ = await self.terminateInstance(instanceId: instanceId!,
waitUntilTerminated: true)
        }

        //======================================================================
        // 18. Delete the security group (DeleteSecurityGroup).
        //======================================================================

        if securityGroupId != nil {
            print("Deleting the security group...")
            _ = await self.deleteSecurityGroup(groupId: securityGroupId!)
        }

        //======================================================================
```

```
        // 19. Delete the key pair (DeleteKeyPair).
        //=====================================================================

        if keyName != nil {
            print("Deleting the key pair...")
            _ = await self.deleteKeyPair(keyPair: keyName!)
        }
    }

    /// Create a new RSA key pair and save the private key to a randomly-named
    /// file in the temporary directory.
    ///
    /// - Parameter name: The name of the key pair to create.
    ///
    /// - Returns: The URL of the newly created `.pem` file or `nil` if unable
    ///   to create the key pair.
    func createKeyPair(name: String) async -> URL? {
        do {
            let output = try await ec2Client.createKeyPair(
                input: CreateKeyPairInput(
                    keyName: name
                )
            )

            guard let keyMaterial = output.keyMaterial else {
                return nil
            }

            // Build the URL of the temporary private key file.

            let fileURL = URL.temporaryDirectory
                                .appendingPathComponent(name)
                                .appendingPathExtension("pem")

            do {
                try keyMaterial.write(to: fileURL, atomically: true, encoding:
 String.Encoding.utf8)
                return fileURL
            } catch {
                print("*** Failed to write the private key.")
                return nil
            }
        } catch {
            print("*** Unable to create the key pair.")
```

```
            return nil
        }
    }

    /// Describe the key pairs associated with the user by outputting each key
    /// pair's name and fingerprint.
    func describeKeyPairs() async {
        do {
            let output = try await ec2Client.describeKeyPairs(
                input: DescribeKeyPairsInput()
            )

            guard let keyPairs = output.keyPairs else {
                print("*** No key pairs list available.")
                return
            }

            for keyPair in keyPairs {
                print(keyPair.keyName ?? "<unknown>", ":", keyPair.keyFingerprint ??
"<unknown>")
            }
        } catch {
            print("*** Error: Unable to obtain a key pair list.")
        }
    }

    /// Delete an EC2 key pair.
    ///
    /// - Parameter keyPair: The name of the key pair to delete.
    ///
    /// - Returns: `true` if the key pair is deleted successfully; otherwise
    ///   `false`.
    func deleteKeyPair(keyPair: String) async -> Bool {
        do {
            _ = try await ec2Client.deleteKeyPair(
                input: DeleteKeyPairInput(
                    keyName: keyPair
                )
            )

            return true
        } catch {
            print("*** Error deleting the key pair: \(error.localizedDescription)")
            return false
```

```
        }
    }

    /// Return a list of AMI names that contain the specified string.
    ///
    /// - Parameter filter: A string that must be contained in all returned
    ///   AMI names.
    ///
    /// - Returns: An array of the parameters matching the specified substring.
    func findAMIsMatchingFilter(_ filter: String) async ->
[SSMClientTypes.Parameter] {
        var parameterList: [SSMClientTypes.Parameter] = []
        var matchingAMIs: [SSMClientTypes.Parameter] = []

        do {
            let pages = ssmClient.getParametersByPathPaginated(
                input: GetParametersByPathInput(
                    path: "/aws/service/ami-amazon-linux-latest"
                )
            )

            for try await page in pages {
                guard let parameters = page.parameters else {
                    return matchingAMIs
                }

                for parameter in parameters {
                    parameterList.append(parameter)
                }
            }

            print("Found \(parameterList.count) images total:")
            for parameter in parameterList {
                guard let name = parameter.name else {
                    continue
                }
                print("    \(name)")

                if name.contains(filter) {
                    matchingAMIs.append(parameter)
                }
            }
        } catch {
            return matchingAMIs
```

```
        }

        return matchingAMIs
    }

    /// Return a list of instance types matching the specified architecture
    /// and instance sizes.
    ///
    /// - Parameters:
    ///   - architecture: The architecture of the instance types to return, as
    ///     a member of `EC2ClientTypes.ArchitectureValues`.
    ///   - sizes: An array of one or more strings identifying sizes of
    ///     instance type to accept.
    ///
    /// - Returns: An array of `EC2ClientTypes.InstanceTypeInfo` records
    ///   describing the instance types matching the given requirements.
    func getMatchingInstanceTypes(architecture: EC2ClientTypes.ArchitectureValues =
EC2ClientTypes.ArchitectureValues.x8664,
                          sizes: [String] = ["*.micro", "*.small"]) async
                          -> [EC2ClientTypes.InstanceTypeInfo] {
        var instanceTypes: [EC2ClientTypes.InstanceTypeInfo] = []

        let archFilter = EC2ClientTypes.Filter(
            name: "processor-info.supported-architecture",
            values: [architecture.rawValue]
        )
        let sizeFilter = EC2ClientTypes.Filter(
            name: "instance-type",
            values: sizes
        )

        do {
            let pages = ec2Client.describeInstanceTypesPaginated(
                input: DescribeInstanceTypesInput(
                    filters: [archFilter, sizeFilter]
                )
            )

            for try await page in pages {
                guard let types = page.instanceTypes else {
                    return []
                }

                instanceTypes += types
```

```
            }
        } catch {
            print("*** Error getting image types: \(error.localizedDescription)")
            return []
        }

        return instanceTypes
    }

    /// Get the latest information about the specified instance and output it
    /// to the screen, returning the instance details to the caller.
    ///
    /// - Parameters:
    ///   - instanceId: The ID of the instance to provide details about.
    ///   - stateFilter: The state to require the instance to be in.
    ///
    /// - Returns: The instance's details as an `EC2ClientTypes.Instance` object.
    func describeInstance(instanceId: String,
                          stateFilter: EC2ClientTypes.InstanceStateName? =
EC2ClientTypes.InstanceStateName.running) async
                          -> EC2ClientTypes.Instance? {
        do {
            let pages = ec2Client.describeInstancesPaginated(
                input: DescribeInstancesInput(
                    instanceIds: [instanceId]
                )
            )

            for try await page in pages {
                guard let reservations = page.reservations else {
                    continue
                }

                for reservation in reservations {
                    guard let instances = reservation.instances else {
                        continue
                    }

                    for instance in instances {
                        guard let state = instance.state else {
                            print("*** Instance is missing its state...")
                            continue
                        }
                        let instanceState = state.name
```

```
                        if stateFilter != nil && (instanceState != stateFilter) {
                            continue
                        }

                        let instanceTypeName: String
                        if instance.instanceType == nil {
                            instanceTypeName = "<N/A>"
                        } else {
                            instanceTypeName = instance.instanceType?.rawValue ??
"<N/A>"
                        }

                        let instanceStateName: String
                        if instanceState == nil {
                            instanceStateName = "<N/A>"
                        } else {
                            instanceStateName = instanceState?.rawValue ?? "<N/A>"
                        }

                        print("""
                        Instance: \(instance.instanceId ?? "<N/A>")
                                • Image ID: \(instance.imageId ?? "<N/A>")
                                • Instance type: \(instanceTypeName)
                                • Key name: \(instance.keyName ?? "<N/A>")
                                • VPC ID: \(instance.vpcId ?? "<N/A>")
                                • Public IP: \(instance.publicIpAddress ?? "N/A")
                                • State: \(instanceStateName)
                        """)

                        return instance
                    }
                }
            }
        } catch {
            print("*** Error retrieving instance information to display:
\(error.localizedDescription)")
            return nil
        }

        return nil
    }

    /// Stop the specified instance.
```

```
    ///
    /// - Parameters:
    ///   - instanceId: The ID of the instance to stop.
    ///   - waitUntilStopped: If `true`, execution waits until the instance
    ///     has stopped. Otherwise, execution continues and the instance stops
    ///     asynchronously.
    ///
    /// - Returns: `true` if the image is successfully stopped (or is left to
    ///   stop asynchronously). `false` if the instance doesn't stop.
    func stopInstance(instanceId: String, waitUntilStopped: Bool = false) async ->
Bool {
        let instanceList = [instanceId]

        do {
            _ = try await ec2Client.stopInstances(
                input: StopInstancesInput(
                    instanceIds: instanceList
                )
            )

            if waitUntilStopped {
                print("Waiting for the instance to stop. Please be patient!")

                let waitOptions = WaiterOptions(maxWaitTime: 600)
                let output = try await ec2Client.waitUntilInstanceStopped(
                    options: waitOptions,
                    input: DescribeInstancesInput(
                        instanceIds: instanceList
                    )
                )

                switch output.result {
                case .success:
                    return true
                case .failure:
                    return false
                }
            } else {
                return true
            }
        } catch {
            print("*** Unable to stop the instance: \(error.localizedDescription)")
            return false
        }
```

```
    }

    /// Start the specified instance.
    ///
    /// - Parameters:
    ///   - instanceId: The ID of the instance to start.
    ///   - waitUntilStarted: If `true`, execution waits until the instance
    ///     has started. Otherwise, execution continues and the instance starts
    ///     asynchronously.
    ///
    /// - Returns: `true` if the image is successfully started (or is left to
    ///   start asynchronously). `false` if the instance doesn't start.
    func startInstance(instanceId: String, waitUntilStarted: Bool = false) async ->
 Bool {
        let instanceList = [instanceId]

        do {
            _ = try await ec2Client.startInstances(
                input: StartInstancesInput(
                    instanceIds: instanceList
                )
            )

            if waitUntilStarted {
                print("Waiting for the instance to start...")

                let waitOptions = WaiterOptions(maxWaitTime: 60.0)
                let output = try await ec2Client.waitUntilInstanceRunning(
                    options: waitOptions,
                    input: DescribeInstancesInput(
                        instanceIds: instanceList
                    )
                )
                switch output.result {
                case .success:
                    return true
                case .failure:
                    return false
                }
            } else {
                return true
            }
        } catch {
            print("*** Unable to start the instance: \(error.localizedDescription)")
```

```
            return false
        }
    }

    /// Terminate the specified instance.
    ///
    /// - Parameters:
    ///   - instanceId: The instance to terminate.
    ///   - waitUntilTerminated: Whether or not to wait until the instance is
    ///     terminated before returning.
    ///
    /// - Returns: `true` if terminated successfully. `false` if not or if an
    ///   error occurs.
    func terminateInstance(instanceId: String, waitUntilTerminated: Bool = false)
 async -> Bool {
        let instanceList = [instanceId]

        do {
            _ = try await ec2Client.terminateInstances(
                input: TerminateInstancesInput(
                    instanceIds: instanceList
                )
            )

            if waitUntilTerminated {
                print("Waiting for the instance to terminate...")

                let waitOptions = WaiterOptions(maxWaitTime: 600.0)
                let output = try await ec2Client.waitUntilInstanceTerminated(
                    options: waitOptions,
                    input: DescribeInstancesInput(
                        instanceIds: instanceList
                    )
                )

                switch output.result {
                case .success:
                    return true
                case .failure:
                    return false
                }
            } else {
                return true
            }
```

```
        } catch {
            print("*** Unable to terminate the instance:
\(error.localizedDescription)")
            return false
        }
    }

    /// Return an array of `EC2ClientTypes.Image` objects describing all of
    /// the images in the specified array.
    ///
    /// - Parameter idList: A list of image ID strings indicating the images
    ///   to return details about.
    ///
    /// - Returns: An array of the images.
    func describeImages(_ idList: [String]) async -> [EC2ClientTypes.Image] {
        do {
            let output = try await ec2Client.describeImages(
                input: DescribeImagesInput(
                    imageIds: idList
                )
            )

            guard let images = output.images else {
                print("*** No images found.")
                return []
            }

            for image in images {
                guard let id = image.imageId else {
                    continue
                }
                print("   \(id): \(image.description ?? "<no description>")")
            }

            return images
        } catch {
            print("*** Error getting image descriptions:
\(error.localizedDescription)")
            return []
        }
    }

    /// Create and return a new EC2 instance.
    ///
```

```
    /// - Parameters:
    ///   - imageId: The image ID of the AMI to use when creating the instance.
    ///   - instanceType: The type of instance to create.
    ///   - keyPairName: The RSA key pair's name to use to secure the instance.
    ///   - securityGroups: The security group or groups to add the instance
    ///     to.
    ///
    /// - Returns: The EC2 instance as an `EC2ClientTypes.Instance` object.
    func runInstance(imageId: String, instanceType: EC2ClientTypes.InstanceType,
                     keyPairName: String, securityGroups: [String]?) async ->
EC2ClientTypes.Instance? {
        do {
            let output = try await ec2Client.runInstances(
                input: RunInstancesInput(
                    imageId: imageId,
                    instanceType: instanceType,
                    keyName: keyPairName,
                    maxCount: 1,
                    minCount: 1,
                    securityGroupIds: securityGroups
                )
            )

            guard let instances = output.instances else {
                print("*** Unable to create the instance.")
                return nil
            }

            return instances[0]
        } catch {
            print("*** Error creating the instance: \(error.localizedDescription)")
            return nil
        }
    }

    /// Return the device's external IP address.
    ///
    /// - Returns: A string containing the device's IP address.
    func getMyIPAddress() -> String? {
        guard let url = URL(string: "http://checkip.amazonaws.com") else {
            print("Couldn't create the URL")
            return nil
        }
```

```
        do {
            print("Getting the IP address...")
            return try String(contentsOf: url, encoding:
String.Encoding.utf8).trim()
        } catch {
            print("*** Unable to get your public IP address.")
            return nil
        }
    }

    /// Create a new security group.
    ///
    /// - Parameters:
    ///   - groupName: The name of the group to create.
    ///   - groupDescription: A description of the new security group.
    ///
    /// - Returns: The ID string of the new security group.
    func createSecurityGroup(name groupName: String, description groupDescription:
String) async -> String? {
        do {
            let output = try await ec2Client.createSecurityGroup(
                input: CreateSecurityGroupInput(
                    description: groupDescription,
                    groupName: groupName
                )
            )

            return output.groupId
        } catch {
            print("*** Error creating the security group:
\(error.localizedDescription)")
            return nil
        }
    }

    /// Authorize ingress of connections for the security group.
    ///
    /// - Parameters:
    ///   - groupId: The group ID of the security group to authorize access for.
    ///   - ipAddress: The IP address of the device to grant access to.
    ///
    /// - Returns: `true` if access is successfully granted; otherwise `false`.
    func authorizeSecurityGroupIngress(groupId: String, ipAddress: String) async ->
Bool {
```

```
        let ipRange = EC2ClientTypes.IpRange(cidrIp: "\(ipAddress)/0")
        let httpPermission = EC2ClientTypes.IpPermission(
            fromPort: 80,
            ipProtocol: "tcp",
            ipRanges: [ipRange],
            toPort: 80
        )

        let sshPermission = EC2ClientTypes.IpPermission(
            fromPort: 22,
            ipProtocol: "tcp",
            ipRanges: [ipRange],
            toPort: 22
        )

        do {
            _ = try await ec2Client.authorizeSecurityGroupIngress(
                input: AuthorizeSecurityGroupIngressInput(
                    groupId: groupId,
                    ipPermissions: [httpPermission, sshPermission]
                )
            )

            return true
        } catch {
            print("*** Error authorizing ingress for the security group:
\(error.localizedDescription)")
            return false
        }
    }

    func describeSecurityGroups(groupId: String) async -> Bool {
        do {
            let output = try await ec2Client.describeSecurityGroups(
                input: DescribeSecurityGroupsInput(
                    groupIds: [groupId]
                )
            )

            guard let securityGroups = output.securityGroups else {
                print("No security groups found.")
                return true
            }
```

```
            for group in securityGroups {
                print("Group \(group.groupId ?? "<unknown>") found with VPC
\(group.vpcId ?? "<unknown>")")
            }
            return true
        } catch {
            print("*** Error getting security group details:
\(error.localizedDescription)")
            return false
        }
    }

    /// Delete a security group.
    ///
    /// - Parameter groupId: The ID of the security group to delete.
    ///
    /// - Returns: `true` on successful deletion; `false` on error.
    func deleteSecurityGroup(groupId: String) async -> Bool {
        do {
            _ = try await ec2Client.deleteSecurityGroup(
                input: DeleteSecurityGroupInput(
                    groupId: groupId
                )
            )

            return true
        } catch {
            print("*** Error deleting the security group:
\(error.localizedDescription)")
            return false
        }
    }

    /// Allocate an Elastic IP address.
    ///
    /// - Returns: A string containing the ID of the Elastic IP.
    func allocateAddress() async -> String? {
        do {
            let output = try await ec2Client.allocateAddress(
                input: AllocateAddressInput(
                    domain: EC2ClientTypes.DomainType.vpc
                )
            )
```

```
            guard let allocationId = output.allocationId else {
                return nil
            }

            return allocationId
        } catch {
            print("*** Unable to allocate the IP address:
\(error.localizedDescription)")
            return nil
        }
    }

    /// Associate the specified allocated Elastic IP to a given instance.
    ///
    /// - Parameters:
    ///   - instanceId: The instance to associate the Elastic IP with.
    ///   - allocationId: The ID of the allocated Elastic IP to associate with
    ///     the instance.
    ///
    /// - Returns: The association ID of the association.
    func associateAddress(instanceId: String?, allocationId: String?) async ->
String? {
        do {
            let output = try await ec2Client.associateAddress(
                input: AssociateAddressInput(
                    allocationId: allocationId,
                    instanceId: instanceId
                )
            )

            return output.associationId
        } catch {
            print("*** Unable to associate the IP address:
\(error.localizedDescription)")
            return nil
        }
    }

    /// Disassociate an Elastic IP.
    ///
    /// - Parameter associationId: The ID of the association to end.
    func disassociateAddress(associationId: String?) async {
        do {
            _ = try await ec2Client.disassociateAddress(
```

```
                input: DisassociateAddressInput(
                    associationId: associationId
                )
            )
        } catch {
            print("*** Unable to disassociate the IP address:
 \(error.localizedDescription)")
        }
    }

    /// Release an allocated Elastic IP.
    ///
    /// - Parameter allocationId: The allocation ID of the Elastic IP to
    ///   release.
    func releaseAddress(allocationId: String?) async {
        do {
            _ = try await ec2Client.releaseAddress(
                input: ReleaseAddressInput(
                    allocationId: allocationId
                )
            )
        } catch {
            print("*** Unable to release the IP address:
 \(error.localizedDescription)")
        }
    }

    /// Generate and return a unique file name that begins with the specified
    /// string.
    ///
    /// - Parameters:
    ///   - prefix: Text to use at the beginning of the returned name.
    ///
    /// - Returns: A string containing a unique filename that begins with the
    ///   specified `prefix`.
    ///
    /// The returned name uses a random number between 1 million and 1 billion to
    /// provide reasonable certainty of uniqueness for the purposes of this
    /// example.
    func tempName(prefix: String) -> String {
        return "\(prefix)-\(Int.random(in: 1000000..<1000000000))"
    }
}
```

```
/// The program's asynchronous entry point.
@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.runAsync()
        } catch {
            ExampleCommand.exit(withError: error)
        }
    }
}
```

- For API details, see the following topics in *AWS SDK for Swift API reference*.
  - AllocateAddress
  - AssociateAddress
  - AuthorizeSecurityGroupIngress
  - CreateKeyPair
  - CreateSecurityGroup
  - DeleteKeyPair
  - DeleteSecurityGroup
  - DescribeImages
  - DescribeInstanceTypes
  - DescribeInstances
  - DescribeKeyPairs
  - DescribeSecurityGroups
  - DisassociateAddress
  - ReleaseAddress
  - RunInstances
  - StartInstances
  - StopInstances
  - TerminateInstances

- UnmonitorInstances

Actions

`AllocateAddress`

The following code example shows how to use `AllocateAddress`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Allocate an Elastic IP address.
    ///
    /// - Returns: A string containing the ID of the Elastic IP.
    func allocateAddress() async -> String? {
        do {
            let output = try await ec2Client.allocateAddress(
                input: AllocateAddressInput(
                    domain: EC2ClientTypes.DomainType.vpc
                )
            )

            guard let allocationId = output.allocationId else {
                return nil
            }

            return allocationId
        } catch {
            print("*** Unable to allocate the IP address:
 \(error.localizedDescription)")
            return nil
        }
    }
```

- For API details, see AllocateAddress in *AWS SDK for Swift API reference*.

`AssociateAddress`

The following code example shows how to use `AssociateAddress`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Associate the specified allocated Elastic IP to a given instance.
    ///
    /// - Parameters:
    ///   - instanceId: The instance to associate the Elastic IP with.
    ///   - allocationId: The ID of the allocated Elastic IP to associate with
    ///     the instance.
    ///
    /// - Returns: The association ID of the association.
    func associateAddress(instanceId: String?, allocationId: String?) async ->
 String? {
        do {
            let output = try await ec2Client.associateAddress(
                input: AssociateAddressInput(
                    allocationId: allocationId,
                    instanceId: instanceId
                )
            )

            return output.associationId
        } catch {
            print("*** Unable to associate the IP address:
 \(error.localizedDescription)")
            return nil
        }
    }
```

- For API details, see AssociateAddress in *AWS SDK for Swift API reference*.

`AuthorizeSecurityGroupIngress`

The following code example shows how to use `AuthorizeSecurityGroupIngress`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Authorize ingress of connections for the security group.
    ///
    /// - Parameters:
    ///   - groupId: The group ID of the security group to authorize access for.
    ///   - ipAddress: The IP address of the device to grant access to.
    ///
    /// - Returns: `true` if access is successfully granted; otherwise `false`.
    func authorizeSecurityGroupIngress(groupId: String, ipAddress: String) async ->
 Bool {
        let ipRange = EC2ClientTypes.IpRange(cidrIp: "\(ipAddress)/0")
        let httpPermission = EC2ClientTypes.IpPermission(
            fromPort: 80,
            ipProtocol: "tcp",
            ipRanges: [ipRange],
            toPort: 80
        )

        let sshPermission = EC2ClientTypes.IpPermission(
            fromPort: 22,
            ipProtocol: "tcp",
            ipRanges: [ipRange],
            toPort: 22
        )
```

```
        do {
            _ = try await ec2Client.authorizeSecurityGroupIngress(
                input: AuthorizeSecurityGroupIngressInput(
                    groupId: groupId,
                    ipPermissions: [httpPermission, sshPermission]
                )
            )

            return true
        } catch {
            print("*** Error authorizing ingress for the security group:
\(error.localizedDescription)")
            return false
        }
    }
```

- For API details, see AuthorizeSecurityGroupIngress in *AWS SDK for Swift API reference*.

`CreateKeyPair`

The following code example shows how to use `CreateKeyPair`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Create a new RSA key pair and save the private key to a randomly-named
    /// file in the temporary directory.
    ///
    /// - Parameter name: The name of the key pair to create.
    ///
    /// - Returns: The URL of the newly created `.pem` file or `nil` if unable
    ///   to create the key pair.
```

```
    func createKeyPair(name: String) async -> URL? {
        do {
            let output = try await ec2Client.createKeyPair(
                input: CreateKeyPairInput(
                    keyName: name
                )
            )

            guard let keyMaterial = output.keyMaterial else {
                return nil
            }

            // Build the URL of the temporary private key file.

            let fileURL = URL.temporaryDirectory
                                 .appendingPathComponent(name)
                                 .appendingPathExtension("pem")

            do {
                try keyMaterial.write(to: fileURL, atomically: true, encoding:
String.Encoding.utf8)
                return fileURL
            } catch {
                print("*** Failed to write the private key.")
                return nil
            }
        } catch {
            print("*** Unable to create the key pair.")
            return nil
        }
    }
```

- For API details, see CreateKeyPair in *AWS SDK for Swift API reference*.

`CreateSecurityGroup`

The following code example shows how to use `CreateSecurityGroup`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Create a new security group.
    ///
    /// - Parameters:
    ///   - groupName: The name of the group to create.
    ///   - groupDescription: A description of the new security group.
    ///
    /// - Returns: The ID string of the new security group.
    func createSecurityGroup(name groupName: String, description groupDescription:
 String) async -> String? {
        do {
            let output = try await ec2Client.createSecurityGroup(
                input: CreateSecurityGroupInput(
                    description: groupDescription,
                    groupName: groupName
                )
            )

            return output.groupId
        } catch {
            print("*** Error creating the security group:
 \(error.localizedDescription)")
            return nil
        }
    }
```

- For API details, see CreateSecurityGroup in *AWS SDK for Swift API reference*.

`DeleteKeyPair`

The following code example shows how to use `DeleteKeyPair`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Delete an EC2 key pair.
    ///
    /// - Parameter keyPair: The name of the key pair to delete.
    ///
    /// - Returns: `true` if the key pair is deleted successfully; otherwise
    ///   `false`.
    func deleteKeyPair(keyPair: String) async -> Bool {
        do {
            _ = try await ec2Client.deleteKeyPair(
                input: DeleteKeyPairInput(
                    keyName: keyPair
                )
            )

            return true
        } catch {
            print("*** Error deleting the key pair: \(error.localizedDescription)")
            return false
        }
    }
```

- For API details, see DeleteKeyPair in *AWS SDK for Swift API reference*.

`DeleteSecurityGroup`

The following code example shows how to use `DeleteSecurityGroup`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Delete a security group.
    ///
    /// - Parameter groupId: The ID of the security group to delete.
    ///
    /// - Returns: `true` on successful deletion; `false` on error.
    func deleteSecurityGroup(groupId: String) async -> Bool {
        do {
            _ = try await ec2Client.deleteSecurityGroup(
                input: DeleteSecurityGroupInput(
                    groupId: groupId
                )
            )

            return true
        } catch {
            print("*** Error deleting the security group:
 \(error.localizedDescription)")
            return false
        }
    }
```

- For API details, see DeleteSecurityGroup in *AWS SDK for Swift API reference*.

`DescribeImages`

The following code example shows how to use `DescribeImages`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Return an array of `EC2ClientTypes.Image` objects describing all of
    /// the images in the specified array.
    ///
    /// - Parameter idList: A list of image ID strings indicating the images
    ///   to return details about.
    ///
    /// - Returns: An array of the images.
    func describeImages(_ idList: [String]) async -> [EC2ClientTypes.Image] {
        do {
            let output = try await ec2Client.describeImages(
                input: DescribeImagesInput(
                    imageIds: idList
                )
            )

            guard let images = output.images else {
                print("*** No images found.")
                return []
            }

            for image in images {
                guard let id = image.imageId else {
                    continue
                }
                print("   \(id): \(image.description ?? "<no description>")")
            }

            return images
        } catch {
            print("*** Error getting image descriptions:
 \(error.localizedDescription)")
            return []
```

```
        }
    }
```

- For API details, see DescribeImages in *AWS SDK for Swift API reference*.

`DescribeInstanceTypes`

The following code example shows how to use `DescribeInstanceTypes`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Return a list of instance types matching the specified architecture
    /// and instance sizes.
    ///
    /// - Parameters:
    ///   - architecture: The architecture of the instance types to return, as
    ///     a member of `EC2ClientTypes.ArchitectureValues`.
    ///   - sizes: An array of one or more strings identifying sizes of
    ///     instance type to accept.
    ///
    /// - Returns: An array of `EC2ClientTypes.InstanceTypeInfo` records
    ///   describing the instance types matching the given requirements.
    func getMatchingInstanceTypes(architecture: EC2ClientTypes.ArchitectureValues =
 EC2ClientTypes.ArchitectureValues.x8664,
                          sizes: [String] = ["*.micro", "*.small"]) async
                          -> [EC2ClientTypes.InstanceTypeInfo] {
        var instanceTypes: [EC2ClientTypes.InstanceTypeInfo] = []

        let archFilter = EC2ClientTypes.Filter(
            name: "processor-info.supported-architecture",
            values: [architecture.rawValue]
        )
        let sizeFilter = EC2ClientTypes.Filter(
```

```
            name: "instance-type",
            values: sizes
        )

        do {
            let pages = ec2Client.describeInstanceTypesPaginated(
                input: DescribeInstanceTypesInput(
                    filters: [archFilter, sizeFilter]
                )
            )

            for try await page in pages {
                guard let types = page.instanceTypes else {
                    return []
                }

                instanceTypes += types
            }
        } catch {
            print("*** Error getting image types: \(error.localizedDescription)")
            return []
        }

        return instanceTypes
    }
```

- For API details, see DescribeInstanceTypes in *AWS SDK for Swift API reference*.

`DescribeKeyPairs`

The following code example shows how to use `DescribeKeyPairs`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2
```

```
    /// Describe the key pairs associated with the user by outputting each key
    /// pair's name and fingerprint.
    func describeKeyPairs() async {
        do {
            let output = try await ec2Client.describeKeyPairs(
                input: DescribeKeyPairsInput()
            )

            guard let keyPairs = output.keyPairs else {
                print("*** No key pairs list available.")
                return
            }

            for keyPair in keyPairs {
                print(keyPair.keyName ?? "<unknown>", ":", keyPair.keyFingerprint ??
 "<unknown>")
            }
        } catch {
            print("*** Error: Unable to obtain a key pair list.")
        }
    }
```

- For API details, see DescribeKeyPairs in *AWS SDK for Swift API reference*.

`DescribeSecurityGroups`

The following code example shows how to use `DescribeSecurityGroups`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

Using pagination with `describeSecurityGroupsPaginated()`.

```
import AWSEC2
```

```
    /// Return an array of strings giving the names of every security group
    /// the user is a member of.
    ///
    /// - Parameter ec2Client: The `EC2Client` to use when calling
    ///   `describeSecurityGroupsPaginated()`.
    ///
    /// - Returns: An array of strings giving the names of every security
    ///   group the user is a member of.
    func getSecurityGroupNames(ec2Client: EC2Client) async -> [String] {
        let pages = ec2Client.describeSecurityGroupsPaginated(
            input: DescribeSecurityGroupsInput()
        )

        var groupNames: [String] = []

        do {
            for try await page in pages {
                guard let groups = page.securityGroups else {
                    print("*** Error: No groups returned.")
                    continue
                }

                for group in groups {
                    groupNames.append(group.groupName ?? "<unknown>")
                }
            }
        } catch {
            print("*** Error: \(error.localizedDescription)")
        }

        return groupNames
    }
```

Without pagination.

```
import AWSEC2

    func describeSecurityGroups(groupId: String) async -> Bool {
        do {
            let output = try await ec2Client.describeSecurityGroups(
                input: DescribeSecurityGroupsInput(
                    groupIds: [groupId]
```

```
                )
            )

            guard let securityGroups = output.securityGroups else {
                print("No security groups found.")
                return true
            }

            for group in securityGroups {
                print("Group \(group.groupId ?? "<unknown>") found with VPC
\(group.vpcId ?? "<unknown>")")
            }
            return true
        } catch {
            print("*** Error getting security group details:
\(error.localizedDescription)")
            return false
        }
    }
```

- For API details, see DescribeSecurityGroups in *AWS SDK for Swift API reference*.

`DisassociateAddress`

The following code example shows how to use `DisassociateAddress`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Disassociate an Elastic IP.
    ///
    /// - Parameter associationId: The ID of the association to end.
    func disassociateAddress(associationId: String?) async {
        do {
```

```
            _ = try await ec2Client.disassociateAddress(
                input: DisassociateAddressInput(
                    associationId: associationId
                )
            )
        } catch {
            print("*** Unable to disassociate the IP address:
\(error.localizedDescription)")
        }
    }
```

- For API details, see DisassociateAddress in *AWS SDK for Swift API reference*.

`ReleaseAddress`

The following code example shows how to use `ReleaseAddress`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Release an allocated Elastic IP.
    ///
    /// - Parameter allocationId: The allocation ID of the Elastic IP to
    ///   release.
    func releaseAddress(allocationId: String?) async {
        do {
            _ = try await ec2Client.releaseAddress(
                input: ReleaseAddressInput(
                    allocationId: allocationId
                )
            )
        } catch {
            print("*** Unable to release the IP address:
\(error.localizedDescription)")
```

```
        }
    }
```

- For API details, see ReleaseAddress in *AWS SDK for Swift API reference*.

`RunInstances`

The following code example shows how to use `RunInstances`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Create and return a new EC2 instance.
    ///
    /// - Parameters:
    ///   - imageId: The image ID of the AMI to use when creating the instance.
    ///   - instanceType: The type of instance to create.
    ///   - keyPairName: The RSA key pair's name to use to secure the instance.
    ///   - securityGroups: The security group or groups to add the instance
    ///     to.
    ///
    /// - Returns: The EC2 instance as an `EC2ClientTypes.Instance` object.
    func runInstance(imageId: String, instanceType: EC2ClientTypes.InstanceType,
                         keyPairName: String, securityGroups: [String]?) async ->
 EC2ClientTypes.Instance? {
        do {
            let output = try await ec2Client.runInstances(
                input: RunInstancesInput(
                    imageId: imageId,
                    instanceType: instanceType,
                    keyName: keyPairName,
                    maxCount: 1,
                    minCount: 1,
                    securityGroupIds: securityGroups
```

```
                )
            )

            guard let instances = output.instances else {
                print("*** Unable to create the instance.")
                return nil
            }

            return instances[0]
        } catch {
            print("*** Error creating the instance: \(error.localizedDescription)")
            return nil
        }
    }
```

- For API details, see RunInstances in *AWS SDK for Swift API reference*.

`StartInstances`

The following code example shows how to use `StartInstances`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Start the specified instance.
    ///
    /// - Parameters:
    ///   - instanceId: The ID of the instance to start.
    ///   - waitUntilStarted: If `true`, execution waits until the instance
    ///     has started. Otherwise, execution continues and the instance starts
    ///     asynchronously.
    ///
    /// - Returns: `true` if the image is successfully started (or is left to
```

```
    ///   start asynchronously). `false` if the instance doesn't start.
    func startInstance(instanceId: String, waitUntilStarted: Bool = false) async ->
Bool {
        let instanceList = [instanceId]

        do {
            _ = try await ec2Client.startInstances(
                input: StartInstancesInput(
                    instanceIds: instanceList
                )
            )

            if waitUntilStarted {
                print("Waiting for the instance to start...")

                let waitOptions = WaiterOptions(maxWaitTime: 60.0)
                let output = try await ec2Client.waitUntilInstanceRunning(
                    options: waitOptions,
                    input: DescribeInstancesInput(
                        instanceIds: instanceList
                    )
                )
                switch output.result {
                case .success:
                    return true
                case .failure:
                    return false
                }
            } else {
                return true
            }
        } catch {
            print("*** Unable to start the instance: \(error.localizedDescription)")
            return false
        }
    }
```

- For API details, see StartInstances in *AWS SDK for Swift API reference*.

`StopInstances`

The following code example shows how to use `StopInstances`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Stop the specified instance.
    ///
    /// - Parameters:
    ///   - instanceId: The ID of the instance to stop.
    ///   - waitUntilStopped: If `true`, execution waits until the instance
    ///     has stopped. Otherwise, execution continues and the instance stops
    ///     asynchronously.
    ///
    /// - Returns: `true` if the image is successfully stopped (or is left to
    ///   stop asynchronously). `false` if the instance doesn't stop.
    func stopInstance(instanceId: String, waitUntilStopped: Bool = false) async ->
 Bool {
        let instanceList = [instanceId]

        do {
            _ = try await ec2Client.stopInstances(
                input: StopInstancesInput(
                    instanceIds: instanceList
                )
            )

            if waitUntilStopped {
                print("Waiting for the instance to stop. Please be patient!")

                let waitOptions = WaiterOptions(maxWaitTime: 600)
                let output = try await ec2Client.waitUntilInstanceStopped(
                    options: waitOptions,
                    input: DescribeInstancesInput(
                        instanceIds: instanceList
                    )
                )
```

```
                switch output.result {
                case .success:
                    return true
                case .failure:
                    return false
                }
            } else {
                return true
            }
        } catch {
            print("*** Unable to stop the instance: \(error.localizedDescription)")
            return false
        }
    }
```

- For API details, see StopInstances in *AWS SDK for Swift API reference*.

`TerminateInstances`

The following code example shows how to use `TerminateInstances`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSEC2

    /// Terminate the specified instance.
    ///
    /// - Parameters:
    ///   - instanceId: The instance to terminate.
    ///   - waitUntilTerminated: Whether or not to wait until the instance is
    ///     terminated before returning.
    ///
    /// - Returns: `true` if terminated successfully. `false` if not or if an
    ///   error occurs.
```

```
   func terminateInstance(instanceId: String, waitUntilTerminated: Bool = false)
async -> Bool {
       let instanceList = [instanceId]

       do {
           _ = try await ec2Client.terminateInstances(
               input: TerminateInstancesInput(
                   instanceIds: instanceList
               )
           )

           if waitUntilTerminated {
               print("Waiting for the instance to terminate...")

               let waitOptions = WaiterOptions(maxWaitTime: 600.0)
               let output = try await ec2Client.waitUntilInstanceTerminated(
                   options: waitOptions,
                   input: DescribeInstancesInput(
                       instanceIds: instanceList
                   )
               )

               switch output.result {
               case .success:
                   return true
               case .failure:
                   return false
               }
           } else {
               return true
           }
       } catch {
           print("*** Unable to terminate the instance:
\(error.localizedDescription)")
           return false
       }
   }
```

- For API details, see TerminateInstances in *AWS SDK for Swift API reference*.

AWS Glue examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with AWS Glue.

Basics are code examples that show you how to perform the essential operations within a service.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Basics
- Actions

Basics

Learn the basics

The following code example shows how to:

- Create a crawler that crawls a public Amazon S3 bucket and generates a database of CSV-formatted metadata.
- List information about databases and tables in your AWS Glue Data Catalog.
- Create a job to extract CSV data from the S3 bucket, transform the data, and load JSON-formatted output into another S3 bucket.
- List information about job runs, view transformed data, and clean up resources.

For more information, see Tutorial: Getting started with AWS Glue Studio.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

The `Package.swift` file.

```
// swift-tools-version: 5.9
//
// The swift-tools-version declares the minimum version of Swift required to
// build this package.

import PackageDescription

let package = Package(
    name: "glue-scenario",
    // Let Xcode know the minimum Apple platforms supported.
    platforms: [
        .macOS(.v13),
        .iOS(.v15)
    ],
    dependencies: [
        // Dependencies declare other packages that this package depends on.
        .package(
            url: "https://github.com/awslabs/aws-sdk-swift",
            from: "1.0.0"),
        .package(
            url: "https://github.com/apple/swift-argument-parser.git",
            branch: "main"
        )
    ],
    targets: [
        // Targets are the basic building blocks of a package, defining a module or
 a test suite.
        // Targets can depend on other targets in this package and products
        // from dependencies.
        .executableTarget(
            name: "glue-scenario",
            dependencies: [
                .product(name: "AWSGlue", package: "aws-sdk-swift"),
                .product(name: "AWSS3", package: "aws-sdk-swift"),
                .product(name: "ArgumentParser", package: "swift-argument-parser")
            ],
            path: "Sources")

    ]
)
```

The Swift code file, `entry.swift`.

```
// An example that shows how to use the AWS SDK for Swift to demonstrate
// creating and using crawlers and jobs using AWS Glue.
//
// 0. Upload the Python job script to Amazon S3 so it can be used when
//    calling `startJobRun()` later.
// 1. Create a crawler, pass it the IAM role and the URL of the public Amazon
//    S3 bucket that contains the source data:
//    s3://crawler-public-us-east-1/flight/2016/csv.
// 2. Start the crawler. This takes time, so after starting it, use a loop
//    that calls `getCrawler()` until the state is "READY".
// 3. Get the database created by the crawler, and the tables in the
//    database. Display them to the user.
// 4. Create a job. Pass it the IAM role and the URL to a Python ETL script
//    previously uploaded to the user's S3 bucket.
// 5. Start a job run, passing the following custom arguments. These are
//    expected by the ETL script, so must exactly match.
//    * `--input_database: <name of the database created by the crawler>`
//    * `--input_table: <name of the table created by the crawler>`
//    * `--output_bucket_url: <URL to the scaffold bucket created for the
//      user>`
// 6. Loop and get the job run until it returns one of the following states:
//    "SUCCEEDED", "STOPPED", "FAILED", or "TIMEOUT".
// 7. Output data is stored in a group of files in the user's S3 bucket.
//    Either direct the user to their location or download a file and display
//    the results inline.
// 8. List the jobs for the user's account.
// 9. Get job run details for a job run.
// 10. Delete the demo job.
// 11. Delete the database and tables created by the example.
// 12. Delete the crawler created by the example.

import ArgumentParser
import AWSS3
import Foundation
import Smithy

import AWSClientRuntime
import AWSGlue

struct ExampleCommand: ParsableCommand {
    @Option(help: "The AWS IAM role to use for AWS Glue calls.")
    var role: String
```

```
    @Option(help: "The Amazon S3 bucket to use for this example.")
    var bucket: String

    @Option(help: "The Amazon S3 URL of the data to crawl.")
    var s3url: String = "s3://crawler-public-us-east-1/flight/2016/csv"

    @Option(help: "The Python script to run as a job with AWS Glue.")
    var script: String = "./flight_etl_job_script.py"

    @Option(help: "The AWS Region to run AWS API calls in.")
    var awsRegion = "us-east-1"

    @Option(help: "A prefix string to use when naming tables.")
    var tablePrefix = "swift-glue-basics-table"

    @Option(
        help: ArgumentHelp("The level of logging for the Swift SDK to perform."),
        completion: .list([
            "critical",
            "debug",
            "error",
            "info",
            "notice",
            "trace",
            "warning"
        ])
    )
    var logLevel: String = "error"

    static var configuration = CommandConfiguration(
        commandName: "glue-scenario",
        abstract: """
        Demonstrates various features of AWS Glue.
        """,
        discussion: """
        An example showing how to use AWS Glue to create, run, and monitor
        crawlers and jobs.
        """
    )

    /// Generate and return a unique file name that begins with the specified
    /// string.
    ///
```

```
    /// - Parameters:
    ///   - prefix: Text to use at the beginning of the returned name.
    ///
    /// - Returns: A string containing a unique filename that begins with the
    ///   specified `prefix`.
    ///
    /// The returned name uses a random number between 1 million and 1 billion to
    /// provide reasonable certainty of uniqueness for the purposes of this
    /// example.
    func tempName(prefix: String) -> String {
        return "\(prefix)-\(Int.random(in: 1000000..<1000000000))"
    }

    /// Upload a file to an Amazon S3 bucket.
    ///
    /// - Parameters:
    ///   - s3Client: The S3 client to use when uploading the file.
    ///   - path: The local path of the source file to upload.
    ///   - toBucket: The name of the S3 bucket into which to upload the file.
    ///   - key: The key (name) to give the file in the S3 bucket.
    ///
    /// - Returns: `true` if the file is uploaded successfully, otherwise `false`.
    func uploadFile(s3Client: S3Client, path: String, toBucket: String, key: String)
 async -> Bool {
        do {
            let fileData: Data = try Data(contentsOf: URL(fileURLWithPath: path))
            let dataStream = ByteStream.data(fileData)
            _ = try await s3Client.putObject(
                input: PutObjectInput(
                    body: dataStream,
                    bucket: toBucket,
                    key: key
                )
            )
        } catch {
            print("*** An unexpected error occurred uploading the script to the
 Amazon S3 bucket \"\(bucket)\".")
            return false
        }

        return true
    }

    /// Create a new AWS Glue crawler.
```

```
    ///
    /// - Parameters:
    ///   - glueClient: An AWS Glue client to use for the crawler.
    ///   - crawlerName: A name for the new crawler.
    ///   - iamRole: The name of an Amazon IAM role for the crawler to use.
    ///   - s3Path: The path of an Amazon S3 folder to use as a target location.
    ///   - cronSchedule: A `cron` schedule indicating when to run the crawler.
    ///   - databaseName: The name of an AWS Glue database to operate on.
    ///
    /// - Returns: `true` if the crawler is created successfully, otherwise `false`.
    func createCrawler(glueClient: GlueClient, crawlerName: String, iamRole: String,
                       s3Path: String, cronSchedule: String, databaseName: String)
async -> Bool {
        let s3Target = GlueClientTypes.S3Target(path: s3url)
        let targetList = GlueClientTypes.CrawlerTargets(s3Targets: [s3Target])

        do {
            _ = try await glueClient.createCrawler(
                input: CreateCrawlerInput(
                    databaseName: databaseName,
                    description: "Created by the AWS SDK for Swift Scenario Example
for AWS Glue.",
                    name: crawlerName,
                    role: iamRole,
                    schedule: cronSchedule,
                    tablePrefix: tablePrefix,
                    targets: targetList
                )
            )
        } catch _ as AlreadyExistsException {
            print("*** A crawler named \"\(crawlerName)\" already exists.")
            return false
        } catch _ as OperationTimeoutException {
            print("*** The attempt to create the AWS Glue crawler timed out.")
            return false
        } catch {
            print("*** An unexpected error occurred creating the AWS Glue crawler:
\(error.localizedDescription)")
            return false
        }

        return true
    }
```

```
    /// Delete an AWS Glue crawler.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - name: The name of the crawler to delete.
    ///
    /// - Returns: `true` if successful, otherwise `false`.
    func deleteCrawler(glueClient: GlueClient, name: String) async -> Bool {
        do {
            _ = try await glueClient.deleteCrawler(
                input: DeleteCrawlerInput(name: name)
            )
        } catch {
            return false
        }
        return true
    }

    /// Start running an AWS Glue crawler.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use when starting the crawler.
    ///   - name: The name of the crawler to start running.
    ///
    /// - Returns: `true` if the crawler is started successfully, otherwise `false`.
    func startCrawler(glueClient: GlueClient, name: String) async -> Bool {
        do {
            _ = try await glueClient.startCrawler(
                input: StartCrawlerInput(name: name)
            )
        } catch {
            print("*** An unexpected error occurred starting the crawler.")
            return false
        }

        return true
    }

    /// Get the state of the specified AWS Glue crawler.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - name: The name of the crawler whose state should be returned.
    ///
```

```
    /// - Returns: A `GlueClientTypes.CrawlerState` value describing the
    ///   state of the crawler.
    func getCrawlerState(glueClient: GlueClient, name: String) async ->
GlueClientTypes.CrawlerState {
        do {
            let output = try await glueClient.getCrawler(
                input: GetCrawlerInput(name: name)
            )

            // If the crawler or its state is `nil`, report that the crawler
            // is stopping. This may not be what you want for your
            // application but it works for this one!

            guard let crawler = output.crawler else {
                return GlueClientTypes.CrawlerState.stopping
            }
            guard let state = crawler.state else {
                return GlueClientTypes.CrawlerState.stopping
            }
            return state
        } catch {
            return GlueClientTypes.CrawlerState.stopping
        }
    }

    /// Wait until the specified crawler is ready to run.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - name: The name of the crawler to wait for.
    ///
    /// - Returns: `true` if the crawler is ready, `false` if the client is
    ///   stopping (and will therefore never be ready).
    func waitUntilCrawlerReady(glueClient: GlueClient, name: String) async -> Bool {
        while true {
            let state = await getCrawlerState(glueClient: glueClient, name: name)

            if state == .ready {
                return true
            } else if state == .stopping {
                return false
            }

            // Wait four seconds before trying again.
```

```
            do {
                try await Task.sleep(for: .seconds(4))
            } catch {
                print("*** Error pausing the task.")
            }
        }
    }

    /// Create a new AWS Glue job.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - jobName: The name to give the new job.
    ///   - role: The IAM role for the job to use when accessing AWS services.
    ///   - scriptLocation: The AWS S3 URI of the script to be run by the job.
    ///
    /// - Returns: `true` if the job is created successfully, otherwise `false`.
    func createJob(glueClient: GlueClient, name jobName: String, role: String,
                   scriptLocation: String) async -> Bool {
        let command = GlueClientTypes.JobCommand(
            name: "glueetl",
            pythonVersion: "3",
            scriptLocation: scriptLocation
        )

        do {
            _ = try await glueClient.createJob(
                input: CreateJobInput(
                    command: command,
                    description: "Created by the AWS SDK for Swift Glue basic
scenario example.",
                    glueVersion: "3.0",
                    name: jobName,
                    numberOfWorkers: 10,
                    role: role,
                    workerType: .g1x
                )
            )
        } catch {
            return false
        }
        return true
    }
```

```
/// Return a list of the AWS Glue jobs listed on the user's account.
///
/// - Parameters:
///   - glueClient: The AWS Glue client to use.
///   - maxJobs: The maximum number of jobs to return (default: 100).
///
/// - Returns: An array of strings listing the names of all available AWS
///   Glue jobs.
func listJobs(glueClient: GlueClient, maxJobs: Int = 100) async -> [String] {
    var jobList: [String] = []
    var nextToken: String?

    repeat {
        do {
            let output = try await glueClient.listJobs(
                input: ListJobsInput(
                    maxResults: maxJobs,
                    nextToken: nextToken
                )
            )

            guard let jobs = output.jobNames else {
                return jobList
            }

            jobList = jobList + jobs
            nextToken = output.nextToken
        } catch {
            return jobList
        }
    } while (nextToken != nil)

    return jobList
}

/// Delete an AWS Glue job.
///
/// - Parameters:
///   - glueClient: The AWS Glue client to use.
///   - jobName: The name of the job to delete.
///
/// - Returns: `true` if the job is successfully deleted, otherwise `false`.
func deleteJob(glueClient: GlueClient, name jobName: String) async -> Bool {
```

```
        do {
            _ = try await glueClient.deleteJob(
                input: DeleteJobInput(jobName: jobName)
            )
        } catch {
            return false
        }
        return true
    }

    /// Create an AWS Glue database.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - databaseName: The name to give the new database.
    ///   - location: The URL of the source data to use with AWS Glue.
    ///
    /// - Returns: `true` if the database is created successfully, otherwise
`false`.
    func createDatabase(glueClient: GlueClient, name databaseName: String, location:
String) async -> Bool {
        let databaseInput = GlueClientTypes.DatabaseInput(
            description: "Created by the AWS SDK for Swift Glue basic scenario
example.",
            locationUri: location,
            name: databaseName
        )

        do {
            _ = try await glueClient.createDatabase(
                input: CreateDatabaseInput(
                    databaseInput: databaseInput
                )
            )
        } catch {
            return false
        }

        return true
    }

    /// Get the AWS Glue database with the specified name.
    ///
    /// - Parameters:
```

```
    ///   - glueClient: The AWS Glue client to use.
    ///   - name: The name of the database to return.
    ///
    /// - Returns: The `GlueClientTypes.Database` object describing the
    ///   specified database, or `nil` if an error occurs or the database
    ///   isn't found.
    func getDatabase(glueClient: GlueClient, name: String) async ->
GlueClientTypes.Database? {
        do {
            let output = try await glueClient.getDatabase(
                input: GetDatabaseInput(name: name)
            )

            return output.database
        } catch {
            return nil
        }
    }

    /// Returns a list of the tables in the specified database.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - databaseName: The name of the database whose tables are to be
    ///     returned.
    ///
    /// - Returns: An array of `GlueClientTypes.Table` objects, each
    ///   describing one table in the named database. An empty array indicates
    ///   that there are either no tables in the database, or an error
    ///   occurred before any tables could be found.
    func getTablesInDatabase(glueClient: GlueClient, databaseName: String) async ->
[GlueClientTypes.Table] {
        var tables: [GlueClientTypes.Table] = []
        var nextToken: String?

        repeat {
            do {
                let output = try await glueClient.getTables(
                    input: GetTablesInput(
                        databaseName: databaseName,
                        nextToken: nextToken
                    )
                )
```

```
                guard let tableList = output.tableList else {
                    return tables
                }

                tables = tables + tableList
                nextToken = output.nextToken
            } catch {
                return tables
            }
        } while nextToken != nil

        return tables
    }

    /// Delete the specified database.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - databaseName: The name of the database to delete.
    ///   - deleteTables: A Bool indicating whether or not to delete the
    ///     tables in the database before attempting to delete the database.
    ///
    /// - Returns: `true` if the database (and optionally its tables) are
    ///   deleted, otherwise `false`.
    func deleteDatabase(glueClient: GlueClient, name databaseName: String,
                        withTables deleteTables: Bool = false) async -> Bool {
        if deleteTables {
            var tableNames: [String] = []

            // Get a list of the names of all of the tables in the database.

            let tableList = await self.getTablesInDatabase(glueClient: glueClient,
databaseName: databaseName)
            for table in tableList {
                guard let name = table.name else {
                    continue
                }
                tableNames.append(name)
            }

            // Delete the tables. If there's only one table, use
            // `deleteTable()`, otherwise, use `batchDeleteTable()`. You can
            // use `batchDeleteTable()` for a single table, but this
            // demonstrates the use of `deleteTable()`.
```

```
            if tableNames.count == 1 {
                do {
                    print("    Deleting table...")
                    _ = try await glueClient.deleteTable(
                        input: DeleteTableInput(
                            databaseName: databaseName,
                            name: tableNames[0]
                        )
                    )
                } catch {
                    print("*** Unable to delete the table.")
                }
            } else {
                do {
                    print("    Deleting tables...")
                    _ = try await glueClient.batchDeleteTable(
                        input: BatchDeleteTableInput(
                            databaseName: databaseName,
                            tablesToDelete: tableNames
                        )
                    )
                } catch {
                    print("*** Unable to delete the tables.")
                }
            }
        }

        // Delete the database itself.

        do {
            print("    Deleting the database itself...")
            _ = try await glueClient.deleteDatabase(
                input: DeleteDatabaseInput(name: databaseName)
            )
        } catch {
            print("*** Unable to delete the database.")
            return false
        }
        return true
    }

    /// Start an AWS Glue job run.
    ///
```

```
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - jobName: The name of the job to run.
    ///   - databaseName: The name of the AWS Glue database to run the job against.
    ///   - tableName: The name of the table in the database to run the job against.
    ///   - outputURL: The AWS S3 URI of the bucket location into which to
    ///     write the resulting output.
    ///
    /// - Returns: `true` if the job run is started successfully, otherwise `false`.
    func startJobRun(glueClient: GlueClient, name jobName: String, databaseName:
String,
                     tableName: String, outputURL: String) async -> String? {
        do {
            let output = try await glueClient.startJobRun(
                input: StartJobRunInput(
                    arguments: [
                        "--input_database": databaseName,
                        "--input_table": tableName,
                        "--output_bucket_url": outputURL
                    ],
                    jobName: jobName,
                    numberOfWorkers: 10,
                    workerType: .g1x
                )
            )

            guard let id = output.jobRunId else {
                return nil
            }

            return id
        } catch {
            return nil
        }
    }

    /// Return a list of the job runs for the specified job.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - jobName: The name of the job for which to return its job runs.
    ///   - maxResults: The maximum number of job runs to return (default:
    ///     1000).
    ///
```

```
    /// - Returns: An array of `GlueClientTypes.JobRun` objects describing
    ///   each job run.
    func getJobRuns(glueClient: GlueClient, name jobName: String, maxResults: Int? =
 nil) async -> [GlueClientTypes.JobRun] {
        do {
            let output = try await glueClient.getJobRuns(
                input: GetJobRunsInput(
                    jobName: jobName,
                    maxResults: maxResults
                )
            )

            guard let jobRuns = output.jobRuns else {
                print("*** No job runs found.")
                return []
            }

            return jobRuns
        } catch is EntityNotFoundException {
            print("*** The specified job name, \(jobName), doesn't exist.")
            return []
        } catch {
            print("*** Unexpected error getting job runs:")
            dump(error)
            return []
        }
    }

    /// Get information about a specific AWS Glue job run.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - jobName: The name of the job to return job run data for.
    ///   - id: The run ID of the specific job run to return.
    ///
    /// - Returns: A `GlueClientTypes.JobRun` object describing the state of
    ///   the job run, or `nil` if an error occurs.
    func getJobRun(glueClient: GlueClient, name jobName: String, id: String) async -
> GlueClientTypes.JobRun? {
        do {
            let output = try await glueClient.getJobRun(
                input: GetJobRunInput(
                    jobName: jobName,
                    runId: id
```

```
                )
            )

            return output.jobRun
        } catch {
            return nil
        }
    }

    /// Called by ``main()`` to run the bulk of the example.
    func runAsync() async throws {
        // A name to give the Python script upon upload to the Amazon S3
        // bucket.
        let scriptName = "jobscript.py"

        // Schedule string in `cron` format, as described here:
        // https://docs.aws.amazon.com/glue/latest/dg/monitor-data-warehouse-
schedule.html
        let cron = "cron(15 12 * * ? *)"

        let glueConfig = try await GlueClient.GlueClientConfig(region: awsRegion)
        let glueClient = GlueClient(config: glueConfig)

        let s3Config = try await S3Client.S3ClientConfig(region: awsRegion)
        let s3Client = S3Client(config: s3Config)

        // Create random names for things that need them.

        let crawlerName = tempName(prefix: "swift-glue-basics-crawler")
        let databaseName = tempName(prefix: "swift-glue-basics-db")

        // Create a name for the AWS Glue job.

        let jobName = tempName(prefix: "scenario-job")

        // The URL of the Python script on S3.

        let scriptURL = "s3://\(bucket)/\(scriptName)"

        print("Welcome to the AWS SDK for Swift basic scenario for AWS Glue!")

        //======================================================================
        // 0. Upload the Python script to the target bucket so it's available
        //    for use by the Amazon Glue service.
```

```
        //=====================================================================

        print("Uploading the Python script: \(script) as key \(scriptName)")
        print("Destination bucket: \(bucket)")
        if !(await uploadFile(s3Client: s3Client, path: script, toBucket: bucket,
key: scriptName)) {
            return
        }

        //=====================================================================
        // 1. Create the database and crawler using the randomized names
        //    generated previously.
        //=====================================================================

        print("Creating database \"\(databaseName)\"...")
        if !(await createDatabase(glueClient: glueClient, name: databaseName,
location: s3url)) {
            print("*** Unable to create the database.")
            return
        }

        print("Creating crawler \"\(crawlerName)\"...")
        if !(await createCrawler(glueClient: glueClient, crawlerName: crawlerName,
                                 iamRole: role, s3Path: s3url, cronSchedule: cron,
                                 databaseName: databaseName)) {
            return
        }

        //=====================================================================
        // 2. Start the crawler, then wait for it to be ready.
        //=====================================================================

        print("Starting the crawler and waiting until it's ready...")
        if !(await startCrawler(glueClient: glueClient, name: crawlerName)) {
            _ = await deleteCrawler(glueClient: glueClient, name: crawlerName)
            return
        }

        if !(await waitUntilCrawlerReady(glueClient: glueClient, name: crawlerName))
{
            _ = await deleteCrawler(glueClient: glueClient, name: crawlerName)
        }

        //=====================================================================
```

```
        // 3. Get the database and table created by the crawler.
        //=====================================================================

        print("Getting the crawler's database...")
        let database = await getDatabase(glueClient: glueClient, name: databaseName)

        guard let database else {
            print("*** Unable to get the database.")
            return
        }
        print("Database URI: \(database.locationUri ?? "<unknown>")")

        let tableList = await getTablesInDatabase(glueClient: glueClient,
databaseName: databaseName)

        print("Found \(tableList.count) table(s):")
        for table in tableList {
            print("  \(table.name ?? "<unnamed>")")
        }

        if tableList.count != 1 {
            print("*** Incorrect number of tables found. There should only be one.")
            _ = await deleteDatabase(glueClient: glueClient, name: databaseName,
withTables: true)
            _ = await deleteCrawler(glueClient: glueClient, name: crawlerName)
            return
        }

        guard let tableName = tableList[0].name else {
            print("*** Table is unnamed.")
            _ = await deleteDatabase(glueClient: glueClient, name: databaseName,
withTables: true)
            _ = await deleteCrawler(glueClient: glueClient, name: crawlerName)
            return
        }

        //=====================================================================
        // 4. Create a job.
        //=====================================================================

        print("Creating a job...")
        if !(await createJob(glueClient: glueClient, name: jobName, role: role,
                             scriptLocation: scriptURL)) {
```

```
            _ = await deleteDatabase(glueClient: glueClient, name: databaseName,
withTables: true)
            _ = await deleteCrawler(glueClient: glueClient, name: crawlerName)
            return
        }

        //======================================================================
        // 5. Start a job run.
        //======================================================================

        print("Starting the job...")

        // Construct the Amazon S3 URL for the job run's output. This is in
        // the bucket specified on the command line, with a folder name that's
        // unique for this job run.

        let timeStamp = Date().timeIntervalSince1970
        let jobPath = "\(jobName)-\(Int(timeStamp))"
        let outputURL = "s3://\(bucket)/\(jobPath)"

        // Start the job run.

        let jobRunID = await startJobRun(glueClient: glueClient, name: jobName,
                                         databaseName: databaseName,
                                         tableName: tableName,
                                         outputURL: outputURL)

        guard let jobRunID else {
            print("*** Job run ID is invalid.")
            _ = await deleteJob(glueClient: glueClient, name: jobName)
            _ = await deleteDatabase(glueClient: glueClient, name: databaseName,
withTables: true)
            _ = await deleteCrawler(glueClient: glueClient, name: crawlerName)
            return
        }

        //======================================================================
        // 6. Wait for the job run to indicate that the run is complete.
        //======================================================================

        print("Waiting for job run to end...")

        var jobRunFinished = false
        var jobRunState: GlueClientTypes.JobRunState
```

```
        repeat {
            let jobRun = await getJobRun(glueClient: glueClient, name: jobName, id:
jobRunID)
            guard let jobRun else {
                print("*** Unable to get the job run.")
                _ = await deleteJob(glueClient: glueClient, name: jobName)
                _ = await deleteDatabase(glueClient: glueClient, name: databaseName,
withTables: true)
                _ = await deleteCrawler(glueClient: glueClient, name: crawlerName)
                return
            }
            jobRunState = jobRun.jobRunState ?? .failed

            //=====================================================================
            // 7. Output where to find the data if the job run was successful.
            //    If the job run failed for any reason, output an appropriate
            //    error message.
            //=====================================================================

            switch jobRunState {
                case .succeeded:
                    print("Job run succeeded. JSON files are in the Amazon S3
path:")
                    print("    \(outputURL)")
                    jobRunFinished = true
                case .stopped:
                    jobRunFinished = true
                case .error:
                    print("*** Error: Job run ended in an error.
\(jobRun.errorMessage ?? "")")
                    jobRunFinished = true
                case .failed:
                    print("*** Error: Job run failed. \(jobRun.errorMessage ?? "")")
                    jobRunFinished = true
                case .timeout:
                    print("*** Warning: Job run timed out.")
                    jobRunFinished = true
                default:
                    do {
                        try await Task.sleep(for: .milliseconds(250))
                    } catch {
                        print("*** Error pausing the task.")
                    }
```

```
            }
        } while jobRunFinished != true

        //=====================================================================
        // 7.5. List the job runs for this job, showing each job run's ID and
        // its execution time.
        //=====================================================================

        print("Getting all job runs for the job \(jobName):")
        let jobRuns = await getJobRuns(glueClient: glueClient, name: jobName)

        if jobRuns.count == 0 {
            print("    <no job runs found>")
        } else {
            print("Found \(jobRuns.count) job runs... listing execution times:")
            for jobRun in jobRuns {
                print("    \(jobRun.id ?? "<unnamed>"): \(jobRun.executionTime)
seconds")
            }
        }

        //=====================================================================
        // 8. List the jobs for the user's account.
        //=====================================================================

        print("\nThe account has the following jobs:")
        let jobs = await listJobs(glueClient: glueClient)

        if jobs.count == 0 {
            print("    <no jobs found>")
        } else {
            for job in jobs {
                print("    \(job)")
            }
        }

        //=====================================================================
        // 9. Get the job run details for a job run.
        //=====================================================================

        print("Information about the job run:")
        let jobRun = await getJobRun(glueClient: glueClient, name: jobName, id:
jobRunID)
```

```
        guard let jobRun else {
            print("*** Unable to retrieve the job run.")
            _ = await deleteJob(glueClient: glueClient, name: jobName)
            _ = await deleteDatabase(glueClient: glueClient, name: databaseName,
 withTables: true)
            _ = await deleteCrawler(glueClient: glueClient, name: crawlerName)
            return
        }

        let startDate = jobRun.startedOn ?? Date(timeIntervalSince1970: 0)
        let endDate = jobRun.completedOn ?? Date(timeIntervalSince1970: 0)
        let dateFormatter: DateFormatter = DateFormatter()
        dateFormatter.dateStyle = .long
        dateFormatter.timeStyle = .long

        print("    Started at: \(dateFormatter.string(from: startDate))")
        print("  Completed at: \(dateFormatter.string(from: endDate))")

        //=====================================================================
        // 10. Delete the job.
        //=====================================================================

        print("\nDeleting the job...")
        _ = await deleteJob(glueClient: glueClient, name: jobName)

        //=====================================================================
        // 11. Delete the database and tables created by this example.
        //=====================================================================

        print("Deleting the database...")
        _ = await deleteDatabase(glueClient: glueClient, name: databaseName,
 withTables: true)

        //=====================================================================
        // 12. Delete the crawler.
        //=====================================================================

        print("Deleting the crawler...")
        if !(await deleteCrawler(glueClient: glueClient, name: crawlerName)) {
            return
        }
    }
}
```

```
/// The program's asynchronous entry point.
@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.runAsync()
        } catch {
            ExampleCommand.exit(withError: error)
        }
    }
}
```

- For API details, see the following topics in *AWS SDK for Swift API reference*.
  - CreateCrawler
  - CreateJob
  - DeleteCrawler
  - DeleteDatabase
  - DeleteJob
  - DeleteTable
  - GetCrawler
  - GetDatabase
  - GetDatabases
  - GetJob
  - GetJobRun
  - GetJobRuns
  - GetTables
  - ListJobs
  - StartCrawler
  - StartJobRun

Actions

`CreateCrawler`

The following code example shows how to use `CreateCrawler`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue

    /// Create a new AWS Glue crawler.
    ///
    /// - Parameters:
    ///   - glueClient: An AWS Glue client to use for the crawler.
    ///   - crawlerName: A name for the new crawler.
    ///   - iamRole: The name of an Amazon IAM role for the crawler to use.
    ///   - s3Path: The path of an Amazon S3 folder to use as a target location.
    ///   - cronSchedule: A `cron` schedule indicating when to run the crawler.
    ///   - databaseName: The name of an AWS Glue database to operate on.
    ///
    /// - Returns: `true` if the crawler is created successfully, otherwise `false`.
    func createCrawler(glueClient: GlueClient, crawlerName: String, iamRole: String,
                       s3Path: String, cronSchedule: String, databaseName: String)
 async -> Bool {
        let s3Target = GlueClientTypes.S3Target(path: s3url)
        let targetList = GlueClientTypes.CrawlerTargets(s3Targets: [s3Target])

        do {
            _ = try await glueClient.createCrawler(
                input: CreateCrawlerInput(
                    databaseName: databaseName,
                    description: "Created by the AWS SDK for Swift Scenario Example
 for AWS Glue.",
                    name: crawlerName,
                    role: iamRole,
```

```
                    schedule: cronSchedule,
                    tablePrefix: tablePrefix,
                    targets: targetList
                )
            )
        } catch _ as AlreadyExistsException {
            print("*** A crawler named \"\(crawlerName)\" already exists.")
            return false
        } catch _ as OperationTimeoutException {
            print("*** The attempt to create the AWS Glue crawler timed out.")
            return false
        } catch {
            print("*** An unexpected error occurred creating the AWS Glue crawler:
 \(error.localizedDescription)")
            return false
        }

        return true
    }
```

- For API details, see CreateCrawler in *AWS SDK for Swift API reference*.

`CreateJob`

The following code example shows how to use `CreateJob`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue

    /// Create a new AWS Glue job.
    ///
    /// - Parameters:
```

```
    ///   - glueClient: The AWS Glue client to use.
    ///   - jobName: The name to give the new job.
    ///   - role: The IAM role for the job to use when accessing AWS services.
    ///   - scriptLocation: The AWS S3 URI of the script to be run by the job.
    ///
    /// - Returns: `true` if the job is created successfully, otherwise `false`.
    func createJob(glueClient: GlueClient, name jobName: String, role: String,
                   scriptLocation: String) async -> Bool {
        let command = GlueClientTypes.JobCommand(
            name: "glueetl",
            pythonVersion: "3",
            scriptLocation: scriptLocation
        )

        do {
            _ = try await glueClient.createJob(
                input: CreateJobInput(
                    command: command,
                    description: "Created by the AWS SDK for Swift Glue basic
 scenario example.",
                    glueVersion: "3.0",
                    name: jobName,
                    numberOfWorkers: 10,
                    role: role,
                    workerType: .g1x
                )
            )
        } catch {
            return false
        }
        return true
    }
```

- For API details, see CreateJob in *AWS SDK for Swift API reference*.

DeleteCrawler

The following code example shows how to use `DeleteCrawler`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue

    /// Delete an AWS Glue crawler.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - name: The name of the crawler to delete.
    ///
    /// - Returns: `true` if successful, otherwise `false`.
    func deleteCrawler(glueClient: GlueClient, name: String) async -> Bool {
        do {
            _ = try await glueClient.deleteCrawler(
                input: DeleteCrawlerInput(name: name)
            )
        } catch {
            return false
        }
        return true
    }
```

- For API details, see DeleteCrawler in *AWS SDK for Swift API reference*.

`DeleteDatabase`

The following code example shows how to use `DeleteDatabase`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue

    /// Delete the specified database.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - databaseName: The name of the database to delete.
    ///   - deleteTables: A Bool indicating whether or not to delete the
    ///     tables in the database before attempting to delete the database.
    ///
    /// - Returns: `true` if the database (and optionally its tables) are
    ///   deleted, otherwise `false`.
    func deleteDatabase(glueClient: GlueClient, name databaseName: String,
                        withTables deleteTables: Bool = false) async -> Bool {
        if deleteTables {
            var tableNames: [String] = []

            // Get a list of the names of all of the tables in the database.

            let tableList = await self.getTablesInDatabase(glueClient: glueClient,
 databaseName: databaseName)
            for table in tableList {
                guard let name = table.name else {
                    continue
                }
                tableNames.append(name)
            }

            // Delete the tables. If there's only one table, use
            // `deleteTable()`, otherwise, use `batchDeleteTable()`. You can
            // use `batchDeleteTable()` for a single table, but this
            // demonstrates the use of `deleteTable()`.
```

```
            if tableNames.count == 1 {
                do {
                    print("    Deleting table...")
                    _ = try await glueClient.deleteTable(
                        input: DeleteTableInput(
                            databaseName: databaseName,
                            name: tableNames[0]
                        )
                    )
                } catch {
                    print("*** Unable to delete the table.")
                }
            } else {
                do {
                    print("    Deleting tables...")
                    _ = try await glueClient.batchDeleteTable(
                        input: BatchDeleteTableInput(
                            databaseName: databaseName,
                            tablesToDelete: tableNames
                        )
                    )
                } catch {
                    print("*** Unable to delete the tables.")
                }
            }
        }

        // Delete the database itself.

        do {
            print("    Deleting the database itself...")
            _ = try await glueClient.deleteDatabase(
                input: DeleteDatabaseInput(name: databaseName)
            )
        } catch {
            print("*** Unable to delete the database.")
            return false
        }
        return true
    }
```

- For API details, see DeleteDatabase in *AWS SDK for Swift API reference*.

`DeleteJob`

The following code example shows how to use `DeleteJob`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue

    /// Delete an AWS Glue job.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - jobName: The name of the job to delete.
    ///
    /// - Returns: `true` if the job is successfully deleted, otherwise `false`.
    func deleteJob(glueClient: GlueClient, name jobName: String) async -> Bool {
        do {
            _ = try await glueClient.deleteJob(
                input: DeleteJobInput(jobName: jobName)
            )
        } catch {
            return false
        }
        return true
    }
```

- For API details, see DeleteJob in *AWS SDK for Swift API reference*.

`DeleteTable`

The following code example shows how to use `DeleteTable`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue

                do {
                    print("    Deleting table...")
                    _ = try await glueClient.deleteTable(
                        input: DeleteTableInput(
                            databaseName: databaseName,
                            name: tableNames[0]
                        )
                    )
                } catch {
                    print("*** Unable to delete the table.")
                }
```

- For API details, see DeleteTable in *AWS SDK for Swift API reference*.

`GetCrawler`

The following code example shows how to use `GetCrawler`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
```

```
import AWSGlue

    /// Get the state of the specified AWS Glue crawler.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - name: The name of the crawler whose state should be returned.
    ///
    /// - Returns: A `GlueClientTypes.CrawlerState` value describing the
    ///   state of the crawler.
    func getCrawlerState(glueClient: GlueClient, name: String) async ->
 GlueClientTypes.CrawlerState {
        do {
            let output = try await glueClient.getCrawler(
                input: GetCrawlerInput(name: name)
            )

            // If the crawler or its state is `nil`, report that the crawler
            // is stopping. This may not be what you want for your
            // application but it works for this one!

            guard let crawler = output.crawler else {
                return GlueClientTypes.CrawlerState.stopping
            }
            guard let state = crawler.state else {
                return GlueClientTypes.CrawlerState.stopping
            }
            return state
        } catch {
            return GlueClientTypes.CrawlerState.stopping
        }
    }
```

- For API details, see GetCrawler in *AWS SDK for Swift API reference*.

GetDatabase

The following code example shows how to use `GetDatabase`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue

    /// Get the AWS Glue database with the specified name.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - name: The name of the database to return.
    ///
    /// - Returns: The `GlueClientTypes.Database` object describing the
    ///   specified database, or `nil` if an error occurs or the database
    ///   isn't found.
    func getDatabase(glueClient: GlueClient, name: String) async ->
 GlueClientTypes.Database? {
        do {
            let output = try await glueClient.getDatabase(
                input: GetDatabaseInput(name: name)
            )

            return output.database
        } catch {
            return nil
        }
    }
```

- For API details, see GetDatabase in *AWS SDK for Swift API reference*.

`GetJobRun`

The following code example shows how to use `GetJobRun`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue

    /// Get information about a specific AWS Glue job run.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - jobName: The name of the job to return job run data for.
    ///   - id: The run ID of the specific job run to return.
    ///
    /// - Returns: A `GlueClientTypes.JobRun` object describing the state of
    ///   the job run, or `nil` if an error occurs.
    func getJobRun(glueClient: GlueClient, name jobName: String, id: String) async -
> GlueClientTypes.JobRun? {
        do {
            let output = try await glueClient.getJobRun(
                input: GetJobRunInput(
                    jobName: jobName,
                    runId: id
                )
            )

            return output.jobRun
        } catch {
            return nil
        }
    }
```

- For API details, see GetJobRun in *AWS SDK for Swift API reference*.

GetJobRuns

The following code example shows how to use GetJobRuns.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue

    /// Return a list of the job runs for the specified job.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - jobName: The name of the job for which to return its job runs.
    ///   - maxResults: The maximum number of job runs to return (default:
    ///     1000).
    ///
    /// - Returns: An array of `GlueClientTypes.JobRun` objects describing
    ///   each job run.
    func getJobRuns(glueClient: GlueClient, name jobName: String, maxResults: Int? =
 nil) async -> [GlueClientTypes.JobRun] {
        do {
            let output = try await glueClient.getJobRuns(
                input: GetJobRunsInput(
                    jobName: jobName,
                    maxResults: maxResults
                )
            )

            guard let jobRuns = output.jobRuns else {
                print("*** No job runs found.")
                return []
            }

            return jobRuns
        } catch is EntityNotFoundException {
            print("*** The specified job name, \(jobName), doesn't exist.")
```

```
            return []
        } catch {
            print("*** Unexpected error getting job runs:")
            dump(error)
            return []
        }
    }
```

- For API details, see GetJobRuns in *AWS SDK for Swift API reference*.

`GetTables`

The following code example shows how to use `GetTables`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue

    /// Returns a list of the tables in the specified database.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - databaseName: The name of the database whose tables are to be
    ///     returned.
    ///
    /// - Returns: An array of `GlueClientTypes.Table` objects, each
    ///   describing one table in the named database. An empty array indicates
    ///   that there are either no tables in the database, or an error
    ///   occurred before any tables could be found.
    func getTablesInDatabase(glueClient: GlueClient, databaseName: String) async ->
 [GlueClientTypes.Table] {
        var tables: [GlueClientTypes.Table] = []
        var nextToken: String?
```

```
        repeat {
            do {
                let output = try await glueClient.getTables(
                    input: GetTablesInput(
                        databaseName: databaseName,
                        nextToken: nextToken
                    )
                )

                guard let tableList = output.tableList else {
                    return tables
                }

                tables = tables + tableList
                nextToken = output.nextToken
            } catch {
                return tables
            }
        } while nextToken != nil

        return tables
    }
```

- For API details, see GetTables in *AWS SDK for Swift API reference*.

`ListJobs`

The following code example shows how to use `ListJobs`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue
```

```
/// Return a list of the AWS Glue jobs listed on the user's account.
///
/// - Parameters:
///   - glueClient: The AWS Glue client to use.
///   - maxJobs: The maximum number of jobs to return (default: 100).
///
/// - Returns: An array of strings listing the names of all available AWS
///   Glue jobs.
func listJobs(glueClient: GlueClient, maxJobs: Int = 100) async -> [String] {
    var jobList: [String] = []
    var nextToken: String?

    repeat {
        do {
            let output = try await glueClient.listJobs(
                input: ListJobsInput(
                    maxResults: maxJobs,
                    nextToken: nextToken
                )
            )

            guard let jobs = output.jobNames else {
                return jobList
            }

            jobList = jobList + jobs
            nextToken = output.nextToken
        } catch {
            return jobList
        }
    } while (nextToken != nil)

    return jobList
}
```

- For API details, see ListJobs in *AWS SDK for Swift API reference*.

`StartCrawler`

The following code example shows how to use `StartCrawler`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue

    /// Start running an AWS Glue crawler.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use when starting the crawler.
    ///   - name: The name of the crawler to start running.
    ///
    /// - Returns: `true` if the crawler is started successfully, otherwise `false`.
    func startCrawler(glueClient: GlueClient, name: String) async -> Bool {
        do {
            _ = try await glueClient.startCrawler(
                input: StartCrawlerInput(name: name)
            )
        } catch {
            print("*** An unexpected error occurred starting the crawler.")
            return false
        }

        return true
    }
```

- For API details, see StartCrawler in *AWS SDK for Swift API reference*.

`StartJobRun`

The following code example shows how to use `StartJobRun`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSGlue

    /// Start an AWS Glue job run.
    ///
    /// - Parameters:
    ///   - glueClient: The AWS Glue client to use.
    ///   - jobName: The name of the job to run.
    ///   - databaseName: The name of the AWS Glue database to run the job against.
    ///   - tableName: The name of the table in the database to run the job against.
    ///   - outputURL: The AWS S3 URI of the bucket location into which to
    ///     write the resulting output.
    ///
    /// - Returns: `true` if the job run is started successfully, otherwise `false`.
    func startJobRun(glueClient: GlueClient, name jobName: String, databaseName:
 String,
                     tableName: String, outputURL: String) async -> String? {
        do {
            let output = try await glueClient.startJobRun(
                input: StartJobRunInput(
                    arguments: [
                        "--input_database": databaseName,
                        "--input_table": tableName,
                        "--output_bucket_url": outputURL
                    ],
                    jobName: jobName,
                    numberOfWorkers: 10,
                    workerType: .g1x
                )
            )

            guard let id = output.jobRunId else {
                return nil
            }
```

```
            return id
        } catch {
            return nil
        }
    }
```

- For API details, see StartJobRun in *AWS SDK for Swift API reference*.

IAM examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with IAM.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Actions

Actions

`AttachRolePolicy`

The following code example shows how to use `AttachRolePolicy`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
```

```
import AWSS3


    public func attachRolePolicy(role: String, policyArn: String) async throws {
        let input = AttachRolePolicyInput(
            policyArn: policyArn,
            roleName: role
        )
        do {
            _ = try await client.attachRolePolicy(input: input)
        } catch {
            print("ERROR: Attaching a role policy:", dump(error))
            throw error
        }
    }
```

- For API details, see AttachRolePolicy in *AWS SDK for Swift API reference*.

`CreateAccessKey`

The following code example shows how to use `CreateAccessKey`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func createAccessKey(userName: String) async throws ->
 IAMClientTypes.AccessKey {
        let input = CreateAccessKeyInput(
            userName: userName
        )
        do {
```

```
            let output = try await iamClient.createAccessKey(input: input)
            guard let accessKey = output.accessKey else {
                throw ServiceHandlerError.keyError
            }
            return accessKey
        } catch {
            print("ERROR: createAccessKey:", dump(error))
            throw error
        }
    }
```

- For API details, see CreateAccessKey in *AWS SDK for Swift API reference*.

`CreatePolicy`

The following code example shows how to use `CreatePolicy`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func createPolicy(name: String, policyDocument: String) async throws ->
 IAMClientTypes.Policy {
        let input = CreatePolicyInput(
            policyDocument: policyDocument,
            policyName: name
        )
        do {
            let output = try await iamClient.createPolicy(input: input)
            guard let policy = output.policy else {
                throw ServiceHandlerError.noSuchPolicy
```

```
            }
            return policy
        } catch {
            print("ERROR: createPolicy:", dump(error))
            throw error
        }
    }
```

- For API details, see CreatePolicy in *AWS SDK for Swift API reference*.

`CreateRole`

The following code example shows how to use `CreateRole`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func createRole(name: String, policyDocument: String) async throws ->
 String {
        let input = CreateRoleInput(
            assumeRolePolicyDocument: policyDocument,
            roleName: name
        )
        do {
            let output = try await client.createRole(input: input)
            guard let role = output.role else {
                throw ServiceHandlerError.noSuchRole
            }
            guard let id = role.roleId else {
                throw ServiceHandlerError.noSuchRole
```

```
            }
            return id
        } catch {
            print("ERROR: createRole:", dump(error))
            throw error
        }
    }
```

- For API details, see CreateRole in *AWS SDK for Swift API reference*.

`CreateServiceLinkedRole`

The following code example shows how to use `CreateServiceLinkedRole`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func createServiceLinkedRole(service: String, suffix: String? = nil,
 description: String?)
                    async throws -> IAMClientTypes.Role {
        let input = CreateServiceLinkedRoleInput(
            awsServiceName: service,
            customSuffix: suffix,
            description: description
        )
        do {
            let output = try await client.createServiceLinkedRole(input: input)
            guard let role = output.role else {
                throw ServiceHandlerError.noSuchRole
            }
            return role
        } catch {
```

```
            print("ERROR: createServiceLinkedRole:", dump(error))
            throw error
        }
    }
```

- For API details, see CreateServiceLinkedRole in *AWS SDK for Swift API reference*.

`CreateUser`

The following code example shows how to use `CreateUser`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func createUser(name: String) async throws -> String {
        let input = CreateUserInput(
            userName: name
        )
        do {
            let output = try await client.createUser(input: input)
            guard let user = output.user else {
                throw ServiceHandlerError.noSuchUser
            }
            guard let id = user.userId else {
                throw ServiceHandlerError.noSuchUser
            }
            return id
        } catch {
            print("ERROR: createUser:", dump(error))
            throw error
        }
    }
```

- For API details, see CreateUser in *AWS SDK for Swift API reference*.

`DeleteAccessKey`

The following code example shows how to use `DeleteAccessKey`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func deleteAccessKey(user: IAMClientTypes.User? = nil,
                                key: IAMClientTypes.AccessKey) async throws
    {
        let userName: String?

        if user != nil {
            userName = user!.userName
        } else {
            userName = nil
        }

        let input = DeleteAccessKeyInput(
            accessKeyId: key.accessKeyId,
            userName: userName
        )
        do {
            _ = try await iamClient.deleteAccessKey(input: input)
        } catch {
            print("ERROR: deleteAccessKey:", dump(error))
            throw error
        }
    }
```

- For API details, see DeleteAccessKey in *AWS SDK for Swift API reference*.

`DeletePolicy`

The following code example shows how to use `DeletePolicy`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func deletePolicy(policy: IAMClientTypes.Policy) async throws {
        let input = DeletePolicyInput(
            policyArn: policy.arn
        )
        do {
            _ = try await iamClient.deletePolicy(input: input)
        } catch {
            print("ERROR: deletePolicy:", dump(error))
            throw error
        }
    }
```

- For API details, see DeletePolicy in *AWS SDK for Swift API reference*.

`DeleteRole`

The following code example shows how to use `DeleteRole`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func deleteRole(role: IAMClientTypes.Role) async throws {
        let input = DeleteRoleInput(
            roleName: role.roleName
        )
        do {
            _ = try await iamClient.deleteRole(input: input)
        } catch {
            print("ERROR: deleteRole:", dump(error))
            throw error
        }
    }
```

- For API details, see DeleteRole in *AWS SDK for Swift API reference*.

`DeleteUser`

The following code example shows how to use `DeleteUser`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func deleteUser(user: IAMClientTypes.User) async throws {
        let input = DeleteUserInput(
            userName: user.userName
        )
        do {
            _ = try await iamClient.deleteUser(input: input)
        } catch {
            print("ERROR: deleteUser:", dump(error))
            throw error
        }
    }
```

- For API details, see DeleteUser in *AWS SDK for Swift API reference*.

`DeleteUserPolicy`

The following code example shows how to use `DeleteUserPolicy`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    func deleteUserPolicy(user: IAMClientTypes.User, policyName: String) async
 throws {
        let input = DeleteUserPolicyInput(
            policyName: policyName,
            userName: user.userName
```

```
        )
        do {
            _ = try await iamClient.deleteUserPolicy(input: input)
        } catch {
            print("ERROR: deleteUserPolicy:", dump(error))
            throw error
        }
    }
```

- For API details, see DeleteUserPolicy in *AWS SDK for Swift API reference*.

`DetachRolePolicy`

The following code example shows how to use `DetachRolePolicy`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func detachRolePolicy(policy: IAMClientTypes.Policy, role:
 IAMClientTypes.Role) async throws {
        let input = DetachRolePolicyInput(
            policyArn: policy.arn,
            roleName: role.roleName
        )

        do {
            _ = try await iamClient.detachRolePolicy(input: input)
        } catch {
            print("ERROR: detachRolePolicy:", dump(error))
            throw error
```

```
        }
    }
```

- For API details, see DetachRolePolicy in *AWS SDK for Swift API reference*.

`GetPolicy`

The following code example shows how to use `GetPolicy`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func getPolicy(arn: String) async throws -> IAMClientTypes.Policy {
        let input = GetPolicyInput(
            policyArn: arn
        )
        do {
            let output = try await client.getPolicy(input: input)
            guard let policy = output.policy else {
                throw ServiceHandlerError.noSuchPolicy
            }
            return policy
        } catch {
            print("ERROR: getPolicy:", dump(error))
            throw error
        }
    }
```

- For API details, see GetPolicy in *AWS SDK for Swift API reference*.

GetRole

The following code example shows how to use `GetRole`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func getRole(name: String) async throws -> IAMClientTypes.Role {
        let input = GetRoleInput(
            roleName: name
        )
        do {
            let output = try await client.getRole(input: input)
            guard let role = output.role else {
                throw ServiceHandlerError.noSuchRole
            }
            return role
        } catch {
            print("ERROR: getRole:", dump(error))
            throw error
        }
    }
```

- For API details, see GetRole in *AWS SDK for Swift API reference*.

GetUser

The following code example shows how to use `GetUser`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func getUser(name: String? = nil) async throws -> IAMClientTypes.User {
        let input = GetUserInput(
            userName: name
        )
        do {
            let output = try await iamClient.getUser(input: input)
            guard let user = output.user else {
                throw ServiceHandlerError.noSuchUser
            }
            return user
        } catch {
            print("ERROR: getUser:", dump(error))
            throw error
        }
    }
```

- For API details, see GetUser in *AWS SDK for Swift API reference*.

`ListAttachedRolePolicies`

The following code example shows how to use `ListAttachedRolePolicies`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3



    /// Returns a list of AWS Identity and Access Management (IAM) policies
    /// that are attached to the role.
    ///
    /// - Parameter role: The IAM role to return the policy list for.
    ///
    /// - Returns: An array of `IAMClientTypes.AttachedPolicy` objects
    ///   describing each managed policy that's attached to the role.
    public func listAttachedRolePolicies(role: String) async throws ->
 [IAMClientTypes.AttachedPolicy] {
        var policyList: [IAMClientTypes.AttachedPolicy] = []

        // Use "Paginated" to get all the attached role polices.
        // This lets the SDK handle the 'isTruncated' in
 "ListAttachedRolePoliciesOutput".
        let input = ListAttachedRolePoliciesInput(
            roleName: role
        )
        let output = client.listAttachedRolePoliciesPaginated(input: input)

        do {
            for try await page in output {
                guard let attachedPolicies = page.attachedPolicies else {
                    print("Error: no attached policies returned.")
                    continue
                }
                for attachedPolicy in attachedPolicies {
                    policyList.append(attachedPolicy)
                }
            }
```

```
        } catch {
            print("ERROR: listAttachedRolePolicies:", dump(error))
            throw error
        }

        return policyList
    }
```

- For API details, see ListAttachedRolePolicies in *AWS SDK for Swift API reference*.

`ListGroups`

The following code example shows how to use `ListGroups`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func listGroups() async throws -> [String] {
        var groupList: [String] = []

        // Use "Paginated" to get all the groups.
        // This lets the SDK handle the 'isTruncated' property in
 "ListGroupsOutput".
        let input = ListGroupsInput()

        let pages = client.listGroupsPaginated(input: input)
        do {
            for try await page in pages {
                guard let groups = page.groups else {
                    print("Error: no groups returned.")
                    continue
                }
```

```
                for group in groups {
                    if let name = group.groupName {
                        groupList.append(name)
                    }
                }
            }
        } catch {
            print("ERROR: listGroups:", dump(error))
            throw error
        }
        return groupList
    }
```

- For API details, see ListGroups in *AWS SDK for Swift API reference*.

`ListPolicies`

The following code example shows how to use `ListPolicies`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func listPolicies() async throws -> [MyPolicyRecord] {
        var policyList: [MyPolicyRecord] = []

        // Use "Paginated" to get all the policies.
        // This lets the SDK handle the 'isTruncated' in "ListPoliciesOutput".
        let input = ListPoliciesInput()
        let output = client.listPoliciesPaginated(input: input)

        do {
```

```
            for try await page in output {
                guard let policies = page.policies else {
                    print("Error: no policies returned.")
                    continue
                }

                for policy in policies {
                    guard let name = policy.policyName,
                          let id = policy.policyId,
                          let arn = policy.arn
                    else {
                        throw ServiceHandlerError.noSuchPolicy
                    }
                    policyList.append(MyPolicyRecord(name: name, id: id, arn: arn))
                }
            }
        } catch {
            print("ERROR: listPolicies:", dump(error))
            throw error
        }

        return policyList
    }
```

- For API details, see ListPolicies in *AWS SDK for Swift API reference*.

`ListRolePolicies`

The following code example shows how to use `ListRolePolicies`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3
```

```
    public func listRolePolicies(role: String) async throws -> [String] {
        var policyList: [String] = []

        // Use "Paginated" to get all the role policies.
        // This lets the SDK handle the 'isTruncated' in "ListRolePoliciesOutput".
        let input = ListRolePoliciesInput(
            roleName: role
        )
        let pages = client.listRolePoliciesPaginated(input: input)

        do {
            for try await page in pages {
                guard let policies = page.policyNames else {
                    print("Error: no role policies returned.")
                    continue
                }

                for policy in policies {
                    policyList.append(policy)
                }
            }
        } catch {
            print("ERROR: listRolePolicies:", dump(error))
            throw error
        }
        return policyList
    }
```

- For API details, see ListRolePolicies in *AWS SDK for Swift API reference*.

`ListRoles`

The following code example shows how to use `ListRoles`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func listRoles() async throws -> [String] {
        var roleList: [String] = []

        // Use "Paginated" to get all the roles.
        // This lets the SDK handle the 'isTruncated' in "ListRolesOutput".
        let input = ListRolesInput()
        let pages = client.listRolesPaginated(input: input)

        do {
            for try await page in pages {
                guard let roles = page.roles else {
                    print("Error: no roles returned.")
                    continue
                }

                for role in roles {
                    if let name = role.roleName {
                        roleList.append(name)
                    }
                }
            }
        } catch {
            print("ERROR: listRoles:", dump(error))
            throw error
        }
        return roleList
    }
```

- For API details, see ListRoles in *AWS SDK for Swift API reference*.

`ListUsers`

The following code example shows how to use `ListUsers`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    public func listUsers() async throws -> [MyUserRecord] {
        var userList: [MyUserRecord] = []

        // Use "Paginated" to get all the users.
        // This lets the SDK handle the 'isTruncated' in "ListUsersOutput".
        let input = ListUsersInput()
        let output = client.listUsersPaginated(input: input)

        do {
            for try await page in output {
                guard let users = page.users else {
                    continue
                }
                for user in users {
                    if let id = user.userId, let name = user.userName {
                        userList.append(MyUserRecord(id: id, name: name))
                    }
                }
            }
        }
        catch {
            print("ERROR: listUsers:", dump(error))
            throw error
        }
       return userList
    }
```

- For API details, see ListUsers in *AWS SDK for Swift API reference*.

`PutUserPolicy`

The following code example shows how to use `PutUserPolicy`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSIAM
import AWSS3


    func putUserPolicy(policyDocument: String, policyName: String, user:
 IAMClientTypes.User) async throws {
        let input = PutUserPolicyInput(
            policyDocument: policyDocument,
            policyName: policyName,
            userName: user.userName
        )
        do {
            _ = try await iamClient.putUserPolicy(input: input)
        } catch {
            print("ERROR: putUserPolicy:", dump(error))
            throw error
        }
    }
```

- For API details, see PutUserPolicy in *AWS SDK for Swift API reference*.

Lambda examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with Lambda.

Basics are code examples that show you how to perform the essential operations within a service.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Basics
- Actions

Basics

Learn the basics

The following code example shows how to:

- Create an IAM role and Lambda function, then upload handler code.
- Invoke the function with a single parameter and get results.
- Update the function code and configure with an environment variable.
- Invoke the function with new parameters and get results. Display the returned execution log.
- List the functions for your account, then clean up resources.

For more information, see Create a Lambda function with the console.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

Define the first Lambda function, which simply increments the specified value.

```
// swift-tools-version: 5.9
// Copyright Amazon.com, Inc. or its affiliates. All Rights Reserved.
// SPDX-License-Identifier: Apache-2.0
//
```

```
// The swift-tools-version declares the minimum version of Swift required to
// build this package.

import PackageDescription

let package = Package(
    name: "increment",
    // Let Xcode know the minimum Apple platforms supported.
    platforms: [
        .macOS(.v13)
    ],
    dependencies: [
        // Dependencies declare other packages that this package depends on.
        .package(
            url: "https://github.com/swift-server/swift-aws-lambda-runtime.git",
            branch: "main"),
    ],
    targets: [
        // Targets are the basic building blocks of a package, defining a module or
 a test suite.
        // Targets can depend on other targets in this package and products
        // from dependencies.
        .executableTarget(
            name: "increment",
            dependencies: [
                .product(name: "AWSLambdaRuntime", package: "swift-aws-lambda-
runtime"),
            ],
            path: "Sources"
        )
    ]
)

import Foundation
import AWSLambdaRuntime

/// Represents the contents of the requests being received from the client.
/// This structure must be `Decodable` to indicate that its initializer
/// converts an external representation into this type.
struct Request: Decodable, Sendable {
    /// The action to perform.
    let action: String
    /// The number to act upon.
    let number: Int
```

```
}

/// The contents of the response sent back to the client. This must be
/// `Encodable`.
struct Response: Encodable, Sendable {
    /// The resulting value after performing the action.
    let answer: Int?
}


/// The Lambda function body.
///
/// - Parameters:
///   - event: The `Request` describing the request made by the
///     client.
///   - context: A `LambdaContext` describing the context in
///     which the lambda function is running.
///
/// - Returns: A `Response` object that will be encoded to JSON and sent
///   to the client by the Lambda runtime.
let incrementLambdaRuntime = LambdaRuntime {
        (event: Request, context: LambdaContext) -> Response in
    let action = event.action
    var answer: Int?

    if action != "increment" {
        context.logger.error("Unrecognized operation: \"\(action)\". The only
 supported action is \"increment\".")
    } else {
        answer = event.number + 1
        context.logger.info("The calculated answer is \(answer!).")
    }

    let response = Response(answer: answer)
    return response
}

// Run the Lambda runtime code.

try await incrementLambdaRuntime.run()
```

Define the second Lambda function, which performs an arithmetic operation on two numbers.

```
// swift-tools-version: 5.9
// Copyright Amazon.com, Inc. or its affiliates. All Rights Reserved.
// SPDX-License-Identifier: Apache-2.0
//
// The swift-tools-version declares the minimum version of Swift required to
// build this package.

import PackageDescription

let package = Package(
    name: "calculator",
    // Let Xcode know the minimum Apple platforms supported.
    platforms: [
        .macOS(.v13)
    ],
    dependencies: [
        // Dependencies declare other packages that this package depends on.
        .package(
            url: "https://github.com/swift-server/swift-aws-lambda-runtime.git",
            branch: "main"),
    ],
    targets: [
        // Targets are the basic building blocks of a package, defining a module or
 a test suite.
        // Targets can depend on other targets in this package and products
        // from dependencies.
        .executableTarget(
            name: "calculator",
            dependencies: [
                .product(name: "AWSLambdaRuntime", package: "swift-aws-lambda-
runtime"),
            ],
            path: "Sources"
        )
    ]
)

import Foundation
import AWSLambdaRuntime

/// Represents the contents of the requests being received from the client.
/// This structure must be `Decodable` to indicate that its initializer
/// converts an external representation into this type.
```

```
struct Request: Decodable, Sendable {
    /// The action to perform.
    let action: String
    /// The first number to act upon.
    let x: Int
    /// The second number to act upon.
    let y: Int
}

/// A dictionary mapping operation names to closures that perform that
/// operation and return the result.
let actions = [
    "plus": { (x: Int, y: Int) -> Int in
        return x + y
    },
    "minus": { (x: Int, y: Int) -> Int in
        return x - y
    },
    "times": { (x: Int, y: Int) -> Int in
        return x * y
    },
    "divided-by": { (x: Int, y: Int) -> Int in
        return x / y
    }
]

/// The contents of the response sent back to the client. This must be
/// `Encodable`.
struct Response: Encodable, Sendable {
    /// The resulting value after performing the action.
    let answer: Int?
}


/// The Lambda function's entry point. Called by the Lambda runtime.
///
/// - Parameters:
///   - event: The `Request` describing the request made by the
///     client.
///   - context: A `LambdaContext` describing the context in
///     which the lambda function is running.
///
/// - Returns: A `Response` object that will be encoded to JSON and sent
///   to the client by the Lambda runtime.
```

```
let calculatorLambdaRuntime = LambdaRuntime {
        (_ event: Request, context: LambdaContext) -> Response in
    let action = event.action
    var answer: Int?
    var actionFunc: ((Int, Int) -> Int)?

    // Get the closure to run to perform the calculation.

    actionFunc = await actions[action]

    guard let actionFunc else {
        context.logger.error("Unrecognized operation '\(action)\'")
        return Response(answer: nil)
    }

    // Perform the calculation and return the answer.

    answer = actionFunc(event.x, event.y)

    guard let answer else {
        context.logger.error("Error computing \(event.x) \(action) \(event.y)")
    }
    context.logger.info("\(event.x) \(action) \(event.y) = \(answer)")

    return Response(answer: answer)
}

try await calculatorLambdaRuntime.run()
```

Define the main program that will invoke the two Lambda functions.

```
// swift-tools-version: 5.9
// Copyright Amazon.com, Inc. or its affiliates. All Rights Reserved.
// SPDX-License-Identifier: Apache-2.0
//
// The swift-tools-version declares the minimum version of Swift required to
// build this package.

import PackageDescription

let package = Package(
    name: "lambda-basics",
```

```
    // Let Xcode know the minimum Apple platforms supported.
    platforms: [
        .macOS(.v13)
    ],
    dependencies: [
        // Dependencies declare other packages that this package depends on.
        .package(
            url: "https://github.com/awslabs/aws-sdk-swift",
            from: "1.0.0"),
        .package(
            url: "https://github.com/apple/swift-argument-parser.git",
            branch: "main"
        )
    ],
    targets: [
        // Targets are the basic building blocks of a package, defining a module or
 a test suite.
        // Targets can depend on other targets in this package and products
        // from dependencies.
        .executableTarget(
            name: "lambda-basics",
            dependencies: [
                .product(name: "AWSLambda", package: "aws-sdk-swift"),
                .product(name: "AWSIAM", package: "aws-sdk-swift"),
                .product(name: "ArgumentParser", package: "swift-argument-parser")
            ],
            path: "Sources"
        )
    ]
)

//
/// An example demonstrating a variety of important AWS Lambda functions.

import ArgumentParser
import AWSIAM
import SmithyWaitersAPI
import AWSClientRuntime
import AWSLambda
import Foundation

/// Represents the contents of the requests being received from the client.
/// This structure must be `Decodable` to indicate that its initializer
/// converts an external representation into this type.
```

```
struct IncrementRequest: Encodable, Decodable, Sendable {
    /// The action to perform.
    let action: String
    /// The number to act upon.
    let number: Int
}

struct Response: Encodable, Decodable, Sendable {
    /// The resulting value after performing the action.
    let answer: Int?
}

struct CalculatorRequest: Encodable, Decodable, Sendable {
    /// The action to perform.
    let action: String
    /// The first number to act upon.
    let x: Int
    /// The second number to act upon.
    let y: Int
}

let exampleName = "SwiftLambdaRoleExample"
let basicsFunctionName = "lambda-basics-function"

/// The ARN of the standard IAM policy for execution of Lambda functions.
let policyARN = "arn:aws:iam::aws:policy/service-role/AWSLambdaBasicExecutionRole"

struct ExampleCommand: ParsableCommand {
    // -MARK: Command arguments
    @Option(help: "Name of the IAM Role to use for the Lambda functions")
    var role = exampleName
    @Option(help: "Zip archive containing the 'increment' lambda function")
    var incpath: String
    @Option(help: "Zip archive containing the 'calculator' lambda function")
    var calcpath: String
    @Option(help: "Name of the Amazon S3 Region to use (default: us-east-1)")
    var region = "us-east-1"

    static var configuration = CommandConfiguration(
        commandName: "lambda-basics",
        abstract: """
        This example demonstrates several common operations using AWS Lambda.
        """,
        discussion: """
```

```
        """
    )

    /// Returns the specified IAM role object.
    ///
    /// - Parameters:
    ///   - iamClient: `IAMClient` to use when looking for the role.
    ///   - roleName: The name of the role to check.
    ///
    /// - Returns: The `IAMClientTypes.Role` representing the specified role.
    func getRole(iamClient: IAMClient, roleName: String) async throws
                 -> IAMClientTypes.Role {
        do {
            let roleOutput = try await iamClient.getRole(
                input: GetRoleInput(
                    roleName: roleName
                )
            )

            guard let role = roleOutput.role else {
                throw ExampleError.roleNotFound
            }
            return role
        } catch {
            throw ExampleError.roleNotFound
        }
    }

    /// Create the AWS IAM role that will be used to access AWS Lambda.
    ///
    /// - Parameters:
    ///   - iamClient: The AWS `IAMClient` to use.
    ///   - roleName: The name of the AWS IAM role to use for Lambda.
    ///
    /// - Throws: `ExampleError.roleCreateError`
    ///
    /// - Returns: The `IAMClientTypes.Role` struct that describes the new role.
    func createRoleForLambda(iamClient: IAMClient, roleName: String) async throws ->
IAMClientTypes.Role {
        let output = try await iamClient.createRole(
            input: CreateRoleInput(
                assumeRolePolicyDocument:
                """
                {
```

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

        guard let role = output.role else {
            throw ExampleError.roleCreateError
        }

        // Wait for the role to be ready for use.

        _ = try await iamClient.waitUntilRoleExists(
            options: WaiterOptions(
                maxWaitTime: 20,
                minDelay: 0.5,
                maxDelay: 2
            ),
            input: GetRoleInput(roleName: roleName)
        )

        return role
    }

    /// Detect whether or not the AWS Lambda function with the specified name
    /// exists, by requesting its function information.
    ///
    /// - Parameters:
    ///   - lambdaClient: The `LambdaClient` to use.
    ///   - name: The name of the AWS Lambda function to find.
    ///
    /// - Returns: `true` if the Lambda function exists. Otherwise `false`.
    func doesLambdaFunctionExist(lambdaClient: LambdaClient, name: String) async ->
Bool {
        do {
            _ = try await lambdaClient.getFunction(
```

```
                input: GetFunctionInput(functionName: name)
            )
        } catch {
            return false
        }

        return true
    }

    /// Create the specified AWS Lambda function.
    ///
    /// - Parameters:
    ///   - lambdaClient: The `LambdaClient` to use.
    ///   - functionName: The name of the AWS Lambda function to create.
    ///   - roleArn: The ARN of the role to apply to the function.
    ///   - path: The path of the Zip archive containing the function.
    ///
    /// - Returns: `true` if the AWS Lambda was successfully created; `false`
    ///   if it wasn't.
    func createFunction(lambdaClient: LambdaClient, functionName: String,
                                roleArn: String?, path: String) async throws -> Bool
{
        do {
            // Read the Zip archive containing the AWS Lambda function.

            let zipUrl = URL(fileURLWithPath: path)
            let zipData = try Data(contentsOf: zipUrl)

            // Create the AWS Lambda function that runs the specified code,
            // using the name given on the command line. The Lambda function
            // will run using the Amazon Linux 2 runtime.

            _ = try await lambdaClient.createFunction(
                input: CreateFunctionInput(
                    code: LambdaClientTypes.FunctionCode(zipFile: zipData),
                    functionName: functionName,
                    handler: "handle",
                    role: roleArn,
                    runtime: .providedal2
                )
            )
        } catch {
            print("*** Error creating Lambda function:")
            dump(error)
```

```
            return false
        }

        // Wait for a while to be sure the function is done being created.

        let output = try await lambdaClient.waitUntilFunctionActiveV2(
            options: WaiterOptions(
                maxWaitTime: 20,
                minDelay: 0.5,
                maxDelay: 2
            ),
            input: GetFunctionInput(functionName: functionName)
        )

        switch output.result {
            case .success:
                return true
            case .failure:
                return false
        }
    }

    /// Update the AWS Lambda function with new code to run when the function
    /// is invoked.
    ///
    /// - Parameters:
    ///   - lambdaClient: The `LambdaClient` to use.
    ///   - functionName: The name of the AWS Lambda function to update.
    ///   - path: The pathname of the Zip file containing the packaged Lambda
    ///     function.
    /// - Throws: `ExampleError.zipFileReadError`
    /// - Returns: `true` if the function's code is updated successfully.
    ///   Otherwise, returns `false`.
    func updateFunctionCode(lambdaClient: LambdaClient, functionName: String,
                            path: String) async throws -> Bool {
        let zipUrl = URL(fileURLWithPath: path)
        let zipData: Data

        // Read the function's Zip file.

        do {
            zipData = try Data(contentsOf: zipUrl)
        } catch {
            throw ExampleError.zipFileReadError
```

```
        }

        // Update the function's code and wait for the updated version to be
        // ready for use.

        do {
            _ = try await lambdaClient.updateFunctionCode(
                input: UpdateFunctionCodeInput(
                    functionName: functionName,
                    zipFile: zipData
                )
            )
        } catch {
            return false
        }

        let output = try await lambdaClient.waitUntilFunctionUpdatedV2(
            options: WaiterOptions(
                maxWaitTime: 20,
                minDelay: 0.5,
                maxDelay: 2
            ),
            input: GetFunctionInput(
                functionName: functionName
            )
        )

        switch output.result {
            case .success:
                return true
            case .failure:
                return false
        }
    }

    /// Tell the server-side component to log debug output by setting its
    /// environment's `LOG_LEVEL` to `DEBUG`.
    ///
    /// - Parameters:
    ///   - lambdaClient: The `LambdaClient` to use.
    ///   - functionName: The name of the AWS Lambda function to enable debug
    ///     logging for.
    ///
    /// - Throws: `ExampleError.environmentResponseMissingError`,
```

```
    ///   `ExampleError.updateFunctionConfigurationError`,
    ///   `ExampleError.environmentVariablesMissingError`,
    ///   `ExampleError.logLevelIncorrectError`,
    ///   `ExampleError.updateFunctionConfigurationError`
    func enableDebugLogging(lambdaClient: LambdaClient, functionName: String) async
throws {
        let envVariables = [
            "LOG_LEVEL": "DEBUG"
        ]
        let environment = LambdaClientTypes.Environment(variables: envVariables)

        do {
            let output = try await lambdaClient.updateFunctionConfiguration(
                input: UpdateFunctionConfigurationInput(
                    environment: environment,
                    functionName: functionName
                )
            )

            guard let response = output.environment else {
                throw ExampleError.environmentResponseMissingError
            }

            if response.error != nil {
                throw ExampleError.updateFunctionConfigurationError
            }

            guard let retVariables = response.variables else {
                throw ExampleError.environmentVariablesMissingError
            }

            for envVar in retVariables {
                if envVar.key == "LOG_LEVEL" && envVar.value != "DEBUG" {
                    print("*** Log level is not set to DEBUG!")
                    throw ExampleError.logLevelIncorrectError
                }
            }
        } catch {
            throw ExampleError.updateFunctionConfigurationError
        }
    }

    /// Returns an array containing the names of all AWS Lambda functions
    /// available to the user.
```

```
    ///
    /// - Parameter lambdaClient: The `IAMClient` to use.
    ///
    /// - Throws: `ExampleError.listFunctionsError`
    ///
    /// - Returns: An array of lambda function name strings.
    func getFunctionNames(lambdaClient: LambdaClient) async throws -> [String] {
        let pages = lambdaClient.listFunctionsPaginated(
            input: ListFunctionsInput()
        )

        var functionNames: [String] = []

        for try await page in pages {
            guard let functions = page.functions else {
                throw ExampleError.listFunctionsError
            }

            for function in functions {
                functionNames.append(function.functionName ?? "<unknown>")
            }
        }

        return functionNames
    }

    /// Invoke the Lambda function to increment a value.
    ///
    /// - Parameters:
    ///   - lambdaClient: The `IAMClient` to use.
    ///   - number: The number to increment.
    ///
    /// - Throws: `ExampleError.noAnswerReceived`, `ExampleError.invokeError`
    ///
    /// - Returns: An integer number containing the incremented value.
    func invokeIncrement(lambdaClient: LambdaClient, number: Int) async throws ->
 Int {
        do {
            let incRequest = IncrementRequest(action: "increment", number: number)
            let incData = try! JSONEncoder().encode(incRequest)

            // Invoke the lambda function.

            let invokeOutput = try await lambdaClient.invoke(
```

```
                input: InvokeInput(
                    functionName: "lambda-basics-function",
                    payload: incData
                )
            )

            let response = try! JSONDecoder().decode(Response.self,
from:invokeOutput.payload!)

            guard let answer = response.answer else {
                throw ExampleError.noAnswerReceived
            }
            return answer

        } catch {
            throw ExampleError.invokeError
        }
    }

    /// Invoke the calculator Lambda function.
    ///
    /// - Parameters:
    ///   - lambdaClient: The `IAMClient` to use.
    ///   - action: Which arithmetic operation to perform: "plus", "minus",
    ///     "times", or "divided-by".
    ///   - x: The first number to use in the computation.
    ///   - y: The second number to use in the computation.
    ///
    /// - Throws: `ExampleError.noAnswerReceived`, `ExampleError.invokeError`
    ///
    /// - Returns: The computed answer as an `Int`.
    func invokeCalculator(lambdaClient: LambdaClient, action: String, x: Int, y:
Int) async throws -> Int {
        do {
            let calcRequest = CalculatorRequest(action: action, x: x, y: y)
            let calcData = try! JSONEncoder().encode(calcRequest)

            // Invoke the lambda function.

            let invokeOutput = try await lambdaClient.invoke(
                input: InvokeInput(
                    functionName: "lambda-basics-function",
                    payload: calcData
                )
```

```
            )

            let response = try! JSONDecoder().decode(Response.self,
from:invokeOutput.payload!)

            guard let answer = response.answer else {
                throw ExampleError.noAnswerReceived
            }
            return answer

        } catch {
            throw ExampleError.invokeError
        }

    }

    /// Perform the example's tasks.
    func basics() async throws {
        let iamClient = try await IAMClient(
            config: IAMClient.IAMClientConfig(region: region)
        )

        let lambdaClient = try await LambdaClient(
            config: LambdaClient.LambdaClientConfig(region: region)
        )

        /// The IAM role to use for the example.
        var iamRole: IAMClientTypes.Role

        // Look for the specified role. If it already exists, use it. If not,
        // create it and attach the desired policy to it.

        do {
            iamRole = try await getRole(iamClient: iamClient, roleName: role)
        } catch ExampleError.roleNotFound {
            // The role wasn't found, so create it and attach the needed
            // policy.

            iamRole = try await createRoleForLambda(iamClient: iamClient, roleName:
role)

            do {
                _ = try await iamClient.attachRolePolicy(
```

```
                    input: AttachRolePolicyInput(policyArn: policyARN, roleName:
role)
                )
            } catch {
                throw ExampleError.policyError
            }
        }

        // Give the policy time to attach to the role.

        sleep(5)

        // Look to see if the function already exists. If it does, throw an
        // error.

        if await doesLambdaFunctionExist(lambdaClient: lambdaClient, name:
basicsFunctionName) {
            throw ExampleError.functionAlreadyExists
        }

        // Create, then invoke, the "increment" version of the calculator
        // function.

        print("Creating the increment Lambda function...")
        if try await createFunction(lambdaClient: lambdaClient, functionName:
basicsFunctionName,
                                    roleArn: iamRole.arn, path: incpath) {
            print("Running increment function calls...")
            for number in 0...4 {
                do {
                    let answer = try await invokeIncrement(lambdaClient:
lambdaClient, number: number)
                    print("Increment \(number) = \(answer)")
                } catch {
                    print("Error incrementing \(number): ",
error.localizedDescription)
                }
            }
        } else {
            print("*** Failed to create the increment function.")
        }

        // Enable debug logging.
```

```
        print("\nEnabling debug logging...")
        try await enableDebugLogging(lambdaClient: lambdaClient, functionName:
basicsFunctionName)

        // Change it to a basic arithmetic calculator. Then invoke it a few
        // times.

        print("\nReplacing the Lambda function with a calculator...")

        if try await updateFunctionCode(lambdaClient: lambdaClient, functionName:
basicsFunctionName,
                                        path: calcpath) {
            print("Running calculator function calls...")
            for x in [6, 10] {
                for y in [2, 4] {
                    for action in ["plus", "minus", "times", "divided-by"] {
                        do {
                            let answer = try await invokeCalculator(lambdaClient:
lambdaClient, action: action, x: x, y: y)
                            print("\(x) \(action) \(y) = \(answer)")
                        } catch {
                            print("Error calculating \(x) \(action) \(y): ",
error.localizedDescription)
                        }
                    }
                }
            }
        }

        // List all lambda functions.

        let functionNames = try await getFunctionNames(lambdaClient: lambdaClient)

        if functionNames.count > 0 {
            print("\nAWS Lambda functions available on your account:")
            for name in functionNames {
                print("  \(name)")
            }
        }

        // Delete the lambda function.

        print("Deleting lambda function...")
```

```
        do {
            _ = try await lambdaClient.deleteFunction(
                input: DeleteFunctionInput(
                    functionName: "lambda-basics-function"
                )
            )
        } catch {
            print("Error: Unable to delete the function.")
        }

        // Detach the role from the policy, then delete the role.

        print("Deleting the AWS IAM role...")

        do {
            _ = try await iamClient.detachRolePolicy(
                input: DetachRolePolicyInput(
                    policyArn: policyARN,
                    roleName: role
                )
            )
            _ = try await iamClient.deleteRole(
                input: DeleteRoleInput(
                    roleName: role
                )
            )
        } catch {
            throw ExampleError.deleteRoleError
        }
    }
}

// -MARK: - Entry point

/// The program's asynchronous entry point.
@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.basics()
        } catch {
```

```
            ExampleCommand.exit(withError: error)
        }
    }
}


/// Errors thrown by the example's functions.
enum ExampleError: Error {
    /// An AWS Lambda function with the specified name already exists.
    case functionAlreadyExists
    /// The specified role doesn't exist.
    case roleNotFound
    /// Unable to create the role.
    case roleCreateError
    /// Unable to delete the role.
    case deleteRoleError
    /// Unable to attach a policy to the role.
    case policyError
    /// Unable to get the executable directory.
    case executableNotFound
    /// An error occurred creating a lambda function.
    case createLambdaError
    /// An error occurred invoking the lambda function.
    case invokeError
    /// No answer received from the invocation.
    case noAnswerReceived
    /// Unable to list the AWS Lambda functions.
    case listFunctionsError
    /// Unable to update the AWS Lambda function.
    case updateFunctionError
    /// Unable to update the function configuration.
    case updateFunctionConfigurationError
    /// The environment response is missing after an
    /// UpdateEnvironmentConfiguration attempt.
    case environmentResponseMissingError
    /// The environment variables are missing from the EnvironmentResponse and
    /// no errors occurred.
    case environmentVariablesMissingError
    /// The log level is incorrect after attempting to set it.
    case logLevelIncorrectError
    /// Unable to load the AWS Lambda function's Zip file.
    case zipFileReadError

    var errorDescription: String? {
```

```
        switch self {
        case .functionAlreadyExists:
            return "An AWS Lambda function with that name already exists."
        case .roleNotFound:
            return "The specified role doesn't exist."
        case .deleteRoleError:
            return "Unable to delete the AWS IAM role."
        case .roleCreateError:
            return "Unable to create the specified role."
        case .policyError:
            return "An error occurred attaching the policy to the role."
        case .executableNotFound:
            return "Unable to find the executable program directory."
        case .createLambdaError:
            return "An error occurred creating a lambda function."
        case .invokeError:
            return "An error occurred invoking a lambda function."
        case .noAnswerReceived:
            return "No answer received from the lambda function."
        case .listFunctionsError:
            return "Unable to list the AWS Lambda functions."
        case .updateFunctionError:
            return "Unable to update the AWS lambda function."
        case .updateFunctionConfigurationError:
            return "Unable to update the AWS lambda function configuration."
        case .environmentResponseMissingError:
            return "The environment is missing from the response after updating the
 function configuration."
        case .environmentVariablesMissingError:
            return "While no error occurred, no environment variables were returned
 following function configuration."
        case .logLevelIncorrectError:
            return "The log level is incorrect after attempting to set it to DEBUG."
        case .zipFileReadError:
            return "Unable to read the AWS Lambda function."
        }
    }
}
```

- For API details, see the following topics in *AWS SDK for Swift API reference*.
  - CreateFunction
  - DeleteFunction

- GetFunction
- Invoke
- ListFunctions
- UpdateFunctionCode
- UpdateFunctionConfiguration

Actions

`CreateFunction`

The following code example shows how to use `CreateFunction`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSLambda
import Foundation

        do {
            // Read the Zip archive containing the AWS Lambda function.

            let zipUrl = URL(fileURLWithPath: path)
            let zipData = try Data(contentsOf: zipUrl)

            // Create the AWS Lambda function that runs the specified code,
            // using the name given on the command line. The Lambda function
            // will run using the Amazon Linux 2 runtime.

            _ = try await lambdaClient.createFunction(
                input: CreateFunctionInput(
                    code: LambdaClientTypes.FunctionCode(zipFile: zipData),
                    functionName: functionName,
                    handler: "handle",
```

```
                    role: roleArn,
                    runtime: .providedal2
                )
            )
        } catch {
            print("*** Error creating Lambda function:")
            dump(error)
            return false
        }
```

- For API details, see CreateFunction in *AWS SDK for Swift API reference*.

`DeleteFunction`

The following code example shows how to use `DeleteFunction`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSLambda
import Foundation

        do {
            _ = try await lambdaClient.deleteFunction(
                input: DeleteFunctionInput(
                    functionName: "lambda-basics-function"
                )
            )
        } catch {
            print("Error: Unable to delete the function.")
        }
```

- For API details, see DeleteFunction in *AWS SDK for Swift API reference*.

`GetFunction`

The following code example shows how to use `GetFunction`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSLambda
import Foundation

    /// Detect whether or not the AWS Lambda function with the specified name
    /// exists, by requesting its function information.
    ///
    /// - Parameters:
    ///   - lambdaClient: The `LambdaClient` to use.
    ///   - name: The name of the AWS Lambda function to find.
    ///
    /// - Returns: `true` if the Lambda function exists. Otherwise `false`.
    func doesLambdaFunctionExist(lambdaClient: LambdaClient, name: String) async ->
 Bool {
        do {
            _ = try await lambdaClient.getFunction(
                input: GetFunctionInput(functionName: name)
            )
        } catch {
            return false
        }

        return true
    }
```

- For API details, see GetFunction in *AWS SDK for Swift API reference*.

Invoke

The following code example shows how to use `Invoke`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSLambda
import Foundation

    /// Invoke the Lambda function to increment a value.
    ///
    /// - Parameters:
    ///   - lambdaClient: The `IAMClient` to use.
    ///   - number: The number to increment.
    ///
    /// - Throws: `ExampleError.noAnswerReceived`, `ExampleError.invokeError`
    ///
    /// - Returns: An integer number containing the incremented value.
    func invokeIncrement(lambdaClient: LambdaClient, number: Int) async throws ->
 Int {
        do {
            let incRequest = IncrementRequest(action: "increment", number: number)
            let incData = try! JSONEncoder().encode(incRequest)

            // Invoke the lambda function.

            let invokeOutput = try await lambdaClient.invoke(
                input: InvokeInput(
                    functionName: "lambda-basics-function",
                    payload: incData
                )
            )

            let response = try! JSONDecoder().decode(Response.self,
 from:invokeOutput.payload!)
```

```
            guard let answer = response.answer else {
                throw ExampleError.noAnswerReceived
            }
            return answer

        } catch {
            throw ExampleError.invokeError
        }
    }
```

- For API details, see Invoke in *AWS SDK for Swift API reference*.

`ListFunctions`

The following code example shows how to use `ListFunctions`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSLambda
import Foundation

    /// Returns an array containing the names of all AWS Lambda functions
    /// available to the user.
    ///
    /// - Parameter lambdaClient: The `IAMClient` to use.
    ///
    /// - Throws: `ExampleError.listFunctionsError`
    ///
    /// - Returns: An array of lambda function name strings.
    func getFunctionNames(lambdaClient: LambdaClient) async throws -> [String] {
        let pages = lambdaClient.listFunctionsPaginated(
            input: ListFunctionsInput()
        )
```

```
        var functionNames: [String] = []

        for try await page in pages {
            guard let functions = page.functions else {
                throw ExampleError.listFunctionsError
            }

            for function in functions {
                functionNames.append(function.functionName ?? "<unknown>")
            }
        }

        return functionNames
    }
```

- For API details, see ListFunctions in *AWS SDK for Swift API reference*.

`UpdateFunctionCode`

The following code example shows how to use `UpdateFunctionCode`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSLambda
import Foundation

        let zipUrl = URL(fileURLWithPath: path)
        let zipData: Data

        // Read the function's Zip file.

        do {
            zipData = try Data(contentsOf: zipUrl)
        } catch {
```

```
            throw ExampleError.zipFileReadError
        }

        // Update the function's code and wait for the updated version to be
        // ready for use.

        do {
            _ = try await lambdaClient.updateFunctionCode(
                input: UpdateFunctionCodeInput(
                    functionName: functionName,
                    zipFile: zipData
                )
            )
        } catch {
            return false
        }
```

- For API details, see UpdateFunctionCode in *AWS SDK for Swift API reference*.

`UpdateFunctionConfiguration`

The following code example shows how to use `UpdateFunctionConfiguration`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSClientRuntime
import AWSLambda
import Foundation

    /// Tell the server-side component to log debug output by setting its
    /// environment's `LOG_LEVEL` to `DEBUG`.
    ///
    /// - Parameters:
    ///   - lambdaClient: The `LambdaClient` to use.
    ///   - functionName: The name of the AWS Lambda function to enable debug
```

```
    ///     logging for.
    ///
    /// - Throws: `ExampleError.environmentResponseMissingError`,
    ///   `ExampleError.updateFunctionConfigurationError`,
    ///   `ExampleError.environmentVariablesMissingError`,
    ///   `ExampleError.logLevelIncorrectError`,
    ///   `ExampleError.updateFunctionConfigurationError`
    func enableDebugLogging(lambdaClient: LambdaClient, functionName: String) async
throws {
        let envVariables = [
            "LOG_LEVEL": "DEBUG"
        ]
        let environment = LambdaClientTypes.Environment(variables: envVariables)

        do {
            let output = try await lambdaClient.updateFunctionConfiguration(
                input: UpdateFunctionConfigurationInput(
                    environment: environment,
                    functionName: functionName
                )
            )

            guard let response = output.environment else {
                throw ExampleError.environmentResponseMissingError
            }

            if response.error != nil {
                throw ExampleError.updateFunctionConfigurationError
            }

            guard let retVariables = response.variables else {
                throw ExampleError.environmentVariablesMissingError
            }

            for envVar in retVariables {
                if envVar.key == "LOG_LEVEL" && envVar.value != "DEBUG" {
                    print("*** Log level is not set to DEBUG!")
                    throw ExampleError.logLevelIncorrectError
                }
            }
        } catch {
            throw ExampleError.updateFunctionConfigurationError
        }
    }
```

- For API details, see UpdateFunctionConfiguration in *AWS SDK for Swift API reference*.

Amazon RDS examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with Amazon RDS.

Basics are code examples that show you how to perform the essential operations within a service.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Basics
- Actions

Basics

Learn the basics

The following code example shows how to:

- Create a custom DB parameter group and set parameter values.
- Create a DB instance that's configured to use the parameter group. The DB instance also contains a database.
- Take a snapshot of the instance.
- Delete the instance and parameter group.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

The `Package.swift` file.

```
// swift-tools-version: 5.9
//
// The swift-tools-version declares the minimum version of Swift required to
// build this package.

import PackageDescription

let package = Package(
    name: "rds-scenario",
    // Let Xcode know the minimum Apple platforms supported.
    platforms: [
        .macOS(.v13),
        .iOS(.v15)
    ],
    dependencies: [
        // Dependencies declare other packages that this package depends on.
        .package(
            url: "https://github.com/awslabs/aws-sdk-swift",
            from: "1.4.0"),
        .package(
            url: "https://github.com/apple/swift-argument-parser.git",
            branch: "main"
        )
    ],
    targets: [
        // Targets are the basic building blocks of a package, defining a module or
 a test suite.
        // Targets can depend on other targets in this package and products
        // from dependencies.
        .executableTarget(
            name: "rds-scenario",
            dependencies: [
                .product(name: "AWSRDS", package: "aws-sdk-swift"),
```

```
                .product(name: "ArgumentParser", package: "swift-argument-parser")
            ],
            path: "Sources")

    ]
)
```

The Swift code file, `entry.swift`.

```
// An example that shows how to use the AWS SDK for Swift to perform a variety
// of operations using Amazon Relational Database Service (RDS).
//

import ArgumentParser
import Foundation
import AWSRDS

struct ExampleCommand: ParsableCommand {
    @Option(help: "The AWS Region to run AWS API calls in.")
    var awsRegion = "us-east-1"
    @Option(help: "The username to use for the database administrator.")
    var dbUsername = "admin"
    @Option(help: "The password to use for the database administrator.")
    var dbPassword: String

    static var configuration = CommandConfiguration(
        commandName: "rds-scenario",
        abstract: """
        Performs various operations to demonstrate the use of Amazon RDS Instances
        using the AWS SDK for Swift.
        """,
        discussion: """
        """
    )

    /// Called by ``main()`` to run the bulk of the example.
    func runAsync() async throws {
        let example = try await Example(region: awsRegion, username: dbUsername,
 password: dbPassword)

        await example.run()
    }
```

```
}

class Example {
    let rdsClient: RDSClient

    // Storage for AWS RDS properties

    let dbUsername: String
    let dbPassword: String
    var dbInstanceIdentifier: String
    var dbSnapshotIdentifier: String
    var dbParameterGroupName: String
    var dbParameterGroup: RDSClientTypes.DBParameterGroup?
    var selectedEngineVersion: String?

    init(region: String, username: String, password: String) async throws{
        let rdsConfig = try await RDSClient.RDSClientConfig(region: region)
        rdsClient = RDSClient(config: rdsConfig)

        dbUsername = username
        dbPassword = password
        dbParameterGroupName = ""
        dbInstanceIdentifier = ""
        dbSnapshotIdentifier = ""
    }

    /// The example's main body.
    func run() async {
        var parameterGroupFamilies: Set<String> = []

        //======================================================================
        // 1. Get available database engine families for MySQL.
        //======================================================================

        let engineVersions = await getDBEngineVersions(engineName: "mysql")

        for version in engineVersions {
            if version.dbParameterGroupFamily != nil {
                parameterGroupFamilies.insert(version.dbParameterGroupFamily!)
            }
        }

        if engineVersions.count > 0 {
            selectedEngineVersion = engineVersions.last!.engineVersion
```

```
        } else {
            print("*** Unable to find a valid database engine version. Canceling
operations.")
            await cleanUp()
            return
        }

        print("Found \(parameterGroupFamilies.count) parameter group families:")
        for family in parameterGroupFamilies {
            print("    \(family)")
        }

        //=====================================================================
        // 2. Select an engine family and create a custom DB parameter group.
        //    We select a family by sorting the set of family names, then
        //    choosing the last one.
        //=====================================================================

        let sortedFamilies = parameterGroupFamilies.sorted()

        guard let selectedFamily = sortedFamilies.last else {
            print("*** Unable to find a database engine family. Canceling
operations.")
            await cleanUp()
            return
        }

        print("Selected database engine family \(selectedFamily)")

        dbParameterGroupName = tempName(prefix: "rds-example")
        print("Creating a database parameter group named \(dbParameterGroupName)
using \(selectedFamily)")
        dbParameterGroup = await createDBParameterGroup(groupName:
dbParameterGroupName,
                                                        familyName: selectedFamily)

        //=====================================================================
        // 3. Get the parameter group's details.
        //=====================================================================

        print("Getting the database parameter group list...")
        let dbParameterGroupList = await describeDBParameterGroups(groupName:
dbParameterGroupName)
        guard let dbParameterGroupList else {
```

```
            await cleanUp()
            return
        }

        print("Found \(dbParameterGroupList.count) parameter groups...")
        for group in dbParameterGroupList {
            print("    \(group.dbParameterGroupName ?? "<unknown>")")
        }
        print()

        //======================================================================
        // 4. Get a list of the parameter group's parameters. This list is
        //    likely to be long, so use pagination. Find the
        //    auto_increment_offset and auto_increment_increment parameters.
        //======================================================================

        let parameters = await describeDBParameters(groupName: dbParameterGroupName)

        //======================================================================
        // 5. Parse and display each parameter's name, description, and
        //    allowed values.
        //======================================================================

        for parameter in parameters {
            let name = parameter.parameterName
            guard let name else {
                print("*** Unable to get parameter name!")
                continue
            }

            if name == "auto_increment_offset" || name ==
"auto_increment_increment" {
                print("Parameter \(name):")
                print("          Value: \(parameter.parameterValue ??
"<undefined>")")
                print("      Data type: \(parameter.dataType ?? "<unknown>")")
                print("    Description: \(parameter.description ?? "")")
                print(" Allowed values: \(parameter.allowedValues ??
"<unspecified")")
                print(String(repeating: "=", count: 78))
            }
        }

        //======================================================================
```

```
        // 6. Modify both the auto_increment_offset and
        //    auto_increment_increment parameters in one call in the custom
        //    parameter group. Set their parameterValue fields to a new
        //    permitted value.
        //======================================================================

        print("Setting auto_increment_offset and auto_increment_increment both to
5...")
        await modifyDBParameters(groupName: dbParameterGroupName)

        //======================================================================
        // 7. Get and display the updated parameters, specifying a source of
        //    "user" to get only the modified parameters.
        //======================================================================

        let updatedParameters = await describeDBParameters(groupName:
dbParameterGroupName, source: "user")

        for parameter in updatedParameters {
            let name = parameter.parameterName
            guard let name else {
                print("*** Unable to get parameter name!")
                continue
            }

            print("Parameter \(name):")
            print("          Value: \(parameter.parameterValue ?? "<undefined>")")
            print("      Data type: \(parameter.dataType ?? "<unknown>")")
            print("    Description: \(parameter.description ?? "")")
            print(" Allowed values: \(parameter.allowedValues ?? "<unspecified")")
            print(String(repeating: "=", count: 78))
        }

        //======================================================================
        // 8. Get a list of allowed engine versions using
        //    DescribeRDSEngineVersions.
        //======================================================================

        await listAllowedEngines(family: selectedFamily)

        //======================================================================
        // 9. Get a list of micro instance classes available for the selected
        //    engine and engine version.
        //======================================================================
```

```
        let dbInstanceClass = await chooseMicroInstance(engine: "mysql",
engineVersion: selectedEngineVersion)
        guard let dbInstanceClass else {
            print("Did not get a valid instance class. Canceling operations.")
            await cleanUp()
            return
        }

        //=====================================================================
        // 10. Create an RDS database that contains a MySQL database and uses
        //     the parameter group we created.
        //=====================================================================

        print("Creating the database instance...")

        guard let instanceClass = dbInstanceClass.dbInstanceClass else {
            print("Instance class name is unknown. Canceling operations.")
            await cleanUp()
            return
        }

        dbInstanceIdentifier = tempName(prefix: "sample-identifier")
        let dbInstanceArn = await createDBInstance(
            name: "SampleDatabase\(Int.random(in: 1000000..<1000000000))",
            instanceIdentifier: dbInstanceIdentifier,
            parameterGroupName: dbParameterGroupName,
            engine: "mysql",
            engineVersion: selectedEngineVersion!,
            instanceClass: instanceClass,
            username: dbUsername,
            password: dbPassword
        )

        if dbInstanceArn == nil {
            await cleanUp()
            return
        }

        //=====================================================================
        // 11. Wait for the database instance to be ready by calling
        //     DescribeDBInstances repeatedly until it reports
        //     dbInstanceStatus as "available". This can take upwards of 10
        //     minutes, let the user know that.
```

```
        //=====================================================================

        guard let endpoint = await waitUntilDBInstanceReady(instanceIdentifier:
dbInstanceIdentifier) else {
            print("\nDid not get a valid endpoint from AWS RDS.")
            await cleanUp()
            return
        }

        guard let endpointAddress = endpoint.address else {
            print("\nNo endpoint address returned.")
            await cleanUp()
            return
        }
        guard let endpointPort = endpoint.port else {
            print("\nNo endpoint port returned.")
            await cleanUp()
            return
        }

        //=====================================================================
        // 12. Display connection information for the database instance.
        //=====================================================================

        print("\nTo connect to the new database instance using 'mysql' from the
shell:")
        print("    mysql -h \(endpointAddress) -P \(endpointPort) -u
\(self.dbUsername)")

        //=====================================================================
        // 13. Create a snapshot of the database instance.
        //=====================================================================

        dbSnapshotIdentifier = tempName(prefix: "sample-snapshot")
        await createDBSnapshot(instanceIdentifier: dbInstanceIdentifier,
snapshotIdentifier: dbSnapshotIdentifier)

        //=====================================================================
        // 14. Wait for the snapshot to be ready.
        //=====================================================================

        await waitUntilDBSnapshotReady(instanceIdentifier: dbInstanceIdentifier,
snapshotIdentifier: dbSnapshotIdentifier)
```

```
        // That's it! Clean up and exit!

        print("Example complete! Cleaning up...")
        await cleanUp()
    }

    /// Clean up by discarding and closing down all allocated EC2 items.
    func cleanUp() async {
        print("Deleting the database instance \(dbInstanceIdentifier)...")
        await deleteDBInstance(instanceIdentifier: dbInstanceIdentifier)
        await waitUntilDBInstanceDeleted(instanceIdentifier: dbInstanceIdentifier)

        print("Deleting the database parameter group \(dbParameterGroupName)...")
        await deleteDBParameterGroup(groupName: dbParameterGroupName)
    }

    /// Get all the database engine versions available for the specified
    /// database engine.
    ///
    /// - Parameter engineName: The name of the database engine to query.
    ///
    /// - Returns: An array of `RDSClientTypes.DBEngineVersion` structures,
    ///   each describing one supported version of the specified database.
    func getDBEngineVersions(engineName: String) async ->
[RDSClientTypes.DBEngineVersion] {
        do {
            let output = try await rdsClient.describeDBEngineVersions(
                input: DescribeDBEngineVersionsInput(
                    engine: engineName
                )
            )

            return output.dbEngineVersions ?? []
        } catch {
            return []
        }
    }

    /// Create a new database parameter group with the specified name.
    ///
    /// - Parameters:
    ///   - groupName: The name of the new parameter group.
    ///   - familyName: The name of the parameter group family.
    /// - Returns:
```

```
    func createDBParameterGroup(groupName: String, familyName: String) async ->
RDSClientTypes.DBParameterGroup? {
        do {
            let output = try await rdsClient.createDBParameterGroup(
                input: CreateDBParameterGroupInput(
                    dbParameterGroupFamily: familyName,
                    dbParameterGroupName: groupName,
                    description: "Created using the AWS SDK for Swift"
                )
            )
            return output.dbParameterGroup
        } catch {
            print("*** Error creating the parameter group:
\(error.localizedDescription)")
            return nil
        }
    }

    /// Get descriptions of the database parameter groups matching the given
    /// name.
    ///
    /// - Parameter groupName: The name of the parameter group to describe.
    ///
    /// - Returns: An array of [RDSClientTypes.DBParameterGroup] objects
    ///   describing the parameter group.
    func describeDBParameterGroups(groupName: String) async ->
[RDSClientTypes.DBParameterGroup]? {
        do {
            let output = try await rdsClient.describeDBParameterGroups(
                input: DescribeDBParameterGroupsInput(
                    dbParameterGroupName: groupName
                )
            )
            return output.dbParameterGroups
        } catch {
            print("*** Error getting the database parameter group's details:
\(error.localizedDescription)")
            return nil
        }
    }

    /// Returns the detailed parameter list for the specified database
    /// parameter group.
    ///
```

```
    /// - Parameters:
    ///   - groupName: The name of the parameter group to return parameters for.
    ///   - source: The types of parameters to return (`user`, `system`, or
    ///     `engine-default`).
    ///
    /// - Returns: An array of `RdSClientTypes.Parameter` objects, each
    ///   describing one of the group's parameters.
    func describeDBParameters(groupName: String, source: String? = nil) async ->
[RDSClientTypes.Parameter] {
        var parameterList: [RDSClientTypes.Parameter] = []

        do {
            let pages = rdsClient.describeDBParametersPaginated(
                input: DescribeDBParametersInput(
                    dbParameterGroupName: groupName,
                    source: source
                )
            )

            for try await page in pages {
                guard let parameters = page.parameters else {
                    return []
                }

                parameterList += parameters
            }
        } catch {
            print("*** Error getting database parameters:
\(error.localizedDescription)")
            return []
        }

        return parameterList
    }

    /// Demonstrates modifying two of the specified database parameter group's
    /// parameters.
    ///
    /// - Parameter groupName: The name of the parameter group to change
    ///   parameters for.
    func modifyDBParameters(groupName: String) async {
        let parameter1 = RDSClientTypes.Parameter(
            applyMethod: RDSClientTypes.ApplyMethod.immediate,
            parameterName: "auto_increment_offset",
```

```
            parameterValue: "5"
        )
        let parameter2 = RDSClientTypes.Parameter(
            applyMethod: RDSClientTypes.ApplyMethod.immediate,
            parameterName: "auto_increment_increment",
            parameterValue: "5"
        )

        let parameterList = [parameter1, parameter2]

        do {
            _ = try await rdsClient.modifyDBParameterGroup(
                input: ModifyDBParameterGroupInput(
                    dbParameterGroupName: groupName,
                    parameters: parameterList
                )
            )

            print("Successfully modified the parameter group \(groupName).")
        } catch {
            print("*** Error modifying the parameter group \(groupName):
\(error.localizedDescription)")
        }
    }

    /// Output a list of the database engine versions supported by the
    /// specified family.
    ///
    /// - Parameter family: The family for which to list allowed database
    ///   engines.
    func listAllowedEngines(family: String?) async {
        do {
            let output = try await rdsClient.describeDBEngineVersions(
                input: DescribeDBEngineVersionsInput(
                    dbParameterGroupFamily: family,
                    engine: "mysql"
                )
            )

            guard let engineVersions = output.dbEngineVersions else {
                print("No engine versions returned.")
                return
            }
```

```
            print("Found \(engineVersions.count) database engine versions:")
            for version in engineVersions {
                print("    \(version.engineVersion ?? "<unknown>"):
\(version.dbEngineDescription ?? "")")
            }
        } catch {
            print("*** Error getting database engine version list:
\(error.localizedDescription)")
            return
        }
    }

    /// Print a list of available database instances with "micro" in the class
    /// name, then return one of them to be used by other code.
    ///
    /// - Parameters:
    ///   - engine: The database engine for which to list database instance
    ///     classes.
    ///   - engineVersion: The database version for which to list instances.
    ///
    /// - Returns: An `RDSClientTypes.OrderableDBInstanceOption` describing
    ///   the selected instance type.
    func chooseMicroInstance(engine: String = "mysql", engineVersion: String? = nil)
async -> RDSClientTypes.OrderableDBInstanceOption? {
        do {
            let pages = rdsClient.describeOrderableDBInstanceOptionsPaginated(
                input: DescribeOrderableDBInstanceOptionsInput(
                    engine: engine,
                    engineVersion: engineVersion
                )
            )

            var optionsList: [RDSClientTypes.OrderableDBInstanceOption] = []

            for try await page in pages {
                guard let orderableDBInstanceOptions =
page.orderableDBInstanceOptions else {
                    continue
                }

                for dbInstanceOption in orderableDBInstanceOptions {
                    guard let className = dbInstanceOption.dbInstanceClass else {
                        continue
                    }
```

```
                    if className.contains("micro") {
                        optionsList.append(dbInstanceOption)
                    }
                }
            }

            print("Found \(optionsList.count) database instances of 'micro' class
types:")
            for dbInstanceOption in optionsList {
                print("    \(dbInstanceOption.engine ??
"<unknown>") \(dbInstanceOption.engineVersion ?? "<unknown>")
(\(dbInstanceOption.dbInstanceClass ?? "<unknown class>"))")
            }

            return optionsList[0]
        } catch {
            print("*** Error getting a list of orderable instance options:
\(error.localizedDescription)")
            return nil
        }
    }

    /// Create a new database instance.
    ///
    /// - Parameters:
    ///   - name: The name of the database to create.
    ///   - instanceIdentifier: The identifier to give the new database
    ///     instance.
    ///   - parameterGroupName: The name of the parameter group to associate
    ///     with the new database instance.
    ///   - engine: The database engine to use.
    ///   - engineVersion: The version of the database given by `engine` to
    ///     use.
    ///   - instanceClass: The memory and compute capacity of the database
    ///     instance, such as `db.m5.large``.
    ///   - username: The admin user's username to establish for the new
    ///     instance.
    ///   - password: The password to use for the specified user's access.
    ///
    /// - Returns: A string indicating the ARN of the newly created database
    ///   instance, or nil if the instance couldn't be created.
    func createDBInstance(name: String, instanceIdentifier: String,
parameterGroupName: String,
```

```
                        engine: String, engineVersion: String, instanceClass:
String,
                        username: String, password: String) async -> String? {
       do {
           let output = try await rdsClient.createDBInstance(
               input: CreateDBInstanceInput(
                   allocatedStorage: 100,
                   dbInstanceClass: instanceClass,
                   dbInstanceIdentifier: instanceIdentifier,
                   dbName: name,
                   dbParameterGroupName: parameterGroupName,
                   engine: engine,
                   engineVersion: engineVersion,
                   masterUserPassword: password,
                   masterUsername: username,
                   storageType: "gp2"
               )
           )

           guard let dbInstance = output.dbInstance else {
               print("*** Unable to get the database instance.")
               return nil
           }

           return dbInstance.dbInstanceArn
       } catch {
           print("*** An error occurred while creating the database instance:
\(error.localizedDescription)")
           return nil
       }
   }

   /// Wait until the specified database is available to use.
   ///
   /// - Parameter instanceIdentifier: The database instance identifier of the
   ///   database to wait for.
   func waitUntilDBInstanceReady(instanceIdentifier: String) async ->
RDSClientTypes.Endpoint? {
       do {
           putString("Waiting for the database instance to be ready to use. This
may take 10 minutes or more...")
           while true {
               let output = try await rdsClient.describeDBInstances(
                   input: DescribeDBInstancesInput(
```

```
                        dbInstanceIdentifier: instanceIdentifier
                    )
                )

                guard let instanceList = output.dbInstances else {
                    continue
                }

                for instance in instanceList {
                    let status = instance.dbInstanceStatus

                    guard let status else {
                        print("\nUnable to determine the status.")
                        continue
                    }

                    if status.contains("available") {
                        return instance.endpoint
                    } else {
                        putString(".")
                        do {
                            try await Task.sleep(for: .seconds(15))
                        } catch {
                            print("*** Error pausing the task!")
                        }
                    }
                }
            }
        } catch {
            print("*** Unable to wait until the database is ready:
\(error.localizedDescription)")
            return nil
        }
    }

    /// Create a snapshot of the specified name.
    ///
    /// - Parameters:
    ///   - instanceIdentifier: The identifier of the database instance to
    ///     snapshot.
    ///   - snapshotIdentifier: A unique identifier to give the newly-created
    ///     snapshot.
    func createDBSnapshot(instanceIdentifier: String, snapshotIdentifier: String)
async {
```

```
        do {
            let output = try await rdsClient.createDBSnapshot(
                input: CreateDBSnapshotInput(
                    dbInstanceIdentifier: instanceIdentifier,
                    dbSnapshotIdentifier: snapshotIdentifier
                )
            )

            guard let snapshot = output.dbSnapshot else {
                print("No snapshot returned.")
                return
            }

            print("The snapshot has been created with ID \(snapshot.dbiResourceId ??
"<unknown>")")
        } catch {
            print("*** Unable to create the database snapshot named
\(snapshotIdentifier): \(error.localizedDescription)")
        }
    }

    /// Wait until the specified database snapshot is available to use.
    ///
    /// - Parameters:
    ///   - instanceIdentifier: The identifier of the database for which the
    ///     snapshot was taken.
    ///   - snapshotIdentifier: The identifier of the snapshot to wait for.
    func waitUntilDBSnapshotReady(instanceIdentifier: String, snapshotIdentifier:
String) async {
        var snapshotReady = false

        putString("Waiting for the snapshot to be ready...")

        do {
            while !snapshotReady {
                let output = try await rdsClient.describeDBSnapshots(
                    input: DescribeDBSnapshotsInput(
                        dbInstanceIdentifier: instanceIdentifier,
                        dbSnapshotIdentifier: snapshotIdentifier
                    )
                )

                guard let snapshotList = output.dbSnapshots else {
                    return
```

```
                }

                for snapshot in snapshotList {
                    guard let snapshotReadyStr = snapshot.status else {
                        return
                    }

                    if snapshotReadyStr.contains("available") {
                        snapshotReady = true
                        print()
                    } else {
                        putString(".")
                        do {
                            try await Task.sleep(for: .seconds(15))
                        } catch {
                            print("\n*** Error pausing the task!")
                        }
                    }
                }
            }
        } catch {
            print("\n*** Unable to wait for the database snapshot to be ready:
\(error.localizedDescription)")
        }
    }

    /// Delete the specified database instance.
    ///
    /// - Parameter instanceIdentifier: The identifier of the database
    ///   instance to delete.
    func deleteDBInstance(instanceIdentifier: String) async {
        do {
            _ = try await rdsClient.deleteDBInstance(
                input: DeleteDBInstanceInput(
                    dbInstanceIdentifier: instanceIdentifier,
                    deleteAutomatedBackups: true,
                    skipFinalSnapshot: true
                )
            )
        } catch {
            print("*** Error deleting the database instance \(instanceIdentifier):
\(error.localizedDescription)")
        }
    }
```

```
   /// Wait until the specified database instance has been deleted.
   ///
   /// - Parameter instanceIdentifier: The identifier of the database
   ///   instance to wait for.
   func waitUntilDBInstanceDeleted(instanceIdentifier: String) async {
       putString("Waiting for the database instance to be deleted. This may take a
few minutes...")
       do {
           var isDatabaseDeleted = false
           var foundInstance = false

           while !isDatabaseDeleted {
               let output = try await rdsClient.describeDBInstances(input:
DescribeDBInstancesInput())
               guard let instanceList = output.dbInstances else {
                   return
               }

               foundInstance = false

               for instance in instanceList {
                   guard let foundInstanceIdentifier =
instance.dbInstanceIdentifier else {
                       continue
                   }

                   if instanceIdentifier == foundInstanceIdentifier {
                       foundInstance = true
                       break
                   } else {
                       putString(".")
                       do {
                           try await Task.sleep(for: .seconds(15))
                       } catch {
                           print("\n*** Error pausing the task!")
                       }
                   }
               }
               if !foundInstance {
                   isDatabaseDeleted = true
                   print()
               }
           }
```

```
        } catch {
            print("\n*** Error waiting for the database instance to be deleted:
\(error.localizedDescription)")
        }
    }

    /// Delete the specified database parameter group.
    ///
    /// - Parameter groupName: The name of the parameter group to delete.
    func deleteDBParameterGroup(groupName: String) async {
        do {
            _ = try await rdsClient.deleteDBParameterGroup(
                input: DeleteDBParameterGroupInput(
                    dbParameterGroupName: groupName
                )
            )
        } catch {
            print("*** Error deleting the database parameter group \(groupName):
\(error.localizedDescription)")
        }
    }

    /// Generate and return a unique file name that begins with the specified
    /// string.
    ///
    /// - Parameters:
    ///   - prefix: Text to use at the beginning of the returned name.
    ///
    /// - Returns: A string containing a unique filename that begins with the
    ///   specified `prefix`.
    ///
    /// The returned name uses a random number between 1 million and 1 billion to
    /// provide reasonable certainty of uniqueness for the purposes of this
    /// example.
    func tempName(prefix: String) -> String {
        return "\(prefix)-\(Int.random(in: 1000000..<1000000000))"
    }

    /// Print a string to standard output without a trailing newline, and
    /// without buffering.
    ///
    /// - Parameter str: The string to output.
    func putString(_ str: String = "") {
        if str.length >= 1 {
```

```
            let data = str.data(using: .utf8)
            guard let data else {
                return
            }
            FileHandle.standardOutput.write(data)
        }
    }
}

/// The program's asynchronous entry point.
@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.runAsync()
        } catch {
            ExampleCommand.exit(withError: error)
        }
    }
}
```

- For API details, see the following topics in *AWS SDK for Swift API reference*.
  - CreateDBInstance
  - CreateDBParameterGroup
  - CreateDBSnapshot
  - DeleteDBInstance
  - DeleteDBParameterGroup
  - DescribeDBEngineVersions
  - DescribeDBInstances
  - DescribeDBParameterGroups
  - DescribeDBParameters
  - DescribeDBSnapshots
  - DescribeOrderableDBInstanceOptions
  - ModifyDBParameterGroup

Actions

`CreateDBInstance`

The following code example shows how to use `CreateDBInstance`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSRDS

    /// Create a new database instance.
    ///
    /// - Parameters:
    ///   - name: The name of the database to create.
    ///   - instanceIdentifier: The identifier to give the new database
    ///     instance.
    ///   - parameterGroupName: The name of the parameter group to associate
    ///     with the new database instance.
    ///   - engine: The database engine to use.
    ///   - engineVersion: The version of the database given by `engine` to
    ///     use.
    ///   - instanceClass: The memory and compute capacity of the database
    ///     instance, such as `db.m5.large``.
    ///   - username: The admin user's username to establish for the new
    ///     instance.
    ///   - password: The password to use for the specified user's access.
    ///
    /// - Returns: A string indicating the ARN of the newly created database
    ///   instance, or nil if the instance couldn't be created.
    func createDBInstance(name: String, instanceIdentifier: String,
 parameterGroupName: String,
                          engine: String, engineVersion: String, instanceClass:
 String,
                          username: String, password: String) async -> String? {
        do {
            let output = try await rdsClient.createDBInstance(
```

```
                input: CreateDBInstanceInput(
                    allocatedStorage: 100,
                    dbInstanceClass: instanceClass,
                    dbInstanceIdentifier: instanceIdentifier,
                    dbName: name,
                    dbParameterGroupName: parameterGroupName,
                    engine: engine,
                    engineVersion: engineVersion,
                    masterUserPassword: password,
                    masterUsername: username,
                    storageType: "gp2"
                )
            )

            guard let dbInstance = output.dbInstance else {
                print("*** Unable to get the database instance.")
                return nil
            }

            return dbInstance.dbInstanceArn
        } catch {
            print("*** An error occurred while creating the database instance:
\(error.localizedDescription)")
            return nil
        }
    }
```

- For API details, see CreateDBInstance in *AWS SDK for Swift API reference*.

`CreateDBParameterGroup`

The following code example shows how to use `CreateDBParameterGroup`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSRDS

    /// Create a new database parameter group with the specified name.
    ///
    /// - Parameters:
    ///   - groupName: The name of the new parameter group.
    ///   - familyName: The name of the parameter group family.
    /// - Returns:
    func createDBParameterGroup(groupName: String, familyName: String) async ->
 RDSClientTypes.DBParameterGroup? {
        do {
            let output = try await rdsClient.createDBParameterGroup(
                input: CreateDBParameterGroupInput(
                    dbParameterGroupFamily: familyName,
                    dbParameterGroupName: groupName,
                    description: "Created using the AWS SDK for Swift"
                )
            )
            return output.dbParameterGroup
        } catch {
            print("*** Error creating the parameter group:
 \(error.localizedDescription)")
            return nil
        }
    }
```

- For API details, see CreateDBParameterGroup in *AWS SDK for Swift API reference*.

`CreateDBSnapshot`

The following code example shows how to use `CreateDBSnapshot`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSRDS

    /// Create a snapshot of the specified name.
    ///
    /// - Parameters:
    ///   - instanceIdentifier: The identifier of the database instance to
    ///     snapshot.
    ///   - snapshotIdentifier: A unique identifier to give the newly-created
    ///     snapshot.
    func createDBSnapshot(instanceIdentifier: String, snapshotIdentifier: String)
 async {
        do {
            let output = try await rdsClient.createDBSnapshot(
                input: CreateDBSnapshotInput(
                    dbInstanceIdentifier: instanceIdentifier,
                    dbSnapshotIdentifier: snapshotIdentifier
                )
            )

            guard let snapshot = output.dbSnapshot else {
                print("No snapshot returned.")
                return
            }

            print("The snapshot has been created with ID \(snapshot.dbiResourceId ??
 "<unknown>")")
        } catch {
            print("*** Unable to create the database snapshot named
 \(snapshotIdentifier): \(error.localizedDescription)")
        }
    }
```

- For API details, see CreateDBSnapshot in *AWS SDK for Swift API reference*.

`DeleteDBInstance`

The following code example shows how to use `DeleteDBInstance`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSRDS

    /// Delete the specified database instance.
    ///
    /// - Parameter instanceIdentifier: The identifier of the database
    ///   instance to delete.
    func deleteDBInstance(instanceIdentifier: String) async {
        do {
            _ = try await rdsClient.deleteDBInstance(
                input: DeleteDBInstanceInput(
                    dbInstanceIdentifier: instanceIdentifier,
                    deleteAutomatedBackups: true,
                    skipFinalSnapshot: true
                )
            )
        } catch {
            print("*** Error deleting the database instance \(instanceIdentifier):
 \(error.localizedDescription)")
        }
    }
```

- For API details, see DeleteDBInstance in *AWS SDK for Swift API reference*.

`DeleteDBParameterGroup`

The following code example shows how to use `DeleteDBParameterGroup`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSRDS

    /// Delete the specified database parameter group.
    ///
    /// - Parameter groupName: The name of the parameter group to delete.
    func deleteDBParameterGroup(groupName: String) async {
        do {
            _ = try await rdsClient.deleteDBParameterGroup(
                input: DeleteDBParameterGroupInput(
                    dbParameterGroupName: groupName
                )
            )
        } catch {
            print("*** Error deleting the database parameter group \(groupName):
 \(error.localizedDescription)")
        }
    }
```

- For API details, see DeleteDBParameterGroup in *AWS SDK for Swift API reference*.

`DescribeDBEngineVersions`

The following code example shows how to use `DescribeDBEngineVersions`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSRDS

    /// Get all the database engine versions available for the specified
    /// database engine.
    ///
    /// - Parameter engineName: The name of the database engine to query.
    ///
    /// - Returns: An array of `RDSClientTypes.DBEngineVersion` structures,
    ///   each describing one supported version of the specified database.
    func getDBEngineVersions(engineName: String) async ->
 [RDSClientTypes.DBEngineVersion] {
        do {
            let output = try await rdsClient.describeDBEngineVersions(
                input: DescribeDBEngineVersionsInput(
                    engine: engineName
                )
            )

            return output.dbEngineVersions ?? []
        } catch {
            return []
        }
    }
```

- For API details, see DescribeDBEngineVersions in *AWS SDK for Swift API reference*.

`DescribeDBInstances`

The following code example shows how to use `DescribeDBInstances`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSRDS
```

```
   /// Wait until the specified database is available to use.
   ///
   /// - Parameter instanceIdentifier: The database instance identifier of the
   ///   database to wait for.
   func waitUntilDBInstanceReady(instanceIdentifier: String) async ->
RDSClientTypes.Endpoint? {
       do {
           putString("Waiting for the database instance to be ready to use. This
may take 10 minutes or more...")
           while true {
               let output = try await rdsClient.describeDBInstances(
                   input: DescribeDBInstancesInput(
                       dbInstanceIdentifier: instanceIdentifier
                   )
               )

               guard let instanceList = output.dbInstances else {
                   continue
               }

               for instance in instanceList {
                   let status = instance.dbInstanceStatus

                   guard let status else {
                       print("\nUnable to determine the status.")
                       continue
                   }

                   if status.contains("available") {
                       return instance.endpoint
                   } else {
                       putString(".")
                       do {
                           try await Task.sleep(for: .seconds(15))
                       } catch {
                           print("*** Error pausing the task!")
                       }
                   }
               }
           }
       } catch {
           print("*** Unable to wait until the database is ready:
\(error.localizedDescription)")
           return nil
```

```
        }
    }
```

- For API details, see DescribeDBInstances in *AWS SDK for Swift API reference*.

`DescribeDBParameterGroups`

The following code example shows how to use `DescribeDBParameterGroups`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSRDS

    /// Get descriptions of the database parameter groups matching the given
    /// name.
    ///
    /// - Parameter groupName: The name of the parameter group to describe.
    ///
    /// - Returns: An array of [RDSClientTypes.DBParameterGroup] objects
    ///   describing the parameter group.
    func describeDBParameterGroups(groupName: String) async ->
 [RDSClientTypes.DBParameterGroup]? {
        do {
            let output = try await rdsClient.describeDBParameterGroups(
                input: DescribeDBParameterGroupsInput(
                    dbParameterGroupName: groupName
                )
            )
            return output.dbParameterGroups
        } catch {
            print("*** Error getting the database parameter group's details:
 \(error.localizedDescription)")
            return nil
        }
```

```
    }
```

- For API details, see DescribeDBParameterGroups in *AWS SDK for Swift API reference*.

`DescribeDBParameters`

The following code example shows how to use `DescribeDBParameters`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSRDS

    /// Returns the detailed parameter list for the specified database
    /// parameter group.
    ///
    /// - Parameters:
    ///   - groupName: The name of the parameter group to return parameters for.
    ///   - source: The types of parameters to return (`user`, `system`, or
    ///     `engine-default`).
    ///
    /// - Returns: An array of `RdSClientTypes.Parameter` objects, each
    ///   describing one of the group's parameters.
    func describeDBParameters(groupName: String, source: String? = nil) async ->
 [RDSClientTypes.Parameter] {
        var parameterList: [RDSClientTypes.Parameter] = []

        do {
            let pages = rdsClient.describeDBParametersPaginated(
                input: DescribeDBParametersInput(
                    dbParameterGroupName: groupName,
                    source: source
                )
            )
```

```
            for try await page in pages {
                guard let parameters = page.parameters else {
                    return []
                }

                parameterList += parameters
            }
        } catch {
            print("*** Error getting database parameters:
 \(error.localizedDescription)")
            return []
        }

        return parameterList
    }
```

- For API details, see DescribeDBParameters in *AWS SDK for Swift API reference*.

`DescribeDBSnapshots`

The following code example shows how to use `DescribeDBSnapshots`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSRDS

    /// Wait until the specified database snapshot is available to use.
    ///
    /// - Parameters:
    ///   - instanceIdentifier: The identifier of the database for which the
    ///     snapshot was taken.
    ///   - snapshotIdentifier: The identifier of the snapshot to wait for.
    func waitUntilDBSnapshotReady(instanceIdentifier: String, snapshotIdentifier:
 String) async {
```

```
        var snapshotReady = false

        putString("Waiting for the snapshot to be ready...")

        do {
            while !snapshotReady {
                let output = try await rdsClient.describeDBSnapshots(
                    input: DescribeDBSnapshotsInput(
                        dbInstanceIdentifier: instanceIdentifier,
                        dbSnapshotIdentifier: snapshotIdentifier
                    )
                )

                guard let snapshotList = output.dbSnapshots else {
                    return
                }

                for snapshot in snapshotList {
                    guard let snapshotReadyStr = snapshot.status else {
                        return
                    }

                    if snapshotReadyStr.contains("available") {
                        snapshotReady = true
                        print()
                    } else {
                        putString(".")
                        do {
                            try await Task.sleep(for: .seconds(15))
                        } catch {
                            print("\n*** Error pausing the task!")
                        }
                    }
                }
            }
        } catch {
            print("\n*** Unable to wait for the database snapshot to be ready:
\(error.localizedDescription)")
        }
    }
```

- For API details, see DescribeDBSnapshots in *AWS SDK for Swift API reference*.

`DescribeOrderableDBInstanceOptions`

The following code example shows how to use `DescribeOrderableDBInstanceOptions`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSRDS

    /// Print a list of available database instances with "micro" in the class
    /// name, then return one of them to be used by other code.
    ///
    /// - Parameters:
    ///   - engine: The database engine for which to list database instance
    ///     classes.
    ///   - engineVersion: The database version for which to list instances.
    ///
    /// - Returns: An `RDSClientTypes.OrderableDBInstanceOption` describing
    ///   the selected instance type.
    func chooseMicroInstance(engine: String = "mysql", engineVersion: String? = nil)
 async -> RDSClientTypes.OrderableDBInstanceOption? {
        do {
            let pages = rdsClient.describeOrderableDBInstanceOptionsPaginated(
                input: DescribeOrderableDBInstanceOptionsInput(
                    engine: engine,
                    engineVersion: engineVersion
                )
            )

            var optionsList: [RDSClientTypes.OrderableDBInstanceOption] = []

            for try await page in pages {
                guard let orderableDBInstanceOptions =
 page.orderableDBInstanceOptions else {
                    continue
                }

                for dbInstanceOption in orderableDBInstanceOptions {
```

```
                    guard let className = dbInstanceOption.dbInstanceClass else {
                        continue
                    }
                    if className.contains("micro") {
                        optionsList.append(dbInstanceOption)
                    }
                }
            }

            print("Found \(optionsList.count) database instances of 'micro' class
types:")
            for dbInstanceOption in optionsList {
                print("    \(dbInstanceOption.engine ??
"<unknown>") \(dbInstanceOption.engineVersion ?? "<unknown>")
(\(dbInstanceOption.dbInstanceClass ?? "<unknown class>"))")
            }

            return optionsList[0]
        } catch {
            print("*** Error getting a list of orderable instance options:
\(error.localizedDescription)")
            return nil
        }
    }
```

- For API details, see DescribeOrderableDBInstanceOptions in *AWS SDK for Swift API reference*.

`ModifyDBParameterGroup`

The following code example shows how to use `ModifyDBParameterGroup`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSRDS
```

```
    /// Demonstrates modifying two of the specified database parameter group's
    /// parameters.
    ///
    /// - Parameter groupName: The name of the parameter group to change
    ///   parameters for.
    func modifyDBParameters(groupName: String) async {
        let parameter1 = RDSClientTypes.Parameter(
            applyMethod: RDSClientTypes.ApplyMethod.immediate,
            parameterName: "auto_increment_offset",
            parameterValue: "5"
        )
        let parameter2 = RDSClientTypes.Parameter(
            applyMethod: RDSClientTypes.ApplyMethod.immediate,
            parameterName: "auto_increment_increment",
            parameterValue: "5"
        )

        let parameterList = [parameter1, parameter2]

        do {
            _ = try await rdsClient.modifyDBParameterGroup(
                input: ModifyDBParameterGroupInput(
                    dbParameterGroupName: groupName,
                    parameters: parameterList
                )
            )

            print("Successfully modified the parameter group \(groupName).")
        } catch {
            print("*** Error modifying the parameter group \(groupName):
\(error.localizedDescription)")
        }
    }
```

- For API details, see ModifyDBParameterGroup in *AWS SDK for Swift API reference*.

Amazon S3 examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with Amazon S3.

Basics are code examples that show you how to perform the essential operations within a service.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Scenarios are code examples that show you how to accomplish specific tasks by calling multiple functions within a service or combined with other AWS services.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Basics
- Actions
- Scenarios

Basics

Learn the basics

The following code example shows how to:

- Create a bucket and upload a file to it.
- Download an object from a bucket.
- Copy an object to a subfolder in a bucket.
- List the objects in a bucket.
- Delete the bucket objects and the bucket.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSS3
```

```
import Foundation
import AWSS3
import Smithy
import ClientRuntime

/// A class containing all the code that interacts with the AWS SDK for Swift.
public class ServiceHandler {
    let configuration: S3Client.S3ClientConfig
    let client: S3Client

    enum HandlerError: Error {
        case getObjectBody(String)
        case readGetObjectBody(String)
        case missingContents(String)
    }

    /// Initialize and return a new ``ServiceHandler`` object, which is used to
 drive the AWS calls
    /// used for the example.
    ///
    /// - Returns: A new ``ServiceHandler`` object, ready to be called to
    ///            execute AWS operations.
    public init() async throws {
        do {
            // To set the Region programmatically, pass it to the initializer:
            //   configuration = try await S3Client.S3ClientConfig(region: "us-
east-2")
            configuration = try await S3Client.S3ClientConfig()
            client = S3Client(config: configuration)
        }
        catch {
            print("ERROR: ", dump(error, name: "Initializing S3 client"))
            throw error
        }
    }


    /// Create a new user given the specified name.
    ///
    /// - Parameters:
    ///   - name: Name of the bucket to create.
    /// Throws an exception if an error occurs.
    public func createBucket(name: String) async throws {
        var input = CreateBucketInput(
```

```
            bucket: name
        )

        // For regions other than "us-east-1", you must set the locationConstraint
 in the createBucketConfiguration.
        // For more information, see LocationConstraint in the S3 API guide.
        // https://docs.aws.amazon.com/AmazonS3/latest/API/
API_CreateBucket.html#API_CreateBucket_RequestBody
        if let region = configuration.region {
            if region != "us-east-1" {
                input.createBucketConfiguration =
 S3ClientTypes.CreateBucketConfiguration(locationConstraint:
 S3ClientTypes.BucketLocationConstraint(rawValue: region))
            }
        }

        do {
            _ = try await client.createBucket(input: input)
        }
        catch let error as BucketAlreadyOwnedByYou {
            print("The bucket '\(name)' already exists and is owned by you. You may
 wish to ignore this exception.")
            throw error
        }
        catch {
            print("ERROR: ", dump(error, name: "Creating a bucket"))
            throw error
        }
    }

    /// Delete a bucket.
    /// - Parameter name: Name of the bucket to delete.
    public func deleteBucket(name: String) async throws {
        let input = DeleteBucketInput(
            bucket: name
        )
        do {
            _ = try await client.deleteBucket(input: input)
        }
        catch {
            print("ERROR: ", dump(error, name: "Deleting a bucket"))
            throw error
        }
    }
```

```
    /// Upload a file from local storage to the bucket.
    /// - Parameters:
    ///   - bucket: Name of the bucket to upload the file to.
    ///   - key: Name of the file to create.
    ///   - file: Path name of the file to upload.
    public func uploadFile(bucket: String, key: String, file: String) async throws {
        let fileUrl = URL(fileURLWithPath: file)
        do {
            let fileData = try Data(contentsOf: fileUrl)
            let dataStream = ByteStream.data(fileData)

            let input = PutObjectInput(
                body: dataStream,
                bucket: bucket,
                key: key
            )

            _ = try await client.putObject(input: input)
        }
        catch {
            print("ERROR: ", dump(error, name: "Putting an object."))
            throw error
        }
    }

    /// Create a file in the specified bucket with the given name. The new
    /// file's contents are uploaded from a `Data` object.
    ///
    /// - Parameters:
    ///   - bucket: Name of the bucket to create a file in.
    ///   - key: Name of the file to create.
    ///   - data: A `Data` object to write into the new file.
    public func createFile(bucket: String, key: String, withData data: Data) async
throws {
        let dataStream = ByteStream.data(data)

        let input = PutObjectInput(
            body: dataStream,
            bucket: bucket,
            key: key
        )

        do {
```

```
            _ = try await client.putObject(input: input)
        }
        catch {
            print("ERROR: ", dump(error, name: "Putting an object."))
            throw error
        }
    }

    /// Download the named file to the given directory on the local device.
    ///
    /// - Parameters:
    ///   - bucket: Name of the bucket that contains the file to be copied.
    ///   - key: The name of the file to copy from the bucket.
    ///   - to: The path of the directory on the local device where you want to
    ///     download the file.
    public func downloadFile(bucket: String, key: String, to: String) async throws {
        let fileUrl = URL(fileURLWithPath: to).appendingPathComponent(key)

        let input = GetObjectInput(
            bucket: bucket,
            key: key
        )
        do {
            let output = try await client.getObject(input: input)

            guard let body = output.body else {
                throw HandlerError.getObjectBody("GetObjectInput missing body.")
            }

            guard let data = try await body.readData() else {
                throw HandlerError.readGetObjectBody("GetObjectInput unable to read
data.")
            }

            try data.write(to: fileUrl)
        }
        catch {
            print("ERROR: ", dump(error, name: "Downloading a file."))
            throw error
        }
    }

    /// Read the specified file from the given S3 bucket into a Swift
    /// `Data` object.
```

```
    ///
    /// - Parameters:
    ///   - bucket: Name of the bucket containing the file to read.
    ///   - key: Name of the file within the bucket to read.
    ///
    /// - Returns: A `Data` object containing the complete file data.
    public func readFile(bucket: String, key: String) async throws -> Data {
        let input = GetObjectInput(
            bucket: bucket,
            key: key
        )
        do {
            let output = try await client.getObject(input: input)

            guard let body = output.body else {
                throw HandlerError.getObjectBody("GetObjectInput missing body.")
            }

            guard let data = try await body.readData() else {
                throw HandlerError.readGetObjectBody("GetObjectInput unable to read
 data.")
            }

            return data
        }
        catch {
            print("ERROR: ", dump(error, name: "Reading a file."))
            throw error
        }
   }


    /// Copy a file from one bucket to another.
    ///
    /// - Parameters:
    ///   - sourceBucket: Name of the bucket containing the source file.
    ///   - name: Name of the source file.
    ///   - destBucket: Name of the bucket to copy the file into.
    public func copyFile(from sourceBucket: String, name: String, to destBucket:
 String) async throws {
        let srcUrl = ("\(sourceBucket)/
\(name)").addingPercentEncoding(withAllowedCharacters: .urlPathAllowed)

        let input = CopyObjectInput(
```

```
            bucket: destBucket,
            copySource: srcUrl,
            key: name
        )
        do {
            _ = try await client.copyObject(input: input)
        }
        catch {
            print("ERROR: ", dump(error, name: "Copying an object."))
            throw error
        }
    }

    /// Deletes the specified file from Amazon S3.
    ///
    /// - Parameters:
    ///   - bucket: Name of the bucket containing the file to delete.
    ///   - key: Name of the file to delete.
    ///
    public func deleteFile(bucket: String, key: String) async throws {
        let input = DeleteObjectInput(
            bucket: bucket,
            key: key
        )

        do {
            _ = try await client.deleteObject(input: input)
        }
        catch {
            print("ERROR: ", dump(error, name: "Deleting a file."))
            throw error
        }
    }

    /// Returns an array of strings, each naming one file in the
    /// specified bucket.
    ///
    /// - Parameter bucket: Name of the bucket to get a file listing for.
    /// - Returns: An array of `String` objects, each giving the name of
    ///            one file contained in the bucket.
    public func listBucketFiles(bucket: String) async throws -> [String] {
        do {
            let input = ListObjectsV2Input(
                bucket: bucket
```

```
            )

            // Use "Paginated" to get all the objects.
            // This lets the SDK handle the 'continuationToken' in
 "ListObjectsV2Output".
            let output = client.listObjectsV2Paginated(input: input)
            var names: [String] = []

            for try await page in output {
                guard let objList = page.contents else {
                    print("ERROR: listObjectsV2Paginated returned nil contents.")
                    continue
                }

                for obj in objList {
                    if let objName = obj.key {
                        names.append(objName)
                    }
                }
            }


            return names
        }
        catch {
            print("ERROR: ", dump(error, name: "Listing objects."))
            throw error
        }
    }
}
```

```
import AWSS3

import Foundation
import ServiceHandler
import ArgumentParser

/// The command-line arguments and options available for this
/// example command.
struct ExampleCommand: ParsableCommand {
    @Argument(help: "Name of the S3 bucket to create")
    var bucketName: String
```

```
    @Argument(help: "Pathname of the file to upload to the S3 bucket")
    var uploadSource: String

    @Argument(help: "The name (key) to give the file in the S3 bucket")
    var objName: String

    @Argument(help: "S3 bucket to copy the object to")
    var destBucket: String

    @Argument(help: "Directory where you want to download the file from the S3
bucket")
    var downloadDir: String

    static var configuration = CommandConfiguration(
        commandName: "s3-basics",
        abstract: "Demonstrates a series of basic AWS S3 functions.",
        discussion: """
        Performs the following Amazon S3 commands:

        * `CreateBucket`
        * `PutObject`
        * `GetObject`
        * `CopyObject`
        * `ListObjects`
        * `DeleteObjects`
        * `DeleteBucket`
        """
    )

    /// Called by ``main()`` to do the actual running of the AWS
    /// example.
    func runAsync() async throws {
        let serviceHandler = try await ServiceHandler()

        // 1. Create the bucket.
        print("Creating the bucket \(bucketName)...")
        try await serviceHandler.createBucket(name: bucketName)

        // 2. Upload a file to the bucket.
        print("Uploading the file \(uploadSource)...")
        try await serviceHandler.uploadFile(bucket: bucketName, key: objName, file:
uploadSource)
```

```
        // 3. Download the file.
        print("Downloading the file \(objName) to \(downloadDir)...")
        try await serviceHandler.downloadFile(bucket: bucketName, key: objName, to:
 downloadDir)

        // 4. Copy the file to another bucket.
        print("Copying the file to the bucket \(destBucket)...")
        try await serviceHandler.copyFile(from: bucketName, name: objName, to:
 destBucket)

        // 5. List the contents of the bucket.

        print("Getting a list of the files in the bucket \(bucketName)")
        let fileList = try await serviceHandler.listBucketFiles(bucket: bucketName)
        let numFiles = fileList.count
        if numFiles != 0 {
            print("\(numFiles) file\((numFiles > 1) ? "s" : "") in bucket
 \(bucketName):")
            for name in fileList {
                print("  \(name)")
            }
        } else {
            print("No files found in bucket \(bucketName)")
        }

        // 6. Delete the objects from the bucket.

        print("Deleting the file \(objName) from the bucket \(bucketName)...")
        try await serviceHandler.deleteFile(bucket: bucketName, key: objName)
        print("Deleting the file \(objName) from the bucket \(destBucket)...")
        try await serviceHandler.deleteFile(bucket: destBucket, key: objName)

        // 7. Delete the bucket.
        print("Deleting the bucket \(bucketName)...")
        try await serviceHandler.deleteBucket(name: bucketName)

        print("Done.")
    }
}

//
// Main program entry point.
//
@main
```

```
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.runAsync()
        } catch {
            ExampleCommand.exit(withError: error)
        }
    }
}
```

- For API details, see the following topics in *AWS SDK for Swift API reference*.
  - CopyObject
  - CreateBucket
  - DeleteBucket
  - DeleteObjects
  - GetObject
  - ListObjectsV2
  - PutObject

Actions

`CopyObject`

The following code example shows how to use `CopyObject`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSS3
```

```
    public func copyFile(from sourceBucket: String, name: String, to destBucket:
 String) async throws {
        let srcUrl = ("\(sourceBucket)/
\(name)").addingPercentEncoding(withAllowedCharacters: .urlPathAllowed)

        let input = CopyObjectInput(
            bucket: destBucket,
            copySource: srcUrl,
            key: name
        )
        do {
            _ = try await client.copyObject(input: input)
        }
        catch {
            print("ERROR: ", dump(error, name: "Copying an object."))
            throw error
        }
    }
```

- For API details, see CopyObject in *AWS SDK for Swift API reference*.

`CreateBucket`

The following code example shows how to use `CreateBucket`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSS3

    public func createBucket(name: String) async throws {
        var input = CreateBucketInput(
            bucket: name
        )
```

```
        // For regions other than "us-east-1", you must set the locationConstraint
 in the createBucketConfiguration.
        // For more information, see LocationConstraint in the S3 API guide.
        // https://docs.aws.amazon.com/AmazonS3/latest/API/
API_CreateBucket.html#API_CreateBucket_RequestBody
        if let region = configuration.region {
            if region != "us-east-1" {
                input.createBucketConfiguration =
 S3ClientTypes.CreateBucketConfiguration(locationConstraint:
 S3ClientTypes.BucketLocationConstraint(rawValue: region))
            }
        }

        do {
            _ = try await client.createBucket(input: input)
        }
        catch let error as BucketAlreadyOwnedByYou {
            print("The bucket '\(name)' already exists and is owned by you. You may
 wish to ignore this exception.")
            throw error
        }
        catch {
            print("ERROR: ", dump(error, name: "Creating a bucket"))
            throw error
        }
    }
```

- For API details, see CreateBucket in *AWS SDK for Swift API reference*.

`DeleteBucket`

The following code example shows how to use `DeleteBucket`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSS3

    public func deleteBucket(name: String) async throws {
        let input = DeleteBucketInput(
            bucket: name
        )
        do {
            _ = try await client.deleteBucket(input: input)
        }
        catch {
            print("ERROR: ", dump(error, name: "Deleting a bucket"))
            throw error
        }
    }
```

- For API details, see DeleteBucket in *AWS SDK for Swift API reference*.

`DeleteObject`

The following code example shows how to use `DeleteObject`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSS3

    public func deleteFile(bucket: String, key: String) async throws {
        let input = DeleteObjectInput(
            bucket: bucket,
            key: key
        )

        do {
            _ = try await client.deleteObject(input: input)
        }
```

```
        catch {
            print("ERROR: ", dump(error, name: "Deleting a file."))
            throw error
        }
    }
```

- For API details, see DeleteObject in *AWS SDK for Swift API reference*.

`DeleteObjects`

The following code example shows how to use `DeleteObjects`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSS3

    public func deleteObjects(bucket: String, keys: [String]) async throws {
        let input = DeleteObjectsInput(
            bucket: bucket,
            delete: S3ClientTypes.Delete(
                objects: keys.map { S3ClientTypes.ObjectIdentifier(key: $0) },
                quiet: true
            )
        )

        do {
            _ = try await client.deleteObjects(input: input)
        } catch {
            print("ERROR: deleteObjects:", dump(error))
            throw error
        }
    }
```

- For API details, see DeleteObjects in *AWS SDK for Swift API reference*.

`GetObject`

The following code example shows how to use `GetObject`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSS3

    public func downloadFile(bucket: String, key: String, to: String) async throws {
        let fileUrl = URL(fileURLWithPath: to).appendingPathComponent(key)

        let input = GetObjectInput(
            bucket: bucket,
            key: key
        )
        do {
            let output = try await client.getObject(input: input)

            guard let body = output.body else {
                throw HandlerError.getObjectBody("GetObjectInput missing body.")
            }

            guard let data = try await body.readData() else {
                throw HandlerError.readGetObjectBody("GetObjectInput unable to read
 data.")
            }

            try data.write(to: fileUrl)
        }
        catch {
            print("ERROR: ", dump(error, name: "Downloading a file."))
            throw error
        }
    }
```

```
import AWSS3

    public func readFile(bucket: String, key: String) async throws -> Data {
        let input = GetObjectInput(
            bucket: bucket,
            key: key
        )
        do {
            let output = try await client.getObject(input: input)

            guard let body = output.body else {
                throw HandlerError.getObjectBody("GetObjectInput missing body.")
            }

            guard let data = try await body.readData() else {
                throw HandlerError.readGetObjectBody("GetObjectInput unable to read
 data.")
            }

            return data
        }
        catch {
            print("ERROR: ", dump(error, name: "Reading a file."))
            throw error
        }
   }
```

- For API details, see GetObject in *AWS SDK for Swift API reference*.

`ListBuckets`

The following code example shows how to use `ListBuckets`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSS3

    /// Return an array containing information about every available bucket.
    ///
    /// - Returns: An array of ``S3ClientTypes.Bucket`` objects describing
    ///   each bucket.
    public func getAllBuckets() async throws -> [S3ClientTypes.Bucket] {
        return try await client.listBuckets(input: ListBucketsInput())
    }
```

- For API details, see ListBuckets in *AWS SDK for Swift API reference*.

`ListObjectsV2`

The following code example shows how to use `ListObjectsV2`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSS3

    public func listBucketFiles(bucket: String) async throws -> [String] {
        do {
            let input = ListObjectsV2Input(
                bucket: bucket
            )
```

```
            // Use "Paginated" to get all the objects.
            // This lets the SDK handle the 'continuationToken' in
"ListObjectsV2Output".
            let output = client.listObjectsV2Paginated(input: input)
            var names: [String] = []

            for try await page in output {
                guard let objList = page.contents else {
                    print("ERROR: listObjectsV2Paginated returned nil contents.")
                    continue
                }

                for obj in objList {
                    if let objName = obj.key {
                        names.append(objName)
                    }
                }
            }


            return names
        }
        catch {
            print("ERROR: ", dump(error, name: "Listing objects."))
            throw error
        }
    }
```

- For API details, see ListObjectsV2 in *AWS SDK for Swift API reference*.

`PutObject`

The following code example shows how to use `PutObject`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSS3
import Smithy

    public func uploadFile(bucket: String, key: String, file: String) async throws {
        let fileUrl = URL(fileURLWithPath: file)
        do {
            let fileData = try Data(contentsOf: fileUrl)
            let dataStream = ByteStream.data(fileData)

            let input = PutObjectInput(
                body: dataStream,
                bucket: bucket,
                key: key
            )

            _ = try await client.putObject(input: input)
        }
        catch {
            print("ERROR: ", dump(error, name: "Putting an object."))
            throw error
        }
    }
```

```
import AWSS3
import Smithy

    public func createFile(bucket: String, key: String, withData data: Data) async
 throws {
        let dataStream = ByteStream.data(data)

        let input = PutObjectInput(
            body: dataStream,
            bucket: bucket,
            key: key
        )

        do {
            _ = try await client.putObject(input: input)
        }
        catch {
            print("ERROR: ", dump(error, name: "Putting an object."))
            throw error
```

```
        }
    }
```

- For API details, see PutObject in *AWS SDK for Swift API reference*.

Scenarios

Download stream of unknown size

The following code example shows how to download a stream of unknown size from an Amazon S3 object.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import ArgumentParser
import AWSClientRuntime
import AWSS3
import Foundation
import Smithy
import SmithyHTTPAPI
import SmithyStreams


    /// Download a file from the specified bucket.
    ///
    /// - Parameters:
    ///   - bucket: The Amazon S3 bucket name to get the file from.
    ///   - key: The name (or path) of the file to download from the bucket.
    ///   - destPath: The pathname on the local filesystem at which to store
    ///     the downloaded file.
    func downloadFile(bucket: String, key: String, destPath: String?) async throws {
        let fileURL: URL

        // If no destination path was provided, use the key as the name to use
```

```
        // for the file in the downloads folder.

        if destPath == nil {
            do {
                try fileURL = FileManager.default.url(
                    for: .downloadsDirectory,
                    in: .userDomainMask,
                    appropriateFor: URL(string: key),
                    create: true
                ).appendingPathComponent(key)
            } catch {
                throw TransferError.directoryError
            }
        } else {
            fileURL = URL(fileURLWithPath: destPath!)
        }

        let config = try await S3Client.S3ClientConfig(region: region)
        let s3Client = S3Client(config: config)

        // Create a `FileHandle` referencing the local destination. Then
        // create a `ByteStream` from that.

        FileManager.default.createFile(atPath: fileURL.path, contents: nil,
attributes: nil)
        let fileHandle = try FileHandle(forWritingTo: fileURL)

        // Download the file using `GetObject`.

        let getInput = GetObjectInput(
            bucket: bucket,
            key: key
        )

        do {
            let getOutput = try await s3Client.getObject(input: getInput)

            guard let body = getOutput.body else {
                throw TransferError.downloadError("Error: No data returned for
download")
            }

            // If the body is returned as a `Data` object, write that to the
            // file. If it's a stream, read the stream chunk by chunk,
```

```
            // appending each chunk to the destination file.

            switch body {
            case .data:
                guard let data = try await body.readData() else {
                    throw TransferError.downloadError("Download error")
                }

                // Write the `Data` to the file.

                do {
                    try data.write(to: fileURL)
                } catch {
                    throw TransferError.writeError
                }
                break

            case .stream(let stream as ReadableStream):
                while (true) {
                    let chunk = try await stream.readAsync(upToCount: 5 * 1024 *
1024)
                    guard let chunk = chunk else {
                        break
                    }

                    // Write the chunk to the destination file.

                    do {
                        try fileHandle.write(contentsOf: chunk)
                    } catch {
                        throw TransferError.writeError
                    }
                }

                break
            default:
                throw TransferError.downloadError("Received data is unknown object
type")
            }
        } catch {
            throw TransferError.downloadError("Error downloading the file:
\(error)")
        }
```

```
        print("File downloaded to \(fileURL.path).")
    }
```

Upload stream of unknown size

The following code example shows how to upload a stream of unknown size to an Amazon S3 object.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import ArgumentParser
import AWSClientRuntime
import AWSS3
import Foundation
import Smithy
import SmithyHTTPAPI
import SmithyStreams


    /// Upload a file to the specified bucket.
    ///
    /// - Parameters:
    ///   - bucket: The Amazon S3 bucket name to store the file into.
    ///   - key: The name (or path) of the file to upload to in the `bucket`.
    ///   - sourcePath: The pathname on the local filesystem of the file to
    ///     upload.
    func uploadFile(sourcePath: String, bucket: String, key: String?) async throws {
        let fileURL: URL = URL(fileURLWithPath: sourcePath)
        let fileName: String

        // If no key was provided, use the last component of the filename.

        if key == nil {
            fileName = fileURL.lastPathComponent
```

```
        } else {
            fileName = key!
        }

        let s3Client = try await S3Client()

        // Create a FileHandle for the source file.

        let fileHandle = FileHandle(forReadingAtPath: sourcePath)
        guard let fileHandle = fileHandle else {
            throw TransferError.readError
        }

        // Create a byte stream to retrieve the file's contents. This uses the
        // Smithy FileStream and ByteStream types.

        let stream = FileStream(fileHandle: fileHandle)
        let body = ByteStream.stream(stream)

        // Create a `PutObjectInput` with the ByteStream as the body of the
        // request's data. The AWS SDK for Swift will handle sending the
        // entire file in chunks, regardless of its size.

        let putInput = PutObjectInput(
            body: body,
            bucket: bucket,
            key: fileName
        )

        do {
            _ = try await s3Client.putObject(input: putInput)
        } catch {
            throw TransferError.uploadError("Error uploading the file: \(error)")
        }

        print("File uploaded to \(fileURL.path).")
    }
```

Amazon SNS examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with Amazon SNS.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Scenarios are code examples that show you how to accomplish specific tasks by calling multiple functions within a service or combined with other AWS services.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Get started
- Actions
- Scenarios

Get started

Hello Amazon SNS

The following code example shows how to get started using Amazon SNS.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

The Package.swift file.

```
import PackageDescription

let package = Package(
    name: "sns-basics",
```

```
    // Let Xcode know the minimum Apple platforms supported.
    platforms: [
        .macOS(.v13),
        .iOS(.v15)
    ],
    dependencies: [
        // Dependencies declare other packages that this package depends on.
        .package(
            url: "https://github.com/awslabs/aws-sdk-swift",
            from: "1.0.0"),
        .package(
            url: "https://github.com/apple/swift-argument-parser.git",
            branch: "main"
        )
    ],
    targets: [
        // Targets are the basic building blocks of a package, defining a module or
 a test suite.
        // Targets can depend on other targets in this package and products
        // from dependencies.
        .executableTarget(
            name: "sns-basics",
            dependencies: [
                .product(name: "AWSSNS", package: "aws-sdk-swift"),
                .product(name: "ArgumentParser", package: "swift-argument-parser")
            ],
            path: "Sources")

    ]
)
```

The main Swift program.

```
import ArgumentParser
import AWSClientRuntime
import AWSSNS
import Foundation

struct ExampleCommand: ParsableCommand {
    @Option(help: "Name of the Amazon Region to use (default: us-east-1)")
    var region = "us-east-1"
```

```
    static var configuration = CommandConfiguration(
        commandName: "sns-basics",
        abstract: """
        This example shows how to list all of your available Amazon SNS topics.
        """,
        discussion: """
        """
    )

    /// Called by ``main()`` to run the bulk of the example.
    func runAsync() async throws {
        let config = try await SNSClient.SNSClientConfig(region: region)
        let snsClient = SNSClient(config: config)

        var topics: [String] = []
        let outputPages = snsClient.listTopicsPaginated(
            input: ListTopicsInput()
        )

        // Each time a page of results arrives, process its contents.

        for try await output in outputPages {
            guard let topicList = output.topics else {
                print("Unable to get a page of Amazon SNS topics.")
                return
            }

            // Iterate over the topics listed on this page, adding their ARNs
            // to the `topics` array.

            for topic in topicList {
                guard let arn = topic.topicArn else {
                    print("Topic has no ARN.")
                    return
                }
                topics.append(arn)
            }
        }

        print("You have \(topics.count) topics:")
        for topic in topics {
            print("   \(topic)")
        }
    }
```

```
}

/// The program's asynchronous entry point.
@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.runAsync()
        } catch {
            ExampleCommand.exit(withError: error)
        }
    }
}
```

- For API details, see ListTopics in *AWS SDK for Swift API reference*.

Actions

`CreateTopic`

The following code example shows how to use `CreateTopic`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSNS

        let config = try await SNSClient.SNSClientConfig(region: region)
        let snsClient = SNSClient(config: config)

        let output = try await snsClient.createTopic(
            input: CreateTopicInput(name: name)
```

```
        )

        guard let arn = output.topicArn else {
            print("No topic ARN returned by Amazon SNS.")
            return
        }
```

- For API details, see CreateTopic in *AWS SDK for Swift API reference*.

`DeleteTopic`

The following code example shows how to use `DeleteTopic`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSNS

        let config = try await SNSClient.SNSClientConfig(region: region)
        let snsClient = SNSClient(config: config)

        _ = try await snsClient.deleteTopic(
            input: DeleteTopicInput(topicArn: arn)
        )
```

- For API details, see DeleteTopic in *AWS SDK for Swift API reference*.

`ListTopics`

The following code example shows how to use `ListTopics`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSNS

        let config = try await SNSClient.SNSClientConfig(region: region)
        let snsClient = SNSClient(config: config)

        var topics: [String] = []
        let outputPages = snsClient.listTopicsPaginated(
            input: ListTopicsInput()
        )

        // Each time a page of results arrives, process its contents.

        for try await output in outputPages {
            guard let topicList = output.topics else {
                print("Unable to get a page of Amazon SNS topics.")
                return
            }

            // Iterate over the topics listed on this page, adding their ARNs
            // to the `topics` array.

            for topic in topicList {
                guard let arn = topic.topicArn else {
                    print("Topic has no ARN.")
                    return
                }
                topics.append(arn)
            }
        }
```

- For API details, see ListTopics in *AWS SDK for Swift API reference*.

Publish

The following code example shows how to use `Publish`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSNS

        let config = try await SNSClient.SNSClientConfig(region: region)
        let snsClient = SNSClient(config: config)

        let output = try await snsClient.publish(
            input: PublishInput(
                message: message,
                topicArn: arn
            )
        )

        guard let messageId = output.messageId else {
            print("No message ID received from Amazon SNS.")
            return
        }

        print("Published message with ID \(messageId)")
```

- For API details, see Publish in *AWS SDK for Swift API reference*.

Subscribe

The following code example shows how to use `Subscribe`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

Subscribe an email address to a topic.

```
import AWSSNS

        let config = try await SNSClient.SNSClientConfig(region: region)
        let snsClient = SNSClient(config: config)

        let output = try await snsClient.subscribe(
            input: SubscribeInput(
                endpoint: email,
                protocol: "email",
                returnSubscriptionArn: true,
                topicArn: arn
            )
        )

        guard let subscriptionArn = output.subscriptionArn else {
            print("No subscription ARN received from Amazon SNS.")
            return
        }

        print("Subscription \(subscriptionArn) created.")
```

Subscribe a phone number to a topic to receive notifications by SMS.

```
import AWSSNS

        let config = try await SNSClient.SNSClientConfig(region: region)
        let snsClient = SNSClient(config: config)

        let output = try await snsClient.subscribe(
            input: SubscribeInput(
                endpoint: phone,
```

```
                protocol: "sms",
                returnSubscriptionArn: true,
                topicArn: arn
            )
        )

        guard let subscriptionArn = output.subscriptionArn else {
            print("No subscription ARN received from Amazon SNS.")
            return
        }

        print("Subscription \(subscriptionArn) created.")
```

- For API details, see Subscribe in *AWS SDK for Swift API reference*.

`Unsubscribe`

The following code example shows how to use `Unsubscribe`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSNS

        let config = try await SNSClient.SNSClientConfig(region: region)
        let snsClient = SNSClient(config: config)

        _ = try await snsClient.unsubscribe(
            input: UnsubscribeInput(
                subscriptionArn: arn
            )
        )

        print("Unsubscribed.")
```

- For API details, see Unsubscribe in *AWS SDK for Swift API reference*.

Scenarios

Publish messages to queues

The following code example shows how to:

- Create topic (FIFO or non-FIFO).
- Subscribe several queues to the topic with an option to apply a filter.
- Publish messages to the topic.
- Poll the queues for messages received.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import ArgumentParser
import AWSClientRuntime
import AWSSNS
import AWSSQS
import Foundation

struct ExampleCommand: ParsableCommand {
    @Option(help: "Name of the Amazon Region to use")
    var region = "us-east-1"

    static var configuration = CommandConfiguration(
        commandName: "queue-scenario",
        abstract: """
        This example interactively demonstrates how to use Amazon Simple
        Notification Service (Amazon SNS) and Amazon Simple Queue Service
        (Amazon SQS) together to publish and receive messages using queues.
        """,
        discussion: """
```

```
        Supports filtering using a "tone" attribute.
        """
    )

    /// Prompt for an input string. Only non-empty strings are allowed.
    ///
    /// - Parameter prompt: The prompt to display.
    ///
    /// - Returns: The string input by the user.
    func stringRequest(prompt: String) -> String {
        var str: String?

        while str == nil {
            print(prompt, terminator: "")
            str = readLine()

            if str != nil && str?.count == 0 {
                str = nil
            }
        }

        return str!
    }

    /// Ask a yes/no question.
    ///
    /// - Parameter prompt: A prompt string to print.
    ///
    /// - Returns: `true` if the user answered "Y", otherwise `false`.
    func yesNoRequest(prompt: String) -> Bool {
        while true {
            let answer = stringRequest(prompt: prompt).lowercased()
            if answer == "y" || answer == "n" {
                return answer == "y"
            }
        }
    }

    /// Display a menu of options then request a selection.
    ///
    /// - Parameters:
    ///   - prompt: A prompt string to display before the menu.
    ///   - options: An array of strings giving the menu options.
    ///
```

```
    /// - Returns: The index number of the selected option or 0 if no item was
    ///   selected.
    func menuRequest(prompt: String, options: [String]) -> Int {
        let numOptions = options.count

        if numOptions == 0 {
            return 0
        }

        print(prompt)

        for (index, value) in options.enumerated() {
            print("(\(index)) \(value)")
        }

        repeat {
            print("Enter your selection (0 - \(numOptions-1)): ", terminator: "")
            if let answer = readLine() {
                guard let answer = Int(answer) else {
                    print("Please enter the number matching your selection.")
                    continue
                }

                if answer >= 0 && answer < numOptions {
                    return answer
                } else {
                    print("Please enter the number matching your selection.")
                }
            }
        } while true
    }

    /// Ask the user too press RETURN. Accepts any input but ignores it.
    ///
    /// - Parameter prompt: The text prompt to display.
    func returnRequest(prompt: String) {
        print(prompt, terminator: "")
        _ = readLine()
    }

    var attrValues = [
        "<none>",
        "cheerful",
        "funny",
```

```
        "serious",
        "sincere"
    ]

    /// Ask the user to choose one of the attribute values to use as a filter.
    ///
    /// - Parameters:
    ///   - message: A message to display before the menu of values.
    ///   - attrValues: An array of strings giving the values to choose from.
    ///
    /// - Returns: The string corresponding to the selected option.
    func askForFilter(message: String, attrValues: [String]) -> String? {
        print(message)
        for (index, value) in attrValues.enumerated() {
            print("  [\(index)] \(value)")
        }

        var answer: Int?
        repeat {
            answer = Int(stringRequest(prompt: "Select an value for the 'tone'
attribute or 0 to end: "))
        } while answer == nil || answer! < 0 || answer! > attrValues.count + 1

        if answer == 0 {
            return nil
        }
        return attrValues[answer!]
    }

    /// Prompts the user for filter terms and constructs the attribute
    /// record that specifies them.
    ///
    /// - Returns: A mapping of "FilterPolicy" to a JSON string representing
    ///   the user-defined filter.
    func buildFilterAttributes() -> [String:String] {
        var attr: [String:String] = [:]
        var filterString = ""

        var first = true

        while let ans = askForFilter(message: "Choose a value to apply to the 'tone'
attribute.",
                                     attrValues: attrValues) {
            if !first {
```

```
                filterString += ","
            }
            first = false

            filterString += "\"\(ans)\""
        }

        let filterJSON = "{ \"tone\": [\(filterString)]}"
        attr["FilterPolicy"] = filterJSON

        return attr
    }
    /// Create a queue, returning its URL string.
    ///
    /// - Parameters:
    ///   - prompt: A prompt to ask for the queue name.
    ///   - isFIFO: Whether or not to create a FIFO queue.
    ///
    /// - Returns: The URL of the queue.
    func createQueue(prompt: String, sqsClient: SQSClient, isFIFO: Bool) async
throws -> String? {
        repeat {
            var queueName = stringRequest(prompt: prompt)
            var attributes: [String: String] = [:]

            if isFIFO {
                queueName += ".fifo"
                attributes["FifoQueue"] = "true"
            }

            do {
                let output = try await sqsClient.createQueue(
                    input: CreateQueueInput(
                        attributes: attributes,
                        queueName: queueName
                    )
                )
                guard let url = output.queueUrl else {
                    return nil
                }

                return url
            } catch _ as QueueDeletedRecently {
```

```
                print("You need to use a different queue name. A queue by that name
was recently deleted.")
                continue
            }
        } while true
    }

    /// Return the ARN of a queue given its URL.
    ///
    /// - Parameter queueUrl: The URL of the queue for which to return the
    ///   ARN.
    ///
    /// - Returns: The ARN of the specified queue.
    func getQueueARN(sqsClient: SQSClient, queueUrl: String) async throws -> String?
{
        let output = try await sqsClient.getQueueAttributes(
            input: GetQueueAttributesInput(
                attributeNames: [.queuearn],
                queueUrl: queueUrl
            )
        )

        guard let attributes = output.attributes else {
            return nil
        }

        return attributes["QueueArn"]
    }

    /// Applies the needed policy to the specified queue.
    ///
    /// - Parameters:
    ///   - sqsClient: The Amazon SQS client to use.
    ///   - queueUrl: The queue to apply the policy to.
    ///   - queueArn: The ARN of the queue to apply the policy to.
    ///   - topicArn: The topic that should have access via the policy.
    ///
    /// - Throws: Errors from the SQS `SetQueueAttributes` action.
    func setQueuePolicy(sqsClient: SQSClient, queueUrl: String,
                        queueArn: String, topicArn: String) async throws {
        _ = try await sqsClient.setQueueAttributes(
            input: SetQueueAttributesInput(
                attributes: [
                    "Policy":
```

```
                        """
                        {
                            "Statement": [
                                {
                                    "Effect": "Allow",
                                    "Principal": {
                                        "Service": "sns.amazonaws.com"
                                    },
                                    "Action": "sqs:SendMessage",
                                    "Resource": "\(queueArn)",
                                    "Condition": {
                                        "ArnEquals": {
                                            "aws:SourceArn": "\(topicArn)"
                                        }
                                    }
                                }
                            ]
                        }
                        """

                ],
                queueUrl: queueUrl
            )
        )
    }

    /// Receive the available messages on a queue, outputting them to the
    /// screen. Returns a dictionary you pass to DeleteMessageBatch to delete
    /// all the received messages.
    ///
    /// - Parameters:
    ///   - sqsClient: The Amazon SQS client to use.
    ///   - queueUrl: The SQS queue on which to receive messages.
    ///
    /// - Throws: Errors from `SQSClient.receiveMessage()`
    ///
    /// - Returns: An array of SQSClientTypes.DeleteMessageBatchRequestEntry
    ///   items, each describing one received message in the format needed to
    ///   delete it.
    func receiveAndListMessages(sqsClient: SQSClient, queueUrl: String) async throws
                                -> [SQSClientTypes.DeleteMessageBatchRequestEntry] {
        let output = try await sqsClient.receiveMessage(
            input: ReceiveMessageInput(
                maxNumberOfMessages: 10,
```

```
                queueUrl: queueUrl
            )
        )

        guard let messages = output.messages else {
            print("No messages received.")
            return []
        }

        var deleteList: [SQSClientTypes.DeleteMessageBatchRequestEntry] = []

        // Print out all the messages that were received, including their
        // attributes, if any.

        for message in messages {
            print("Message ID:     \(message.messageId ?? "<unknown>")")
            print("Receipt handle: \(message.receiptHandle ?? "<unknown>")")
            print("Message JSON:   \(message.body ?? "<body missing>")")

            if message.receiptHandle != nil {
                deleteList.append(
                    SQSClientTypes.DeleteMessageBatchRequestEntry(
                        id: message.messageId,
                        receiptHandle: message.receiptHandle
                    )
                )
            }
        }

        return deleteList
    }

    /// Delete all the messages in the specified list.
    ///
    /// - Parameters:
    ///   - sqsClient: The Amazon SQS client to use.
    ///   - queueUrl: The SQS queue to delete messages from.
    ///   - deleteList: A list of `DeleteMessageBatchRequestEntry` objects
    ///     describing the messages to delete.
    ///
    /// - Throws: Errors from `SQSClient.deleteMessageBatch()`.
    func deleteMessageList(sqsClient: SQSClient, queueUrl: String,
                           deleteList:
[SQSClientTypes.DeleteMessageBatchRequestEntry]) async throws {
```

```
        let output = try await sqsClient.deleteMessageBatch(
            input: DeleteMessageBatchInput(entries: deleteList, queueUrl: queueUrl)
        )

        if let failed = output.failed {
            print("\(failed.count) errors occurred deleting messages from the
queue.")
            for message in failed {
                print("---> Failed to delete message \(message.id ?? "<unknown ID>")
with error: \(message.code ?? "<unknown>") (\(message.message ?? "..."))")
            }
        }
    }

    /// Called by ``main()`` to run the bulk of the example.
    func runAsync() async throws {
        let rowOfStars = String(repeating: "*", count: 75)

        print("""
              \(rowOfStars)
              Welcome to the cross-service messaging with topics and queues example.
              In this workflow, you'll create an SNS topic, then create two SQS
              queues which will be subscribed to that topic.

              You can specify several options for configuring the topic, as well as
              the queue subscriptions. You can then post messages to the topic and
              receive the results on the queues.
              \(rowOfStars)\n
              """
        )

        // 0. Create SNS and SQS clients.

        let snsConfig = try await SNSClient.SNSClientConfig(region: region)
        let snsClient = SNSClient(config: snsConfig)

        let sqsConfig = try await SQSClient.SQSClientConfig(region: region)
        let sqsClient = SQSClient(config: sqsConfig)

        // 1. Ask the user whether to create a FIFO topic. If so, ask whether
        //    to use content-based deduplication instead of requiring a
        //    deduplication ID.
```

```
        let isFIFO = yesNoRequest(prompt: "Do you want to create a FIFO topic (Y/N)?
")
        var isContentBasedDeduplication = false

        if isFIFO {
            print("""
                  \(rowOfStars)
                  Because you've chosen to create a FIFO topic, deduplication is
                  supported.

                  Deduplication IDs are either set in the message or are
automatically
                  generated from the content using a hash function.

                  If a message is successfully published to an SNS FIFO topic, any
                  message published and found to have the same deduplication ID
                  (within a five-minute deduplication interval), is accepted but
                  not delivered.

                  For more information about deduplication, see:
                  https://docs.aws.amazon.com/sns/latest/dg/fifo-message-dedup.html.
                  """
            )

            isContentBasedDeduplication = yesNoRequest(
                prompt: "Use content-based deduplication instead of entering a
deduplication ID (Y/N)? ")
            print(rowOfStars)
        }

        var topicName = stringRequest(prompt: "Enter the name of the topic to
create: ")

        // 2. Create the topic. Append ".fifo" to the name if FIFO was
        //    requested, and set the "FifoTopic" attribute to "true" if so as
        //    well. Set the "ContentBasedDeduplication" attribute to "true" if
        //    content-based deduplication was requested.

        if isFIFO {
            topicName += ".fifo"
        }

        print("Topic name: \(topicName)")
```

```
        var attributes = [
            "FifoTopic": (isFIFO ? "true" : "false")
        ]

        // If it's a FIFO topic with content-based deduplication, set the
        // "ContentBasedDeduplication" attribute.

        if isContentBasedDeduplication {
            attributes["ContentBasedDeduplication"] = "true"
        }

        // Create the topic and retrieve the ARN.

        let output = try await snsClient.createTopic(
            input: CreateTopicInput(
                attributes: attributes,
                name: topicName
            )
        )

        guard let topicArn = output.topicArn else {
            print("No topic ARN returned!")
            return
        }

        print(""""
              Topic '\(topicName) has been created with the
              topic ARN \(topicArn)."
              """
        )

        print(rowOfStars)

        // 3. Create an SQS queue. Append ".fifo" to the name if one of the
        //    FIFO topic configurations was chosen, and set "FifoQueue" to
        //    "true" if the topic is FIFO.

        print(""""
              Next, you will create two SQS queues that will be subscribed
              to the topic you just created.\n
              """
        )
```

```
        let q1Url = try await createQueue(prompt: "Enter the name of the first
queue: ",
                                          sqsClient: sqsClient, isFIFO: isFIFO)
        guard let q1Url else {
            print("Unable to create queue 1!")
            return
        }

        // 4. Get the SQS queue's ARN attribute using `GetQueueAttributes`.

        let q1Arn = try await getQueueARN(sqsClient: sqsClient, queueUrl: q1Url)

        guard let q1Arn else {
            print("Unable to get ARN of queue 1!")
            return
        }
        print("Got queue 1 ARN: \(q1Arn)")

        // 5. Attach an AWS IAM policy to the queue using
        //    `SetQueueAttributes`.

        try await setQueuePolicy(sqsClient: sqsClient, queueUrl: q1Url,
                                 queueArn: q1Arn, topicArn: topicArn)

        // 6. Subscribe the SQS queue to the SNS topic. Set the topic ARN in
        //    the request. Set the protocol to "sqs". Set the queue ARN to the
        //    ARN just received in step 5. For FIFO topics, give the option to
        //    apply a filter. A filter allows only matching messages to enter
        //    the queue.

        var q1Attributes: [String:String]? = nil

        if isFIFO {
            print(
                """

                If you add a filter to this subscription, then only the filtered
messages will
                be received in the queue. For information about message filtering,
see
                https://docs.aws.amazon.com/sns/latest/dg/sns-message-filtering.html
                For this example, you can filter messages by a 'tone' attribute.

                """
```

```
            )

            let subPrompt = """
                Would you like to filter messages for the first queue's subscription
to the
                topic \(topicName) (Y/N)?
                """
            if (yesNoRequest(prompt: subPrompt)) {
                q1Attributes = buildFilterAttributes()
            }
        }

        let sub1Output = try await snsClient.subscribe(
            input: SubscribeInput(
                attributes: q1Attributes,
                endpoint: q1Arn,
                protocol: "sqs",
                topicArn: topicArn
            )
        )

        guard let q1SubscriptionArn = sub1Output.subscriptionArn else {
            print("Invalid subscription ARN returned for queue 1!")
            return
        }

        // 7. Repeat steps 3-6 for the second queue.

        let q2Url = try await createQueue(prompt: "Enter the name of the second
queue: ",
                                  sqsClient: sqsClient, isFIFO: isFIFO)

        guard let q2Url else {
            print("Unable to create queue 2!")
            return
        }

        let q2Arn = try await getQueueARN(sqsClient: sqsClient, queueUrl: q2Url)

        guard let q2Arn else {
            print("Unable to get ARN of queue 2!")
            return
        }
        print("Got queue 2 ARN: \(q2Arn)")
```

```
        try await setQueuePolicy(sqsClient: sqsClient, queueUrl: q2Url,
                                 queueArn: q2Arn, topicArn: topicArn)

        var q2Attributes: [String:String]? = nil

        if isFIFO {
            let subPrompt = """
                Would you like to filter messages for the second queue's
subscription to the
                topic \(topicName) (Y/N)?
                """
            if (yesNoRequest(prompt: subPrompt)) {
                q2Attributes = buildFilterAttributes()
            }
        }

        let sub2Output = try await snsClient.subscribe(
            input: SubscribeInput(
                attributes: q2Attributes,
                endpoint: q2Arn,
                protocol: "sqs",
                topicArn: topicArn
            )
        )

        guard let q2SubscriptionArn = sub2Output.subscriptionArn else {
            print("Invalid subscription ARN returned for queue 1!")
            return
        }

        // 8. Let the user publish messages to the topic, asking for a message
        //    body for each message. Handle the types of topic correctly (SEE
        //    MVP INFORMATION AND FIX THESE COMMENTS!!!

        print("\n\(rowOfStars)\n")

        var first = true

        repeat {
            var publishInput = PublishInput(
                topicArn: topicArn
            )
```

```
            publishInput.message = stringRequest(prompt: "Enter message text to
publish: ")

            // If using a FIFO topic, a message group ID must be set on the
            // message.

            if isFIFO {
                if first {
                    print("""
                        Because you're using a FIFO topic, you must set a message
                        group ID. All messages within the same group will be
                        received in the same order in which they were published.\n
                        """
                    )
                }
                publishInput.messageGroupId = stringRequest(prompt: "Enter a message
group ID for this message: ")

                if !isContentBasedDeduplication {
                    if first {
                        print("""
                              Because you're not using content-based deduplication,
you
                              must enter a deduplication ID. If other messages with
the
                              same deduplication ID are published within the same
                              deduplication interval, they will not be delivered.
                              """
                        )
                    }
                    publishInput.messageDeduplicationId = stringRequest(prompt:
"Enter a deduplication ID for this message: ")
                }
            }

            // Allow the user to add a value for the "tone" attribute if they
            // wish to do so.

            var messageAttributes: [String:SNSClientTypes.MessageAttributeValue] =
[:]
            let attrValSelection = menuRequest(prompt: "Choose a tone to apply to
this message.", options: attrValues)

            if attrValSelection != 0 {
```

```
                let val = SNSClientTypes.MessageAttributeValue(dataType: "String",
stringValue: attrValues[attrValSelection])
                messageAttributes["tone"] = val
            }

            publishInput.messageAttributes = messageAttributes

            // Publish the message and display its ID.

            let publishOutput = try await snsClient.publish(input: publishInput)

            guard let messageID = publishOutput.messageId else {
                print("Unable to get the published message's ID!")
                return
            }

            print("Message published with ID \(messageID).")
            first = false

            // 9. Repeat step 8 until the user says they don't want to post
            //    another.

        } while (yesNoRequest(prompt: "Post another message (Y/N)? "))

        // 10. Display a list of the messages in each queue by using
        //     `ReceiveMessage`. Show at least the body and the attributes.

        print(rowOfStars)
        print("Contents of queue 1:")
        let q1DeleteList = try await receiveAndListMessages(sqsClient: sqsClient,
queueUrl: q1Url)
        print("\n\nContents of queue 2:")
        let q2DeleteList = try await receiveAndListMessages(sqsClient: sqsClient,
queueUrl: q2Url)
        print(rowOfStars)

        returnRequest(prompt: "\nPress return to clean up: ")

        // 11. Delete the received messages using `DeleteMessageBatch`.

        print("Deleting the messages from queue 1...")
        try await deleteMessageList(sqsClient: sqsClient, queueUrl: q1Url,
deleteList: q1DeleteList)
        print("\nDeleting the messages from queue 2...")
```

```
        try await deleteMessageList(sqsClient: sqsClient, queueUrl: q2Url,
 deleteList: q2DeleteList)

        // 12. Unsubscribe and delete both queues.

        print("\nUnsubscribing from queue 1...")
        _ = try await snsClient.unsubscribe(
            input: UnsubscribeInput(subscriptionArn: q1SubscriptionArn)
        )

        print("Unsubscribing from queue 2...")
        _ = try await snsClient.unsubscribe(
            input: UnsubscribeInput(subscriptionArn: q2SubscriptionArn)
        )

        print("Deleting queue 1...")
        _ = try await sqsClient.deleteQueue(
            input: DeleteQueueInput(queueUrl: q1Url)
        )

        print("Deleting queue 2...")
        _ = try await sqsClient.deleteQueue(
            input: DeleteQueueInput(queueUrl: q2Url)
        )

        // 13. Delete the topic.

        print("Deleting the SNS topic...")
        _ = try await snsClient.deleteTopic(
            input: DeleteTopicInput(topicArn: topicArn)
        )
    }
}

/// The program's asynchronous entry point.
@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.runAsync()
        } catch {
```

```
            ExampleCommand.exit(withError: error)
        }
    }
}
```

- For API details, see the following topics in *AWS SDK for Swift API reference*.
  - CreateQueue
  - CreateTopic
  - DeleteMessageBatch
  - DeleteQueue
  - DeleteTopic
  - GetQueueAttributes
  - Publish
  - ReceiveMessage
  - SetQueueAttributes
  - Subscribe
  - Unsubscribe

Amazon SQS examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with Amazon SQS.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Scenarios are code examples that show you how to accomplish specific tasks by calling multiple functions within a service or combined with other AWS services.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Get started
- Actions

- Scenarios

Get started

Hello Amazon SQS

The following code example shows how to get started using Amazon SQS.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

The `Package.swift` file.

```
import PackageDescription

let package = Package(
    name: "sqs-basics",
    // Let Xcode know the minimum Apple platforms supported.
    platforms: [
        .macOS(.v13),
        .iOS(.v15)
    ],
    dependencies: [
        // Dependencies declare other packages that this package depends on.
        .package(
            url: "https://github.com/awslabs/aws-sdk-swift",
            from: "1.0.0"),
        .package(
            url: "https://github.com/apple/swift-argument-parser.git",
            branch: "main"
        )
    ],
    targets: [
        // Targets are the basic building blocks of a package, defining a module or
 a test suite.
        // Targets can depend on other targets in this package and products
        // from dependencies.
```

```
        .executableTarget(
            name: "sqs-basics",
            dependencies: [
                .product(name: "AWSSQS", package: "aws-sdk-swift"),
                .product(name: "ArgumentParser", package: "swift-argument-parser")
            ],
            path: "Sources")

    ]
)
```

The Swift source code, `entry.swift`.

```
import ArgumentParser
import AWSClientRuntime
import AWSSQS
import Foundation

struct ExampleCommand: ParsableCommand {
    @Option(help: "Name of the Amazon Region to use (default: us-east-1)")
    var region = "us-east-1"

    static var configuration = CommandConfiguration(
        commandName: "sqs-basics",
        abstract: """
        This example shows how to list all of your available Amazon SQS queues.
        """,
        discussion: """
        """
    )

    /// Called by ``main()`` to run the bulk of the example.
    func runAsync() async throws {
        let config = try await SQSClient.SQSClientConfig(region: region)
        let sqsClient = SQSClient(config: config)

        var queues: [String] = []
        let outputPages = sqsClient.listQueuesPaginated(
            input: ListQueuesInput()
        )

        // Each time a page of results arrives, process its contents.
```

```
        for try await output in outputPages {
            guard let urls = output.queueUrls else {
                print("No queues found.")
                return
            }

            // Iterate over the queue URLs listed on this page, adding them
            // to the `queues` array.

            for queueUrl in urls {
                queues.append(queueUrl)
            }
        }

        print("You have \(queues.count) queues:")
        for queue in queues {
            print("   \(queue)")
        }
    }
}

/// The program's asynchronous entry point.
@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.runAsync()
        } catch {
            ExampleCommand.exit(withError: error)
        }
    }
}
```

- For API details, see ListQueues in *AWS SDK for Swift API reference*.

Actions

`CreateQueue`

The following code example shows how to use `CreateQueue`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSQS

        let config = try await SQSClient.SQSClientConfig(region: region)
        let sqsClient = SQSClient(config: config)

        let output = try await sqsClient.createQueue(
            input: CreateQueueInput(
                queueName: queueName
            )
        )

        guard let queueUrl = output.queueUrl else {
            print("No queue URL returned.")
            return
        }
```

- For API details, see CreateQueue in *AWS SDK for Swift API reference*.

`DeleteMessageBatch`

The following code example shows how to use `DeleteMessageBatch`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSQS

        let config = try await SQSClient.SQSClientConfig(region: region)
        let sqsClient = SQSClient(config: config)

        // Create the list of message entries.

        var entries: [SQSClientTypes.DeleteMessageBatchRequestEntry] = []
        var messageNumber = 1

        for handle in handles {
            let entry = SQSClientTypes.DeleteMessageBatchRequestEntry(
                id: "\(messageNumber)",
                receiptHandle: handle
            )
            entries.append(entry)
            messageNumber += 1
        }

        // Delete the messages.

        let output = try await sqsClient.deleteMessageBatch(
            input: DeleteMessageBatchInput(
                entries: entries,
                queueUrl: queue
            )
        )

        // Get the lists of failed and successful deletions from the output.

        guard let failedEntries = output.failed else {
            print("Failed deletion list is missing!")
            return
        }
```

```
        guard let successfulEntries = output.successful else {
            print("Successful deletion list is missing!")
            return
        }

        // Display a list of the failed deletions along with their
        // corresponding explanation messages.

        if failedEntries.count != 0 {
            print("Failed deletions:")

            for entry in failedEntries {
                print("Message #\(entry.id ?? "<unknown>") failed:
\(entry.message ?? "<unknown>")")
            }
        } else {
            print("No failed deletions.")
        }

        // Output a list of the message numbers that were successfully deleted.

        if successfulEntries.count != 0 {
            var successes = ""

            for entry in successfulEntries {
                if successes.count == 0 {
                    successes = entry.id ?? "<unknown>"
                } else {
                    successes = "\(successes), \(entry.id ?? "<unknown>")"
                }
            }
            print("Succeeded: ", successes)
        } else {
            print("No successful deletions.")
        }
```

- For API details, see DeleteMessageBatch in *AWS SDK for Swift API reference*.

`DeleteQueue`

The following code example shows how to use `DeleteQueue`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSQS

        let config = try await SQSClient.SQSClientConfig(region: region)
        let sqsClient = SQSClient(config: config)

        do {
            _ = try await sqsClient.deleteQueue(
                input: DeleteQueueInput(
                    queueUrl: queueUrl
                )
            )
        } catch _ as AWSSQS.QueueDoesNotExist {
            print("Error: The specified queue doesn't exist.")
            return
        }
```

- For API details, see DeleteQueue in *AWS SDK for Swift API reference*.

`GetQueueAttributes`

The following code example shows how to use `GetQueueAttributes`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSQS
```

```
        let config = try await SQSClient.SQSClientConfig(region: region)
        let sqsClient = SQSClient(config: config)

        let output = try await sqsClient.getQueueAttributes(
            input: GetQueueAttributesInput(
                attributeNames: [
                    .approximatenumberofmessages,
                    .maximummessagesize
                ],
                queueUrl: url
            )
        )

        guard let attributes = output.attributes else {
            print("No queue attributes returned.")
            return
        }

        for (attr, value) in attributes {
            switch(attr) {
            case "ApproximateNumberOfMessages":
                print("Approximate message count: \(value)")
            case "MaximumMessageSize":
                print("Maximum message size: \(value)kB")
            default:
                continue
            }
        }
```

- For API details, see GetQueueAttributes in *AWS SDK for Swift API reference*.

`ListQueues`

The following code example shows how to use `ListQueues`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSQS

        let config = try await SQSClient.SQSClientConfig(region: region)
        let sqsClient = SQSClient(config: config)

        var queues: [String] = []
        let outputPages = sqsClient.listQueuesPaginated(
            input: ListQueuesInput()
        )

        // Each time a page of results arrives, process its contents.

        for try await output in outputPages {
            guard let urls = output.queueUrls else {
                print("No queues found.")
                return
            }

            // Iterate over the queue URLs listed on this page, adding them
            // to the `queues` array.

            for queueUrl in urls {
                queues.append(queueUrl)
            }
        }
```

- For API details, see ListQueues in *AWS SDK for Swift API reference*.

`ReceiveMessage`

The following code example shows how to use `ReceiveMessage`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSQS

        let config = try await SQSClient.SQSClientConfig(region: region)
        let sqsClient = SQSClient(config: config)

        let output = try await sqsClient.receiveMessage(
            input: ReceiveMessageInput(
                maxNumberOfMessages: maxMessages,
                queueUrl: url
            )
        )

        guard let messages = output.messages else {
            print("No messages received.")
            return
        }

        for message in messages {
            print("Message ID:     \(message.messageId ?? "<unknown>")")
            print("Receipt handle: \(message.receiptHandle ?? "<unknown>")")
            print(message.body ?? "<body missing>")
            print("---")
        }
```

- For API details, see ReceiveMessage in *AWS SDK for Swift API reference*.

`SetQueueAttributes`

The following code example shows how to use `SetQueueAttributes`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSQS

        let config = try await SQSClient.SQSClientConfig(region: region)
        let sqsClient = SQSClient(config: config)

        do {
            _ = try await sqsClient.setQueueAttributes(
                input: SetQueueAttributesInput(
                    attributes: [
                        "MaximumMessageSize": "\(maxSize)"
                    ],
                    queueUrl: url
                )
            )
        } catch _ as AWSSQS.InvalidAttributeValue {
            print("Invalid maximum message size: \(maxSize) kB.")
        }
```

- For API details, see SetQueueAttributes in *AWS SDK for Swift API reference*.

Scenarios

Publish messages to queues

The following code example shows how to:

- Create topic (FIFO or non-FIFO).
- Subscribe several queues to the topic with an option to apply a filter.
- Publish messages to the topic.
- Poll the queues for messages received.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import ArgumentParser
import AWSClientRuntime
import AWSSNS
import AWSSQS
import Foundation

struct ExampleCommand: ParsableCommand {
    @Option(help: "Name of the Amazon Region to use")
    var region = "us-east-1"

    static var configuration = CommandConfiguration(
        commandName: "queue-scenario",
        abstract: """
        This example interactively demonstrates how to use Amazon Simple
        Notification Service (Amazon SNS) and Amazon Simple Queue Service
        (Amazon SQS) together to publish and receive messages using queues.
        """,
        discussion: """
        Supports filtering using a "tone" attribute.
        """
    )

    /// Prompt for an input string. Only non-empty strings are allowed.
    ///
    /// - Parameter prompt: The prompt to display.
    ///
    /// - Returns: The string input by the user.
    func stringRequest(prompt: String) -> String {
        var str: String?

        while str == nil {
            print(prompt, terminator: "")
            str = readLine()
```

```
            if str != nil && str?.count == 0 {
                str = nil
            }
        }

        return str!
    }

    /// Ask a yes/no question.
    ///
    /// - Parameter prompt: A prompt string to print.
    ///
    /// - Returns: `true` if the user answered "Y", otherwise `false`.
    func yesNoRequest(prompt: String) -> Bool {
        while true {
            let answer = stringRequest(prompt: prompt).lowercased()
            if answer == "y" || answer == "n" {
                return answer == "y"
            }
        }
    }

    /// Display a menu of options then request a selection.
    ///
    /// - Parameters:
    ///   - prompt: A prompt string to display before the menu.
    ///   - options: An array of strings giving the menu options.
    ///
    /// - Returns: The index number of the selected option or 0 if no item was
    ///   selected.
    func menuRequest(prompt: String, options: [String]) -> Int {
        let numOptions = options.count

        if numOptions == 0 {
            return 0
        }

        print(prompt)

        for (index, value) in options.enumerated() {
            print("(\(index)) \(value)")
        }

        repeat {
```

```
            print("Enter your selection (0 - \(numOptions-1)): ", terminator: "")
            if let answer = readLine() {
                guard let answer = Int(answer) else {
                    print("Please enter the number matching your selection.")
                    continue
                }

                if answer >= 0 && answer < numOptions {
                    return answer
                } else {
                    print("Please enter the number matching your selection.")
                }
            }
        } while true
    }

    /// Ask the user too press RETURN. Accepts any input but ignores it.
    ///
    /// - Parameter prompt: The text prompt to display.
    func returnRequest(prompt: String) {
        print(prompt, terminator: "")
        _ = readLine()
    }

    var attrValues = [
        "<none>",
        "cheerful",
        "funny",
        "serious",
        "sincere"
    ]

    /// Ask the user to choose one of the attribute values to use as a filter.
    ///
    /// - Parameters:
    ///   - message: A message to display before the menu of values.
    ///   - attrValues: An array of strings giving the values to choose from.
    ///
    /// - Returns: The string corresponding to the selected option.
    func askForFilter(message: String, attrValues: [String]) -> String? {
        print(message)
        for (index, value) in attrValues.enumerated() {
            print("  [\(index)] \(value)")
        }
```

```
        var answer: Int?
        repeat {
            answer = Int(stringRequest(prompt: "Select an value for the 'tone'
attribute or 0 to end: "))
        } while answer == nil || answer! < 0 || answer! > attrValues.count + 1

        if answer == 0 {
            return nil
        }
        return attrValues[answer!]
    }

    /// Prompts the user for filter terms and constructs the attribute
    /// record that specifies them.
    ///
    /// - Returns: A mapping of "FilterPolicy" to a JSON string representing
    ///   the user-defined filter.
    func buildFilterAttributes() -> [String:String] {
        var attr: [String:String] = [:]
        var filterString = ""

        var first = true

        while let ans = askForFilter(message: "Choose a value to apply to the 'tone'
attribute.",
                                     attrValues: attrValues) {
            if !first {
                filterString += ","
            }
            first = false

            filterString += "\"\(ans)\""
        }

        let filterJSON = "{ \"tone\": [\(filterString)]}"
        attr["FilterPolicy"] = filterJSON

        return attr
    }
    /// Create a queue, returning its URL string.
    ///
    /// - Parameters:
    ///   - prompt: A prompt to ask for the queue name.
```

```
   ///   - isFIFO: Whether or not to create a FIFO queue.
   ///
   /// - Returns: The URL of the queue.
   func createQueue(prompt: String, sqsClient: SQSClient, isFIFO: Bool) async
throws -> String? {
       repeat {
           var queueName = stringRequest(prompt: prompt)
           var attributes: [String: String] = [:]

           if isFIFO {
               queueName += ".fifo"
               attributes["FifoQueue"] = "true"
           }

           do {
               let output = try await sqsClient.createQueue(
                   input: CreateQueueInput(
                       attributes: attributes,
                       queueName: queueName
                   )
               )
               guard let url = output.queueUrl else {
                   return nil
               }

               return url
           } catch _ as QueueDeletedRecently {
               print("You need to use a different queue name. A queue by that name
was recently deleted.")
               continue
           }
       } while true
   }

   /// Return the ARN of a queue given its URL.
   ///
   /// - Parameter queueUrl: The URL of the queue for which to return the
   ///   ARN.
   ///
   /// - Returns: The ARN of the specified queue.
   func getQueueARN(sqsClient: SQSClient, queueUrl: String) async throws -> String?
{
       let output = try await sqsClient.getQueueAttributes(
           input: GetQueueAttributesInput(
```

```
                attributeNames: [.queuearn],
                queueUrl: queueUrl
            )
        )

        guard let attributes = output.attributes else {
            return nil
        }

        return attributes["QueueArn"]
    }

    /// Applies the needed policy to the specified queue.
    ///
    /// - Parameters:
    ///   - sqsClient: The Amazon SQS client to use.
    ///   - queueUrl: The queue to apply the policy to.
    ///   - queueArn: The ARN of the queue to apply the policy to.
    ///   - topicArn: The topic that should have access via the policy.
    ///
    /// - Throws: Errors from the SQS `SetQueueAttributes` action.
    func setQueuePolicy(sqsClient: SQSClient, queueUrl: String,
                        queueArn: String, topicArn: String) async throws {
        _ = try await sqsClient.setQueueAttributes(
            input: SetQueueAttributesInput(
                attributes: [
                    "Policy":
                        """
                        {
                            "Statement": [
                                {
                                    "Effect": "Allow",
                                    "Principal": {
                                        "Service": "sns.amazonaws.com"
                                    },
                                    "Action": "sqs:SendMessage",
                                    "Resource": "\(queueArn)",
                                    "Condition": {
                                        "ArnEquals": {
                                            "aws:SourceArn": "\(topicArn)"
                                        }
                                    }
                                }
                            ]
```

```
                        }
                        """

                ],
                queueUrl: queueUrl
            )
        )
    }

    /// Receive the available messages on a queue, outputting them to the
    /// screen. Returns a dictionary you pass to DeleteMessageBatch to delete
    /// all the received messages.
    ///
    /// - Parameters:
    ///   - sqsClient: The Amazon SQS client to use.
    ///   - queueUrl: The SQS queue on which to receive messages.
    ///
    /// - Throws: Errors from `SQSClient.receiveMessage()`
    ///
    /// - Returns: An array of SQSClientTypes.DeleteMessageBatchRequestEntry
    ///   items, each describing one received message in the format needed to
    ///   delete it.
    func receiveAndListMessages(sqsClient: SQSClient, queueUrl: String) async throws
                                -> [SQSClientTypes.DeleteMessageBatchRequestEntry] {
        let output = try await sqsClient.receiveMessage(
            input: ReceiveMessageInput(
                maxNumberOfMessages: 10,
                queueUrl: queueUrl
            )
        )

        guard let messages = output.messages else {
            print("No messages received.")
            return []
        }

        var deleteList: [SQSClientTypes.DeleteMessageBatchRequestEntry] = []

        // Print out all the messages that were received, including their
        // attributes, if any.

        for message in messages {
            print("Message ID:     \(message.messageId ?? "<unknown>")")
            print("Receipt handle: \(message.receiptHandle ?? "<unknown>")")
```

```
            print("Message JSON:   \(message.body ?? "<body missing>")")

            if message.receiptHandle != nil {
                deleteList.append(
                    SQSClientTypes.DeleteMessageBatchRequestEntry(
                        id: message.messageId,
                        receiptHandle: message.receiptHandle
                    )
                )
            }
        }

        return deleteList
    }

    /// Delete all the messages in the specified list.
    ///
    /// - Parameters:
    ///   - sqsClient: The Amazon SQS client to use.
    ///   - queueUrl: The SQS queue to delete messages from.
    ///   - deleteList: A list of `DeleteMessageBatchRequestEntry` objects
    ///     describing the messages to delete.
    ///
    /// - Throws: Errors from `SQSClient.deleteMessageBatch()`.
    func deleteMessageList(sqsClient: SQSClient, queueUrl: String,
                           deleteList:
[SQSClientTypes.DeleteMessageBatchRequestEntry]) async throws {
        let output = try await sqsClient.deleteMessageBatch(
            input: DeleteMessageBatchInput(entries: deleteList, queueUrl: queueUrl)
        )

        if let failed = output.failed {
            print("\(failed.count) errors occurred deleting messages from the
queue.")
            for message in failed {
                print("---> Failed to delete message \(message.id ?? "<unknown ID>")
with error: \(message.code ?? "<unknown>") (\(message.message ?? "..."))")
            }
        }
    }

    /// Called by ``main()`` to run the bulk of the example.
    func runAsync() async throws {
        let rowOfStars = String(repeating: "*", count: 75)
```

```
        print("""
              \(rowOfStars)
              Welcome to the cross-service messaging with topics and queues example.
              In this workflow, you'll create an SNS topic, then create two SQS
              queues which will be subscribed to that topic.

              You can specify several options for configuring the topic, as well as
              the queue subscriptions. You can then post messages to the topic and
              receive the results on the queues.
              \(rowOfStars)\n
              """
        )

        // 0. Create SNS and SQS clients.

        let snsConfig = try await SNSClient.SNSClientConfig(region: region)
        let snsClient = SNSClient(config: snsConfig)

        let sqsConfig = try await SQSClient.SQSClientConfig(region: region)
        let sqsClient = SQSClient(config: sqsConfig)

        // 1. Ask the user whether to create a FIFO topic. If so, ask whether
        //    to use content-based deduplication instead of requiring a
        //    deduplication ID.

        let isFIFO = yesNoRequest(prompt: "Do you want to create a FIFO topic (Y/N)?
")
        var isContentBasedDeduplication = false

        if isFIFO {
            print("""
                  \(rowOfStars)
                  Because you've chosen to create a FIFO topic, deduplication is
                  supported.

                  Deduplication IDs are either set in the message or are
automatically
                  generated from the content using a hash function.

                  If a message is successfully published to an SNS FIFO topic, any
                  message published and found to have the same deduplication ID
                  (within a five-minute deduplication interval), is accepted but
                  not delivered.
```

```
                 For more information about deduplication, see:
                 https://docs.aws.amazon.com/sns/latest/dg/fifo-message-dedup.html.
                 """
            )

            isContentBasedDeduplication = yesNoRequest(
                prompt: "Use content-based deduplication instead of entering a
deduplication ID (Y/N)? ")
            print(rowOfStars)
        }

        var topicName = stringRequest(prompt: "Enter the name of the topic to
create: ")

        // 2. Create the topic. Append ".fifo" to the name if FIFO was
        //    requested, and set the "FifoTopic" attribute to "true" if so as
        //    well. Set the "ContentBasedDeduplication" attribute to "true" if
        //    content-based deduplication was requested.

        if isFIFO {
            topicName += ".fifo"
        }

        print("Topic name: \(topicName)")

        var attributes = [
            "FifoTopic": (isFIFO ? "true" : "false")
        ]

        // If it's a FIFO topic with content-based deduplication, set the
        // "ContentBasedDeduplication" attribute.

        if isContentBasedDeduplication {
            attributes["ContentBasedDeduplication"] = "true"
        }

        // Create the topic and retrieve the ARN.

        let output = try await snsClient.createTopic(
            input: CreateTopicInput(
                attributes: attributes,
                name: topicName
            )
```

```
        )

        guard let topicArn = output.topicArn else {
            print("No topic ARN returned!")
            return
        }

        print("""
              Topic '\(topicName) has been created with the
              topic ARN \(topicArn)."
              """
        )

        print(rowOfStars)

        // 3. Create an SQS queue. Append ".fifo" to the name if one of the
        //    FIFO topic configurations was chosen, and set "FifoQueue" to
        //    "true" if the topic is FIFO.

        print("""
              Next, you will create two SQS queues that will be subscribed
              to the topic you just created.\n
              """
        )

        let q1Url = try await createQueue(prompt: "Enter the name of the first
queue: ",
                                          sqsClient: sqsClient, isFIFO: isFIFO)
        guard let q1Url else {
            print("Unable to create queue 1!")
            return
        }

        // 4. Get the SQS queue's ARN attribute using `GetQueueAttributes`.

        let q1Arn = try await getQueueARN(sqsClient: sqsClient, queueUrl: q1Url)

        guard let q1Arn else {
            print("Unable to get ARN of queue 1!")
            return
        }
        print("Got queue 1 ARN: \(q1Arn)")

        // 5. Attach an AWS IAM policy to the queue using
```

```
        //    `SetQueueAttributes`.

        try await setQueuePolicy(sqsClient: sqsClient, queueUrl: q1Url,
                                 queueArn: q1Arn, topicArn: topicArn)

        // 6. Subscribe the SQS queue to the SNS topic. Set the topic ARN in
        //    the request. Set the protocol to "sqs". Set the queue ARN to the
        //    ARN just received in step 5. For FIFO topics, give the option to
        //    apply a filter. A filter allows only matching messages to enter
        //    the queue.

        var q1Attributes: [String:String]? = nil

        if isFIFO {
            print(
                """

                If you add a filter to this subscription, then only the filtered
messages will
                be received in the queue. For information about message filtering,
see
                https://docs.aws.amazon.com/sns/latest/dg/sns-message-filtering.html
                For this example, you can filter messages by a 'tone' attribute.

                """
            )

            let subPrompt = """
                Would you like to filter messages for the first queue's subscription
to the
                topic \(topicName) (Y/N)?
                """
            if (yesNoRequest(prompt: subPrompt)) {
                q1Attributes = buildFilterAttributes()
            }
        }

        let sub1Output = try await snsClient.subscribe(
            input: SubscribeInput(
                attributes: q1Attributes,
                endpoint: q1Arn,
                protocol: "sqs",
                topicArn: topicArn
            )
```

```
        )

        guard let q1SubscriptionArn = sub1Output.subscriptionArn else {
            print("Invalid subscription ARN returned for queue 1!")
            return
        }

        // 7. Repeat steps 3-6 for the second queue.

        let q2Url = try await createQueue(prompt: "Enter the name of the second
queue: ",
                                  sqsClient: sqsClient, isFIFO: isFIFO)

        guard let q2Url else {
            print("Unable to create queue 2!")
            return
        }

        let q2Arn = try await getQueueARN(sqsClient: sqsClient, queueUrl: q2Url)

        guard let q2Arn else {
            print("Unable to get ARN of queue 2!")
            return
        }
        print("Got queue 2 ARN: \(q2Arn)")

        try await setQueuePolicy(sqsClient: sqsClient, queueUrl: q2Url,
                                 queueArn: q2Arn, topicArn: topicArn)

        var q2Attributes: [String:String]? = nil

        if isFIFO {
            let subPrompt = """
                Would you like to filter messages for the second queue's
subscription to the
                topic \(topicName) (Y/N)?
                """
            if (yesNoRequest(prompt: subPrompt)) {
                q2Attributes = buildFilterAttributes()
            }
        }

        let sub2Output = try await snsClient.subscribe(
            input: SubscribeInput(
```

```
                attributes: q2Attributes,
                endpoint: q2Arn,
                protocol: "sqs",
                topicArn: topicArn
            )
        )

        guard let q2SubscriptionArn = sub2Output.subscriptionArn else {
            print("Invalid subscription ARN returned for queue 1!")
            return
        }

        // 8. Let the user publish messages to the topic, asking for a message
        //    body for each message. Handle the types of topic correctly (SEE
        //    MVP INFORMATION AND FIX THESE COMMENTS!!!

        print("\n\(rowOfStars)\n")

        var first = true

        repeat {
            var publishInput = PublishInput(
                topicArn: topicArn
            )

            publishInput.message = stringRequest(prompt: "Enter message text to
publish: ")

            // If using a FIFO topic, a message group ID must be set on the
            // message.

            if isFIFO {
                if first {
                    print("""
                        Because you're using a FIFO topic, you must set a message
                        group ID. All messages within the same group will be
                        received in the same order in which they were published.\n
                        """
                    )
                }
                publishInput.messageGroupId = stringRequest(prompt: "Enter a message
group ID for this message: ")

                if !isContentBasedDeduplication {
```

```
                    if first {
                        print("""
                              Because you're not using content-based deduplication,
you
                              must enter a deduplication ID. If other messages with
the
                              same deduplication ID are published within the same
                              deduplication interval, they will not be delivered.
                              """
                        )
                    }
                    publishInput.messageDeduplicationId = stringRequest(prompt:
"Enter a deduplication ID for this message: ")
                }
            }

            // Allow the user to add a value for the "tone" attribute if they
            // wish to do so.

            var messageAttributes: [String:SNSClientTypes.MessageAttributeValue] =
[:]
            let attrValSelection = menuRequest(prompt: "Choose a tone to apply to
this message.", options: attrValues)

            if attrValSelection != 0 {
                let val = SNSClientTypes.MessageAttributeValue(dataType: "String",
stringValue: attrValues[attrValSelection])
                messageAttributes["tone"] = val
            }

            publishInput.messageAttributes = messageAttributes

            // Publish the message and display its ID.

            let publishOutput = try await snsClient.publish(input: publishInput)

            guard let messageID = publishOutput.messageId else {
                print("Unable to get the published message's ID!")
                return
            }

            print("Message published with ID \(messageID).")
            first = false
```

```
            // 9. Repeat step 8 until the user says they don't want to post
            //    another.

        } while (yesNoRequest(prompt: "Post another message (Y/N)? "))

        // 10. Display a list of the messages in each queue by using
        //     `ReceiveMessage`. Show at least the body and the attributes.

        print(rowOfStars)
        print("Contents of queue 1:")
        let q1DeleteList = try await receiveAndListMessages(sqsClient: sqsClient,
queueUrl: q1Url)
        print("\n\nContents of queue 2:")
        let q2DeleteList = try await receiveAndListMessages(sqsClient: sqsClient,
queueUrl: q2Url)
        print(rowOfStars)

        returnRequest(prompt: "\nPress return to clean up: ")

        // 11. Delete the received messages using `DeleteMessageBatch`.

        print("Deleting the messages from queue 1...")
        try await deleteMessageList(sqsClient: sqsClient, queueUrl: q1Url,
deleteList: q1DeleteList)
        print("\nDeleting the messages from queue 2...")
        try await deleteMessageList(sqsClient: sqsClient, queueUrl: q2Url,
deleteList: q2DeleteList)

        // 12. Unsubscribe and delete both queues.

        print("\nUnsubscribing from queue 1...")
        _ = try await snsClient.unsubscribe(
            input: UnsubscribeInput(subscriptionArn: q1SubscriptionArn)
        )

        print("Unsubscribing from queue 2...")
        _ = try await snsClient.unsubscribe(
            input: UnsubscribeInput(subscriptionArn: q2SubscriptionArn)
        )

        print("Deleting queue 1...")
        _ = try await sqsClient.deleteQueue(
            input: DeleteQueueInput(queueUrl: q1Url)
        )
```

```
        print("Deleting queue 2...")
        _ = try await sqsClient.deleteQueue(
            input: DeleteQueueInput(queueUrl: q2Url)
        )

        // 13. Delete the topic.

        print("Deleting the SNS topic...")
        _ = try await snsClient.deleteTopic(
            input: DeleteTopicInput(topicArn: topicArn)
        )
    }
}

/// The program's asynchronous entry point.
@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.runAsync()
        } catch {
            ExampleCommand.exit(withError: error)
        }
    }
}
```

- For API details, see the following topics in *AWS SDK for Swift API reference*.
  - CreateQueue
  - CreateTopic
  - DeleteMessageBatch
  - DeleteQueue
  - DeleteTopic
  - GetQueueAttributes
  - Publish
  - ReceiveMessage

- SetQueueAttributes
- Subscribe
- Unsubscribe

AWS STS examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with AWS STS.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Actions

Actions

`AssumeRole`

The following code example shows how to use `AssumeRole`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
import AWSSTS

    public func assumeRole(role: IAMClientTypes.Role, sessionName: String)
        async throws -> STSClientTypes.Credentials
    {
```

```
        let input = AssumeRoleInput(
            roleArn: role.arn,
            roleSessionName: sessionName
        )
        do {
            let output = try await stsClient.assumeRole(input: input)

            guard let credentials = output.credentials else {
                throw ServiceHandlerError.authError
            }

            return credentials
        } catch {
            print("Error assuming role: ", dump(error))
            throw error
        }
    }
```

- For API details, see AssumeRole in *AWS SDK for Swift API reference*.

Amazon Transcribe Streaming examples using SDK for Swift

The following code examples show you how to perform actions and implement common scenarios by using the AWS SDK for Swift with Amazon Transcribe Streaming.

Actions are code excerpts from larger programs and must be run in context. While actions show you how to call individual service functions, you can see actions in context in their related scenarios.

Scenarios are code examples that show you how to accomplish specific tasks by calling multiple functions within a service or combined with other AWS services.

Each example includes a link to the complete source code, where you can find instructions on how to set up and run the code in context.

Topics

- Actions
- Scenarios

Actions

`StartStreamTranscription`

The following code example shows how to use `StartStreamTranscription`.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

```
        let client = TranscribeStreamingClient(
            config: try await
TranscribeStreamingClient.TranscribeStreamingClientConfig(
                region: region
            )
        )

        // Start the transcription running on the audio stream.

        let output = try await client.startStreamTranscription(
            input: StartStreamTranscriptionInput(
                audioStream: try await createAudioStream(),
                languageCode: TranscribeStreamingClientTypes.LanguageCode(rawValue:
lang),
                mediaEncoding: encoding,
                mediaSampleRateHertz: sampleRate
            )
        )
```

- For API details, see StartStreamTranscription in *AWS SDK for Swift API reference*.

Scenarios

Transcribe an audio file

The following code example shows how to generate a transcription of a source audio file using Amazon Transcribe streaming.

SDK for Swift

> **Note**
>
> There's more on GitHub. Find the complete example and learn how to set up and run in the AWS Code Examples Repository.

Use Amazon Transcribe streaming to transcribe the spoken language in an audio file.

```
/// An example that demonstrates how to watch an transcribe event stream to
/// transcribe audio from a file to the console.

import ArgumentParser
import AWSClientRuntime
import AWSTranscribeStreaming
import Foundation

/// Identify one of the media file formats supported by Amazon Transcribe.
enum TranscribeFormat: String, ExpressibleByArgument {
    case ogg = "ogg"
    case pcm = "pcm"
    case flac = "flac"
}

// -MARK: - Async command line tool

struct ExampleCommand: ParsableCommand {
    // -MARK: Command arguments
    @Flag(help: "Show partial results")
    var showPartial = false
    @Option(help: "Language code to transcribe into")
    var lang: String = "en-US"
    @Option(help: "Format of the source audio file")
    var format: TranscribeFormat
```

```
    @Option(help: "Sample rate of the source audio file in Hertz")
    var sampleRate: Int = 16000
    @Option(help: "Path of the source audio file")
    var path: String
    @Option(help: "Name of the Amazon S3 Region to use (default: us-east-1)")
    var region = "us-east-1"

    static var configuration = CommandConfiguration(
        commandName: "tsevents",
        abstract: """
        This example shows how to use event streaming with Amazon Transcribe.
        """,
        discussion: """
        """
    )

    /// Create and return an Amazon Transcribe audio stream from the file
    /// specified in the arguments.
    ///
    /// - Throws: Errors from `TranscribeError`.
    ///
    /// - Returns: `AsyncThrowingStream<TranscribeStreamingClientTypes.AudioStream,
Error>`
    func createAudioStream() async throws
                -> AsyncThrowingStream<TranscribeStreamingClientTypes.AudioStream,
Error> {

        let fileURL: URL = URL(fileURLWithPath: path)
        let audioData = try Data(contentsOf: fileURL)

        // Properties defining the size of audio chunks and the total size of
        // the audio file in bytes. You should try to send chunks that last on
        // average 125 milliseconds.

        let chunkSizeInMilliseconds = 125.0
        let chunkSize = Int(chunkSizeInMilliseconds  / 1000.0 * Double(sampleRate) *
2.0)
        let audioDataSize = audioData.count

        // Create an audio stream from the source data. The stream's job is
        // to send the audio in chunks to Amazon Transcribe as
        // `AudioStream.audioevent` events.
```

```
        let audioStream =
AsyncThrowingStream<TranscribeStreamingClientTypes.AudioStream,
                                Error> { continuation in
            Task {
                var currentStart = 0
                var currentEnd = min(chunkSize, audioDataSize - currentStart)

                // Generate and send chunks of audio data as `audioevent`
                // events until the entire file has been sent. Each event is
                // yielded to the SDK after being created.

                while currentStart < audioDataSize {
                    let dataChunk = audioData[currentStart ..< currentEnd]

                    let audioEvent =
TranscribeStreamingClientTypes.AudioStream.audioevent(
                        .init(audioChunk: dataChunk)
                    )
                    let yieldResult = continuation.yield(audioEvent)
                    switch yieldResult {
                        case .enqueued(_):
                            // The chunk was successfully enqueued into the
                            // stream. The `remaining` parameter estimates how
                            // much room is left in the queue, but is ignored here.
                            break
                        case .dropped(_):
                            // The chunk was dropped because the queue buffer
                            // is full. This will cause transcription errors.
                            print("Warning: Dropped audio! The transcription will be
incomplete.")
                        case .terminated:
                            print("Audio stream terminated.")
                            continuation.finish()
                            return
                        default:
                            print("Warning: Unrecognized response during audio
streaming.")
                    }

                    currentStart = currentEnd
                    currentEnd = min(currentStart + chunkSize, audioDataSize)
                }

                // Let the SDK's continuation block know the stream is over.
```

```
                continuation.finish()
            }
        }

        return audioStream
    }

    /// Run the transcription process.
    ///
    /// - Throws: An error from `TranscribeError`.
    func transcribe(encoding: TranscribeStreamingClientTypes.MediaEncoding) async
throws {
        // Create the Transcribe Streaming client.

        let client = TranscribeStreamingClient(
            config: try await
TranscribeStreamingClient.TranscribeStreamingClientConfig(
                region: region
            )
        )

        // Start the transcription running on the audio stream.

        let output = try await client.startStreamTranscription(
            input: StartStreamTranscriptionInput(
                audioStream: try await createAudioStream(),
                languageCode: TranscribeStreamingClientTypes.LanguageCode(rawValue:
lang),
                mediaEncoding: encoding,
                mediaSampleRateHertz: sampleRate
            )
        )

        // Iterate over the events in the returned transcript result stream.
        // Each `transcriptevent` contains a list of result fragments which
        // need to be concatenated together to build the final transcript.
        for try await event in output.transcriptResultStream! {
            switch event {
            case .transcriptevent(let event):
            for result in event.transcript?.results ?? [] {
                guard let transcript = result.alternatives?.first?.transcript else {
                    continue
                }
```

```
                // If showing partial results is enabled and the result is
                // partial, show it. Partial results may be incomplete, and
                // may be inaccurate, with upcoming audio making the
                // transcription complete or by giving more context to make
                // transcription make more sense.

                if (result.isPartial && showPartial) {
                    print("[Partial] \(transcript)")
                }

                // When the complete fragment of transcribed text is ready,
                // print it. This could just as easily be used to draw the
                // text as a subtitle over a playing video, though timing
                // would need to be managed.

                if !result.isPartial {
                    if (showPartial) {
                        print("[Final  ] ", terminator: "")
                    }
                    print(transcript)
                }
            }
            default:
                print("Error: Unexpected message from Amazon Transcribe:")
            }
        }
    }

    /// Convert the value of the `--format` command line option into the
    /// corresponding Transcribe Streaming `MediaEncoding` type.
    ///
    /// - Returns: The `MediaEncoding` equivalent of the format specified on
    ///   the command line.
    func getMediaEncoding() -> TranscribeStreamingClientTypes.MediaEncoding {
        let mediaEncoding: TranscribeStreamingClientTypes.MediaEncoding

        switch format {
        case .flac:
            mediaEncoding = .flac
        case .ogg:
            mediaEncoding = .oggOpus
        case .pcm:
            mediaEncoding = .pcm
```

```
        }

        return mediaEncoding
    }
}

// -MARK: - Entry point

/// The program's asynchronous entry point.
@main
struct Main {
    static func main() async {
        let args = Array(CommandLine.arguments.dropFirst())

        do {
            let command = try ExampleCommand.parse(args)
            try await command.transcribe(encoding: command.getMediaEncoding())
        } catch let error as TranscribeError {
            print("ERROR: \(error.errorDescription ?? "Unknown error")")
        } catch {
            ExampleCommand.exit(withError: error)
        }
    }
}

/// Errors thrown by the example's functions.
enum TranscribeError: Error {
    /// No transcription stream available.
    case noTranscriptionStream
    /// The source media file couldn't be read.
    case readError

    var errorDescription: String? {
        switch self {
        case .noTranscriptionStream:
            return "No transcription stream returned by Amazon Transcribe."
        case .readError:
            return "Unable to read the source audio file."
        }
    }
}
```

- For API details, see StartStreamTranscription in *AWS SDK for Swift API reference*.

Security in AWS SDK for Swift

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the .
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using SDK for Swift. The following topics show you how to configure SDK for Swift to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your SDK for Swift resources.

Topics

- Data protection in the AWS SDK for Swift
- Identity and Access Management
- Compliance Validation for this AWS Product or Service
- Resilience for this AWS Product or Service
- Infrastructure Security for this AWS Product or Service

Data protection in the AWS SDK for Swift

The AWS shared responsibility model applies to data protection in AWS SDK for Swift. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see Data Privacy FAQ. For

information about data protection in Europe, see the General Data Protection Regulation (GDPR) Center.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see Working with CloudTrail trails in the *AWS CloudTrail User Guide*.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see Federal Information Processing Standard (FIPS) 140-3.

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with SDK for Swift or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Encryption in transit

The SDK for Swift uses the TLS implementation of the platform that your code runs on:

- On Apple platforms, the SDK sends requests using `URLSession`.
- On Linux, the SDK uses the AWS Common Runtime (AWS CRT).

The TLS versions available to your application depend on that platform's TLS support. To use a specific TLS version, make sure that the platform your application runs on supports it.

TLS 1.3 is a prerequisite for enabling post-quantum cryptography, which might require additional actions or configurations. For information about post-quantum TLS support in the AWS SDKs, see Enabling hybrid post-quantum TLS in the *AWS SDKs and Tools Reference Guide*.

Identity and Access Management

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

Topics

- Audience
- Authenticating with identities
- Managing access using policies
- How AWS services work with IAM
- Troubleshooting AWS identity and access

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in AWS.

Service user – If you use AWS services to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more AWS features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in AWS, see Troubleshooting AWS identity and access or the user guide of the AWS service you are using.

Service administrator – If you're in charge of AWS resources at your company, you probably have full access to AWS. It's your job to determine which AWS features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with AWS, see the user guide of the AWS service you are using.

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to AWS. To view example AWS identity-based policies that you can use in IAM, see the user guide of the AWS service you are using.

Authenticating with identities

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

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

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

AWS account root user

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

Federated identity

As a best practice, require human users to use federation with an identity provider to access AWS services using temporary credentials.

A *federated identity* is a user from your enterprise directory, web identity provider, or Directory Service that accesses AWS services using credentials from an identity source. Federated identities assume roles that provide temporary credentials.

For centralized access management, we recommend AWS IAM Identity Center. For more information, see What is IAM Identity Center? in the *AWS IAM Identity Center User Guide*.

IAM users and groups

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

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

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

IAM roles

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

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

Managing access using policies

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

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

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

Identity-based policies

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

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

managed and inline policies, see Choose between managed policies and inline policies in the *IAM User Guide*.

Resource-based policies

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

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

Access control lists (ACLs)

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

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

Other policy types

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

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

Multiple policy types

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

How AWS services work with IAM

To get a high-level view of how AWS services work with most IAM features, see AWS services that work with IAM in the *IAM User Guide*.

To learn how to use a specific AWS service with IAM, see the security section of the relevant service's User Guide.

Troubleshooting AWS identity and access

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

Topics

- I am not authorized to perform an action in AWS
- I am not authorized to perform iam:PassRole
- I want to allow people outside of my AWS account to access my AWS resources

I am not authorized to perform an action in AWS

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the `mateojackson` IAM user tries to use the console to view details about a fictional *`my-example-widget`* resource but doesn't have the fictional `awes:`*`GetWidget`* permissions.

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

In this case, the policy for the `mateojackson` user must be updated to allow access to the *`my-example-widget`* resource by using the `awes:`*`GetWidget`* action.

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

I am not authorized to perform iam:PassRole

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

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

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

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

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

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

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

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

To learn more, consult the following:

- To learn whether AWS supports these features, see How AWS services work with IAM.
- To learn how to provide access to your resources across AWS accounts that you own, see Providing access to an IAM user in another AWS account that you own in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see Providing access to AWS accounts owned by third parties in the *IAM User Guide*.
- To learn how to provide access through identity federation, see Providing access to externally authenticated users (identity federation) in the *IAM User Guide*.

- To learn the difference between using roles and resource-based policies for cross-account access, see Cross account resource access in IAM in the *IAM User Guide*.

Compliance Validation for this AWS Product or Service

To learn whether an AWS service is within the scope of specific compliance programs, see AWS services in Scope by Compliance Program and choose the compliance program that you are interested in. For general information, see AWS Compliance Programs.

You can download third-party audit reports using AWS Artifact. For more information, see Downloading Reports in AWS Artifact.

Your compliance responsibility when using AWS services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. For more information about your compliance responsibility when using AWS services, see AWS Security Documentation.

This AWS product or service follows the shared responsibility model through the specific Amazon Web Services (AWS) services it supports. For AWS service security information, see the AWS service security documentation page and AWS services that are in scope of AWS compliance efforts by compliance program.

Resilience for this AWS Product or Service

The AWS global infrastructure is built around AWS Regions and Availability Zones.

AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking.

With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure.

This AWS product or service follows the shared responsibility model through the specific Amazon Web Services (AWS) services it supports. For AWS service security information, see the AWS service security documentation page and AWS services that are in scope of AWS compliance efforts by compliance program.

Infrastructure Security for this AWS Product or Service

This AWS product or service uses managed services, and therefore is protected by the AWS global network security. For information about AWS security services and how AWS protects infrastructure, see AWS Cloud Security. To design your AWS environment using the best practices for infrastructure security, see Infrastructure Protection in *Security Pillar AWS Well-Architected Framework*.

You use AWS published API calls to access this AWS Product or Service through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the AWS Security Token Service (AWS STS) to generate temporary security credentials to sign requests.

This AWS product or service follows the shared responsibility model through the specific Amazon Web Services (AWS) services it supports. For AWS service security information, see the AWS service security documentation page and AWS services that are in scope of AWS compliance efforts by compliance program.

Document history for the AWS SDK for Swift Developer Guide

Updates about significant changes or changes of interest to the AWS SDK for Swift Developer Guide.

Change	Description	Date
Client configuration rename	Replaced references to the deprecated client configuration classes, such as `S3Client.S3ClientConfiguration`, with the `Config` structures that supersede them, such as `S3Client.S3ClientConfig`.	October 9, 2026
Updated tool version requirements	Updated the minimum version requirements to Swift 6.0 and Xcode 16.	September 25, 2026
Table of contents reorganization	Reorganized the table of contents to more closely follow the standard used by other AWS SDK guides.	July 30, 2025
HTTP client configuration	Added documentation about configuring HTTP clients.	May 22, 2025
Versioning policy	Added information about the versioning policy and revised the content hierarchy.	February 14, 2025
S3 checksums	Added a section on using checksums with Amazon S3.	February 11, 2025

Event streaming	Added a section on using event streaming, using Amazon Transcribe as an example.	December 26, 2024
Binary streaming	Added a section on using binary streaming for both uploads and downloads.	December 13, 2024
Presigned URLs and multipart uploads	Added to the "Using the SDK for Swift" chapter two new sections: one on using multipart uploads and another on presigned URLs and requests.	December 7, 2024
Credential identity resolvers	Added a new section on credential identity resolvers, with examples for both AWS IAM Identity Center (SSO) and static resolvers.	November 8, 2024
Sign In With Apple and paginators	Added the new section covering how to use Sign In With Apple to authenticate for AWS services. Also added a new section covering the use of Paginators with the SDK. Removed the Lambda runtime warning added on 2024-10-04.	October 10, 2024

Compatibility note in the Lambda content	Added a note to the chapter on using the SDK to write Lambda functions, indicating that you temporarily can't use the SDK with the Swift AWS Lambda Runtime package.	October 4, 2024
General availability	Content has been updated for the general availability release of the SDK.	September 17, 2024
Updated logging information	The Testing and Debugging chapter now provides up-to-date information about configuring logging for both the SDK and the Common RunTime library.	August 30, 2024
Added Lambda function information	Added the section covering how to create Lambda functions using the Swift AWS Lambda Runtime from Apple's server-side Swift project.	August 21, 2024
Assorted small corrections	Fixed assorted small problems with the document.	August 19, 2024
Update package file information	Added information missing about how to set up the `Package.swift` file.	February 19, 2024
Maintenance and updates	Updated and added links to the SDK reference now that it's in place in its final home. Additionally, synchronized content with latest SDK changes.	February 1, 2024

Updated tool version requirements	Updated the minimum version requirements to Swift 5.7 and Xcode 14.	October 12, 2023
Rewrote section on configuring clients	The section named "Client configuration" has been replaced with a mostly rewritten section named "Customize client configurations". This brings obsolete content up-to-date and provides a working example.	October 10, 2023
Content corrections	Replaced some placeholder text with the correct content and fixed a broken link.	September 12, 2023
Content additions	Renamed the chapter on logging to "Testing and debugging," and added new content covering how to mock the SDK for Swift in tests.	August 28, 2023
Added the section on using waiters	Added the new section on using waiters.	August 9, 2023
Updated setup process to use AWS IAM Identity Center	Updated guide to align with the IAM best practices . For more information, see Security best practices in IAM. This update also features general improvements and updates to bring the documentation closer to being in sync with version 0.20.0 of the SDK.	July 11, 2023

New code examples for Amazon S3 and IAM	Added new examples for Amazon S3 and IAM, cleaned up and removed obsolete information, and fixed reported content errors.	November 4, 2022
Documentation update	Minor corrections and organizational improveme nts to the AWS SDK for Swift Developer Guide.	August 19, 2022
AWS SDK for Swift Developer Preview release	AWS SDK for Swift Developer Preview draft documentation.	December 2, 2021