
AWS SDK for C++

Developer Guide

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

Table of Contents

AWS SDK for C++ Developer Guide	1
Additional Documentation and Resources	1
Getting Started	2
Setting Up the AWS SDK for C++	2
Prerequisites	2
Getting the SDK Using NuGet with Visual C++	3
Getting the SDK Using Vcpkg with Visual C++	3
Building the SDK from Source	3
Providing AWS Credentials	5
Using the AWS SDK for C++	6
Initializing and Shutting Down the SDK	6
Setting SDK options	7
More Information	8
Building Your Application with CMake	8
Setting Up a CMake Project	8
Setting CMAKE_PREFIX_PATH (Optional)	8
Building with CMake	9
Configuring the SDK	10
CMake Parameters	10
General CMake Variables and Options	10
Android CMake Variables and Options	14
AWS Client Configuration	15
Configuration Variables	16
Overriding Your HTTP Client	17
Controlling IOStreams Used by the HttpClient and the AWSClient	17
Using the SDK	19
Using Service Clients	19
Using the Default Credential Provider Chain	19
Passing Credentials Manually	19
Using a Custom Credentials Provider	20
Utility Modules	20
HTTP Stack	20
String Utils	20
Hashing Utils	20
JSON Parser	20
XML Parser	21
Memory Management	21
Allocating and Deallocating Memory	21
STL and AWS Strings and Vectors	22
Remaining Issues	23
Native SDK Developers and Memory Controls	23
Logging	23
Error Handling	24
Code Examples	26
Amazon CloudWatch Examples	26
Getting Metrics from CloudWatch	26
Publishing Custom Metric Data	28
Working with CloudWatch Alarms	29
Using Alarm Actions in CloudWatch	32
Sending Events to CloudWatch	34
Amazon DynamoDB Examples	37
Working with Tables in DynamoDB	37
Working with Items in DynamoDB	42
Amazon EC2 Examples	45

Managing Amazon EC2 Instances	45
Using Elastic IP Addresses in Amazon EC2	52
Using Regions and Availability Zones for Amazon EC2	55
Working with Amazon EC2 Key Pairs	57
Working with Security Groups in Amazon EC2	59
AWS Identity and Access Management (IAM) Examples	62
Managing IAM Access Keys	62
Managing IAM Users	66
Using IAM Account Aliases	70
Working with IAM Policies	72
Working with IAM Server Certificates	78
Amazon S3 Examples	81
Creating, Listing, and Deleting Buckets	82
Operations on Objects	84
Managing Amazon S3 Access Permissions for Buckets and Objects	87
Managing Access to Amazon S3 Buckets Using Bucket Policies	90
Configuring an Amazon S3 Bucket as a Website	92
Amazon SQS Examples	95
Working with Amazon SQS Message Queues	95
Sending, Receiving, and Deleting Amazon SQS Messages	98
Enabling Long Polling for Amazon SQS Message Queues	100
Setting Visibility Timeout in Amazon SQS	103
Using Dead Letter Queues in Amazon SQS	104
Document History	107

AWS SDK for C++ Developer Guide

Welcome to the *AWS SDK for C++ Developer Guide*.

The AWS SDK for C++ provides a modern C++ (version C++ 11 or later) interface for Amazon Web Services (AWS). It provides both high-level and low-level APIs for nearly all AWS features, minimizing dependencies and providing platform portability on Windows, macOS, Linux, and mobile.

Additional Documentation and Resources

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

- [AWS SDK for C++ Reference](#)
- [Video: Introducing the AWS SDK for C++ from AWS re:invent 2015](#)
- [AWS C++ Developer Blog](#)
- GitHub:
 - [SDK source](#)
 - [SDK issues](#)
- [SDK License](#)

Getting Started Using the AWS SDK for C++

The topics in this section will help you set up and use the AWS SDK for C++.

Topics

- [Setting Up the AWS SDK for C++ \(p. 2\)](#)
- [Providing AWS Credentials \(p. 5\)](#)
- [Using the AWS SDK for C++ \(p. 6\)](#)
- [Building Your Application with CMake \(p. 8\)](#)

Setting Up the AWS SDK for C++

This section presents information about how to setup the AWS SDK for C++ on your development platform.

Prerequisites

To use the AWS SDK for C++, you need:

- Visual Studio 2013 or later

Note

Visual Studio 2013 doesn't provide default move constructors and operators. Later versions of Visual Studio provide a standards-compliant compiler.

- *or* GNU Compiler Collection (GCC) 4.9 or later
- *or* Clang 3.3 or later
- A minimum of 4 GB of RAM

Note

You need 4 GB of RAM to build some of the larger AWS clients. The SDK might fail to build on Amazon EC2 instance types *t2.micro*, *t2.small*, and other small instance types due to insufficient memory.

Additional Requirements for Linux Systems

To compile on Linux, you must have the header files (`-dev` packages) for `libcurl`, `libopenssl`, `libuuid`, `zlib`, and optionally, `libpulse` for Amazon Polly support. Typically, you'll find the packages in your system's package manager.

To install these packages on *Debian/Ubuntu-based systems*

```
sudo apt-get install libcurl4-openssl-dev libssl-dev uuid-dev zlib1g-dev libpulse-dev
```

To install these packages on *Redhat/Fedora-based systems*

```
sudo dnf install libcurl-devel openssl-devel libuuid-devel pulseaudio-devel
```

Getting the SDK Using NuGet with Visual C++

You can use NuGet to manage dependencies for AWS SDK for C++ projects that you develop with Microsoft Visual C++. To use this procedure, you must have [NuGet](#) installed on your system.

To use the SDK with NuGet

1. Open your project in Visual Studio.
2. In **Solution Explorer**, right-click your project name, and then choose **Manage NuGet Packages**.
3. Select the packages to use by searching for a particular service or library name. For example, you could use a search term such as `aws s3 native`. Or, because AWS SDK for C++ libraries are named consistently, use `AWSSDKCPP-service name` to add a library for a particular service to your project.
4. Choose **Install** to install the libraries and add them to your project.

When you build your project, the correct binaries are automatically included for each runtime/architecture configuration you use—you won't need to manage these dependencies yourself.

Getting the SDK Using Vcpkg with Visual C++

You can use vcpkg to manage dependencies for AWS SDK for C++ projects that you develop with Microsoft Visual C++. To use this procedure, you must have [vcpkg](#) installed on your system.

To use the SDK with vcpkg

1. Open a Windows command prompt and navigate to the vcpkg directory.
2. Integrate vcpkg into Visual Studio. You can [integrate](#) per project or per user (shown below) to avoid manually editing Visual C++ directory paths.:

```
vcpkg integrate install
```

3. Install the AWS SDK for C++ package. This package compiles the SDK and its dependencies. It can take awhile.:

```
vcpkg install aws-sdk-cpp:x86-windows
```

4. Open your project in Visual Studio.
5. `#include` AWS SDK for C++ header files you want in your source code.

Like NuGet, when you build your project, the correct binaries are automatically included for each runtime/architecture configuration you use.

Building the SDK from Source

If you don't use Visual Studio (or don't want to use NuGet), you can build the SDK from source to set it up for your development system. This method also enables you to customize your SDK build—see [CMake Parameters \(p. 10\)](#) for the available options.

To build the SDK from source

1. Download or clone the SDK source from [aws/aws-sdk-cpp](#) on GitHub.
 - Direct download: [aws/aws-sdk-cpp/archive/master.zip](#)

- Clone with Git:

HTTPS

```
git clone https://github.com/aws/aws-sdk-cpp.git
```

SSH

```
git clone git@github.com:aws/aws-sdk-cpp.git
```

2. Install **cmake** (v3.0+) and the relevant build tools for your platform. Ensure these are available in your `PATH`. If you're unable to install **cmake**, you can use **make** or **msbuild**.
3. Create a directory in which to create your buildfiles, and generate the necessary buildfiles within it. This is the recommended approach, referred to as an *out-of-source build*:

```
mkdir sdk_build  
cd sdk_build  
cmake <path/to/sdk/source>
```

Alternatively, create the build files directly in the SDK source directory:

```
cd <path/to/sdk/source>  
cmake .
```

If you don't have **cmake** installed, you can use these alternative commands to set up your build directory:

auto make

```
make
```

Visual Studio

```
msbuild ALL_BUILD.vcxproj
```

4. Build and install the SDK by typing one of the following in the same location where you generated your build files:

auto make

```
make  
sudo make install
```

Visual Studio

```
msbuild INSTALL.vcxproj
```

Note

Building the entire SDK can take awhile. To build only a particular client, such as Amazon S3, you can use the **cmake** `BUILD_ONLY` parameter. For example:

```
cmake -DBUILD_ONLY="s3"
```

See [CMake Parameters \(p. 10\)](#) for more ways to modify the build output.

Building for Android

To build for Android, add `-DTARGET_ARCH=ANDROID` to your `cmake` command line. The AWS SDK for C++ includes a `cmake` toolchain file that should cover what's needed, assuming you've set the appropriate environment variables (`ANDROID_NDK`).

Android on Windows

Building for Android on Windows requires additional setup. In particular, you have to run `cmake` from a Visual Studio (2013 or later) developer command prompt. You'll also need the commands `git` and `patch` in your path. If you have git installed on a Windows system, you'll most likely find `patch` in a sibling directory (`../Git/usr/bin/`). Once you've verified these requirements, your `cmake` command line will change slightly to use `nmake` .:

```
cmake -G "NMake Makefiles" `--DTARGET_ARCH=ANDROID` <other options> ..
```

`nmake` builds targets in serially. To make things go more quickly, we recommend installing JOM as an alternative to `nmake` , and then changing the `cmake` invocation as follows.:

```
cmake -G "NMake Makefiles JOM" `--DTARGET_ARCH=ANDROID` <other options> ..
```

Creating Release Builds

auto make

```
cmake -DCMAKE_BUILD_TYPE=Release <path/to/sdk/source>  
make  
sudo make install
```

Visual Studio

```
cmake <path-to-root-of-this-source-code> -G "Visual Studio 12 Win64"  
msbuild INSTALL.vcxproj /p:Configuration=Release
```

Running Integration Tests

Several directories are appended with `*integration-tests`. After you build your project, you can run these executables to ensure everything works correctly.

Providing AWS Credentials

To connect to any of the supported services with the AWS SDK for C++, you must provide AWS credentials. The AWS SDKs and CLIs use *provider chains* to look for AWS credentials in several different places, including system/user environment variables and local AWS configuration files.

You can set your credentials for use by the AWS SDK for C++ in various ways, but here are the recommended approaches:

- Set credentials in the AWS credentials profile file on your local system, located at:
 - `~/.aws/credentials` on Linux, macOS, or Unix
 - `C:\Users\USERNAME\.aws\credentials` on Windows

This file should contain lines in the following format:

```
[default]
aws_access_key_id = your_access_key_id
aws_secret_access_key = your_secret_access_key
```

Substitute your own AWS credentials values for the values *your_access_key_id* and *your_secret_access_key*.

- Set the `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY` environment variables.

To set these variables on Linux, macOS, or Unix, use **export** :

```
export AWS_ACCESS_KEY_ID=your_access_key_id
export AWS_SECRET_ACCESS_KEY=your_secret_access_key
```

To set these variables on Windows, use **set** :

```
set AWS_ACCESS_KEY_ID=your_access_key_id
set AWS_SECRET_ACCESS_KEY=your_secret_access_key
```

- For an EC2 instance, specify an IAM role and then give your EC2 instance access to that role. See [IAM Roles for Amazon EC2](#) in the *Amazon EC2 User Guide for Linux Instances* for a detailed discussion about how this works.

After you set your AWS credentials using one of these methods, the AWS SDK for C++ loads them automatically by using the default credential provider chain.

You can also supply AWS credentials using your own methods by:

- Providing credentials to an AWS client class constructor.
- Using [Amazon Cognito](#), an AWS identity management solution. You can use the `CognitoCachingCredentialsProviders` classes in the identity-management project. For more information, see the [Amazon Cognito Developer Guide](#).

Using the AWS SDK for C++

To use the AWS SDK for C++, you should properly initialize it with `Aws::InitAPI` before creating service clients and using them. You should then shut down the SDK with `Aws::ShutdownAPI`.

Both of these functions take an instance of `Aws::SDKOptions`, which you can use to set additional run-time options for SDK calls.

Initializing and Shutting Down the SDK

A basic skeleton application looks like this:

```
#include <aws/core/Aws.h>
int main(int argc, char** argv)
{
    Aws::SDKOptions options;
    Aws::InitAPI(options);
    {
        // make your SDK calls here.
    }
}
```

```
Aws::ShutdownAPI(options);  
return 0;  
}
```

<admonition>
<title>best practice</title>

To properly shut down / clean up any service clients that you may initialize before `Aws::ShutdownAPI` is called, it's a *best practice* to make sure that all SDK calls are made within a pair of curly-braces as shown above, or within another function called between `Aws::InitAPI` and `Aws::ShutdownAPI`.
</admonition>

Setting SDK options

The `Aws::SDKOptions` struct shown in [Initializing and Shutting Down the SDK \(p. 6\)](#) provides a number of options you can set. You should send the same options object to both `Aws::InitAPI` and `Aws::ShutdownAPI`.

A few examples:

- Turn logging on using the default logger:

```
Aws::SDKOptions options;  
options.loggingOptions.logLevel = Aws::Utils::Logging::LogLevel::Info;  
Aws::InitAPI(options);  
{  
    // make your SDK calls here.  
}  
Aws::ShutdownAPI(options);
```

- Install a custom memory manager:

```
MyMemoryManager memoryManager;  
Aws::SDKOptions options;  
options.memoryManagementOptions.memoryManager = &memoryManager;  
Aws::InitAPI(options);  
{  
    // make your SDK calls here.  
}  
Aws::ShutdownAPI(options);
```

- Override the default HTTP client factory:

```
Aws::SDKOptions options;  
options.httpOptions.httpClientFactory_create_fn = [](){  
    return Aws::MakeShared<MyCustomHttpClientFactory>(  
        "ALLOC_TAG", arg1);  
};  
Aws::InitAPI(options);  
{  
    // make your SDK calls here.  
}  
Aws::ShutdownAPI(options);
```

Note

`httpOptions` takes a closure instead of a `std::shared_ptr`. The SDK does this for all of its factory functions because the memory manager will not yet be installed at the time you will need to allocate this memory. Pass a closure to the SDK, and it will be called when it is safe to do so. This simplest way to do this is with a Lambda expression.

More Information

For further examples of AWS SDK for C++ application code, view the topics in the [AWS SDK for C++ Code Examples \(p. 26\)](#) section. Each example contains a link to the full source code on GitHub, which you can use as a starting point for your own applications.

Building Your Application with CMake

CMake is a build tool that can manage your application's dependencies and create native makefiles suitable for the platform you're building on. It's an easy way to create and build projects using the AWS SDK for C++.

<highlightlang></highlightlang>

Setting Up a CMake Project

To set up a CMake project for use with the AWS SDK for C++

1. Create a directory to hold your source files.:

```
mkdir my_example_project
```

2. Open the directory and add a `CMakeLists.txt` file that specifies your project's name, executables, source files, and linked libraries. The following is a minimal example:

```
# minimal CMakeLists.txt for the AWS SDK for C++
cmake_minimum_required(VERSION 3.2)

# "my-example" is just an example value.
project(my-example)

# Locate the AWS SDK for C++ package.
# Requires that you build with:
#   -DCMAKE_PREFIX_PATH=/path/to/sdk_install
find_package(AWSSDK REQUIRED COMPONENTS service1 service2 ...)

# The executable name and its sourcefiles
add_executable(my-example my-example.cpp)

# The libraries used by your executable.
# "aws-cpp-sdk-s3" is just an example.
target_link_libraries(my-example ${AWSSDK_LINK_LIBRARIES})
```

Note

You can set many options in your `CMakeLists.txt` build configuration file. For an introduction to the file's features, see the [CMake tutorial](#) on the CMake website.

Setting CMAKE_PREFIX_PATH (Optional)

CMake needs to know the location of the `aws-sdk-cpp-config.cmake` so that it can properly resolve the AWS SDK libraries that your application uses. You can find this file in the build directory that you used to [build the SDK \(p. 2\)](#).

By setting the path in `CMAKE_PREFIX_PATH`, you won't need to type this path every time you rebuild your application.

You can set it on Linux, macOS, or Unix like this:

```
export CMAKE_PREFIX_PATH=/path/to/sdk_build_dir
```

On Windows, use `set` instead:

```
set CMAKE_PREFIX_PATH=C:\path\to\sdk_build_dir
```

Building with CMake

Create a directory into which `cmake` will build your application.:

```
mkdir my_project_build
```

Open the directory and run `cmake` using the path to your project's source directory.:

```
cd my_project_build  
cmake ../my_example_project
```

If you didn't set `CMAKE_PREFIX_PATH`, you must add the path to the SDK's build directory using `-Daws-sdk-cpp_DIR`.:

```
cmake -Daws-sdk-cpp_DIR=/path/to/sdk_build_dir ../my_example_project
```

After `cmake` generates your build directory, you can use `make` (or `nmake` on Windows) to build your application.:

```
make
```

Configuring the AWS SDK for C++

This section presents information about how to configure the AWS SDK for C++.

Topics

- [CMake Parameters \(p. 10\)](#)
- [AWS Client Configuration \(p. 15\)](#)
- [Overriding Your HTTP Client \(p. 17\)](#)
- [Controlling IOStreams Used by the HttpClient and the AWSClient \(p. 17\)](#)

CMake Parameters

You can set these options with CMake GUI tools or the command line by using `-D`. For example:

```
cmake -DENABLE_UNITY_BUILD=ON -DREGENERATE_CLIENTS=1
```

General CMake Variables and Options

The following are general **cmake** variables and options that affect your SDK build.

Note

To use the `ADD_CUSTOM_CLIENTS` or `REGENERATE_CLIENTS` variables, you must have [Python 2.7](#), [Java \(JDK 1.8+\)](#), and [Maven](#) installed and in your `PATH`.

Topics

- [ADD_CUSTOM_CLIENTS \(p. 10\)](#)
- [BUILD_ONLY \(p. 11\)](#)
- [BUILD_SHARED_LIBS \(p. 11\)](#)
- [CPP_STANDARD \(p. 11\)](#)
- [CUSTOM_MEMORY_MANAGEMENT \(p. 11\)](#)
- [ENABLE_RTTI \(p. 12\)](#)
- [ENABLE_TESTING \(p. 12\)](#)
- [ENABLE_UNITY_BUILD \(p. 12\)](#)
- [FORCE_SHARED_CRT \(p. 12\)](#)
- [G \(p. 12\)](#)
- [MINIMIZE_SIZE \(p. 13\)](#)
- [NO_ENCRYPTION \(p. 13\)](#)
- [NO_HTTP_CLIENT \(p. 13\)](#)
- [REGENERATE_CLIENTS \(p. 13\)](#)
- [SIMPLE_INSTALL \(p. 13\)](#)
- [TARGET_ARCH \(p. 14\)](#)

ADD_CUSTOM_CLIENTS

Builds any arbitrary clients based on the API definition. Place your definition in the `code-generation/api-definitions` folder, and then pass this argument to **cmake**. The **cmake** configure step generates your client and includes it as a subdirectory in your build. This is particularly useful to generate a C++ client for using one of your [API Gateway](#) services. For example:

```
-  
DADD_CUSTOM_CLIENTS="serviceName=myCustomService;version=2015-12-21;serviceName=someOtherService;version=2015-12-21"
```

BUILD_ONLY

Builds only the clients you want to use. If set to a high-level SDK such as `aws-cpp-sdk-transfer`, `BUILD_ONLY` resolves any low-level client dependencies. It also builds integration and unit tests related to the projects you select, if they exist. This is a list argument, with values separated by semicolon (;) characters. For example:

```
-DBUILD_ONLY="s3;cognito-identity"
```

Note

The core SDK module, `aws-sdk-cpp-core`, is *always* built, regardless of the value of the `BUILD_ONLY` parameter.

BUILD_SHARED_LIBS

A built-in CMake option, re-exposed here for visibility. If enabled, it builds shared libraries; otherwise, it builds only static libraries.

Note

To dynamically link to the SDK, you must define the `USE_IMPORT_EXPORT` symbol for all build targets using the SDK.

Values

`ON` | `OFF`

Default

`ON`

CPP_STANDARD

Specifies a custom C++ standard for use with C++ 14 and 17 code bases.

Values

`11` | `14` | `17`

Default

`11`

CUSTOM_MEMORY_MANAGEMENT

To use a custom memory manager, set the value to 1. You can install a custom allocator so that all STL types use the custom allocation interface. If you set the value 0, you still might want to use the STL template types to help with DLL safety on Windows.

If static linking is enabled, custom memory management defaults to *off* (0). If dynamic linking is enabled, custom memory management defaults to *on* (1) and avoids cross-DLL allocation and deallocation.

Note

To prevent linker mismatch errors, you must use the same value (0 or 1) throughout your build system.

To install your own memory manager to handle allocations made by the SDK, you must set `DCUSTOM_MEMORY_MANAGEMENT` and define `AWS_CUSTOM_MEMORY_MANAGEMENT` for all build targets that depend on the SDK.

ENABLE_RTTI

Controls whether the SDK is built to enable run-time type information (RTTI).

Values

`ON` | `OFF`

Default

`ON`

ENABLE_TESTING

Controls whether unit and integration test projects are built during the SDK build.

Values

`ON` | `OFF`

Default

`ON`

ENABLE_UNITY_BUILD

If enabled, most SDK libraries are built as a single, generated `.cpp` file. This can significantly reduce static library size and speed up compilation time.

Values

`ON` | `OFF`

Default

`OFF`

FORCE_SHARED_CRT

If enabled, the SDK links to the C runtime *dynamically*; otherwise, it uses the `BUILD_SHARED_LIBS` setting (sometimes necessary for backward compatibility with earlier versions of the SDK).

Values

`ON` | `OFF`

Default

`ON`

G

Generates build artifacts, such as Visual Studio solutions and Xcode projects.

For example, on Windows:

```
-G "Visual Studio 12 Win64"
```

For more information, see the CMake documentation for your platform.

MINIMIZE_SIZE

A superset of [ENABLE_UNITY_BUILD](#) (p. 12). If enabled, this option turns on `ENABLE_UNITY_BUILD` and additional binary size reduction settings.

Values

`ON` | `OFF`

Default

`OFF`

NO_ENCRYPTION

If enabled, prevents the default platform-specific cryptography implementation from being built into the library. Turn this `ON` to inject your own cryptography implementation.

Values

`ON` | `OFF`

Default

`OFF`

NO_HTTP_CLIENT

If enabled, prevents the default platform-specific HTTP client from being built into the library. Turn this `ON` to inject your own HTTP client implementation.

Values

`ON` | `OFF`

Default

`OFF`

REGENERATE_CLIENTS

This argument wipes out all generated code and generates the client directories from the `code-generation/api-definitions` folder. For example:

```
-DREGENERATE_CLIENTS=1
```

SIMPLE_INSTALL

If enabled, the install process does not insert platform-specific intermediate directories underneath `bin/` and `lib/`. Turn `OFF` if you need to make multiplatform releases under a single install directory.

Values

ON | OFF

Default

ON

TARGET_ARCH

To cross-compile or build for a mobile platform, you must specify the target platform. By default, the build detects the host operating system and builds for the detected operating system.

Note

When *TARGET_ARCH* is *ANDROID*, additional options are available. See [Android CMake Variables and Options](#) (p. 14).

Values

WINDOWS | LINUX | APPLE | ANDROID

Android CMake Variables and Options

Use the following variables when you are creating an Android build of the SDK (when *TARGET_ARCH* (p. 14) is set to *ANDROID*).

Topics

- [ANDROID_ABI](#) (p. 14)
- [ANDROID_NATIVE_API_LEVEL](#) (p. 14)
- [ANDROID_STL](#) (p. 15)
- [ANDROID_TOOLCHAIN_NAME](#) (p. 15)
- [DISABLE_ANDROID_STANDALONE_BUILD](#) (p. 15)
- [NDK_DIR](#) (p. 15)

ANDROID_ABI

Controls which Application Binary Interface (ABI) to output code for.

Note

Not all valid Android ABI values are currently supported.

Values

arm64 | armeabi-v7a | x86_64 | x86 | mips64 | mips

Default

armeabi-v7a

ANDROID_NATIVE_API_LEVEL

Controls what API level the SDK builds against. If you set [ANDROID_STL](#) (p. 15) to *gnustl*, you can choose any API level. If you use *libc++*, you must use an API level of at least 21.

Default

Varies by STL choice.

ANDROID_STL

Controls what flavor of the C++ standard library the SDK uses.

Important

Performance problems can occur within the SDK if the `gnustl` options are used; we strongly recommend using `libc++_shared` or `libc++_static`.

Values

`libc++_shared` | `libc++_static` | `gnustl_shared` | `gnustl_static`

Default

`libc++_shared`

ANDROID_TOOLCHAIN_NAME

Controls which compiler is used to build the SDK.

Note

With GCC being deprecated by the Android NDK, we recommend using the default value.

Default

`standalone-clang`

DISABLE_ANDROID_STANDALONE_BUILD

By default, Android builds use a standalone clang-based toolchain constructed via NDK scripts. To use your own toolchain, turn this option *ON*.

Values

`ON` | `OFF`

Default

`OFF`

NDK_DIR

Specifies an override path where the build system should find the Android NDK. By default, the build system checks environment variables (`ANDROID_NDK`) if this variable is not set.

AWS Client Configuration

You can use the client configuration to control most functionality in the AWS SDK for C++.

ClientConfiguration declaration:

```
struct AWS_CORE_API ClientConfiguration
```

```
{
    ClientConfiguration();

    Aws::String userAgent;
    Aws::Http::Scheme scheme;
    Aws::Region region;
    Aws::String authenticationRegion;
    bool useDualStack;
    unsigned maxConnections;
    long requestTimeoutMs;
    long connectTimeoutMs;
    std::shared_ptr<RetryStrategy> retryStrategy;
    Aws::String endpointOverride;
    Aws::Http::Scheme proxyScheme;
    Aws::String proxyHost;
    unsigned proxyPort;
    Aws::String proxyUserName;
    Aws::String proxyPassword;
    std::shared_ptr<Aws::Utils::Threading::Executor> executor;
    bool verifySSL;
    Aws::String caPath;
    Aws::String caFile;
    std::shared_ptr<Aws::Utils::RateLimits::RateLimiterInterface> writeRateLimiter;
    std::shared_ptr<Aws::Utils::RateLimits::RateLimiterInterface> readRateLimiter;
    Aws::Http::TransferLibType httpLibOverride;
    bool followRedirects;
};
```

Configuration Variables

userAgent

Built in the constructor and pulls information from your operating system. Do not alter the user agent.

scheme

The default value is HTTPS. You can set this value to HTTP if the information you are passing is not sensitive and the service to which you want to connect supports an HTTP endpoint. AWS Auth protects you from tampering.

region

Specifies where you want the client to communicate. Examples include *us-east-1* or *us-west-1*. You must ensure that the service you want to use has an endpoint in the region you configure.

authenticationRegion

If you specify an authentication region, you must also specify an endpoint override.

useDualStack

Use dual stack endpoint in the endpoint calculation. You must ensure that the service you want to use supports ipv6 in the region you select.

maxConnections

The maximum number of allowed connections to a single server for your HTTP communications. The default value is 25. You can set this value as high as you can support the bandwidth. We recommend a value around 25.

requestTimeoutMs and connectTimeoutMs

Values that determine the length of time, in milliseconds, to wait before timing out a request. You can increase this value if you need to transfer large files, such as in Amazon S3 or Amazon CloudFront.

retryStrategy

Defaults to exponential backoff. You can override this default by implementing a subclass of `RetryStrategy` and passing an instance.

endpointOverride

Use this to override the http endpoint used to talk to a service. If you set this, you must also set `authenticationRegion`.

proxyScheme, proxyHost, proxyPort, proxyUserName, and proxyPassword

These settings allow you to configure a proxy for all communication with AWS. Examples of when this functionality might be useful include debugging in conjunction with the Burp suite, or using a proxy to connect to the Internet.

executor

The default behavior is to create and detach a thread for each async call. You can change this behavior by implementing a subclass of `Executor` and passing an instance.

verifySSL

Specifies whether to enable SSL certificate verification. If necessary, you can disable SSL certificate verification by setting `verifySSL` to `false`.

caPath, caFile

Enables you to tell the HTTP client where to find your SSL certificate trust store (for example, a directory prepared with OpenSSL's `c_rehash` utility). You shouldn't need to do this unless you are using symlinks in your environment. This has no effect on Windows or OS X.

writeRateLimiter and readRateLimiter

Used to throttle the bandwidth used by the transport layer. The default for these limiters is open. You can use the default implementation with the rates you want, or you can create your own instance by implementing a subclass of `RateLimiterInterface`.

httpLibOverride

Override the http implementation the default factory returns. The default HTTP client for Windows is `WinHTTP`. The default HTTP client for all other platforms is `curl`.

followRedirects

If set to true the http stack will follow 300 redirect codes

Overriding Your HTTP Client

The default HTTP client for Windows is `WinHTTP`. The default HTTP client for all other platforms is `curl`. If needed, you can create a custom `HttpClientFactory` to pass to any service client's constructor.

Controlling IOStreams Used by the HttpClient and the AWSClient

By default, all responses use an input stream backed by a `stringbuf`. If needed, you can override the default behavior. For example, if you are using an `AmazonS3GetObject` and don't want to load the entire file into memory, you can use `IOStreamFactory` in `AmazonWebServiceRequest` to pass a lambda to create a file stream.

Example file stream request

```
GetObjectRequest getObjectRequest;  
getObjectRequest.SetBucket(fullBucketName);  
getObjectRequest.SetKey(keyName);  
getObjectRequest.SetResponseStreamFactory([](){  
    return Aws::New<Aws::FStream>(  
        ALLOCATION_TAG, DOWNLOADED_FILENAME, std::ios_base::out); });  
  
auto getObjectOutcome = s3Client->GetObject(getObjectRequest);
```

Using the AWS SDK for C++

This section provides information about general use of the AWS SDK for C++, beyond that covered in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#).

For service-specific programming examples, see [AWS SDK for C++ Code Examples \(p. 26\)](#).

Topics

- [Using Service Clients \(p. 19\)](#)
- [Utility Modules \(p. 20\)](#)
- [Memory Management \(p. 21\)](#)
- [Logging \(p. 23\)](#)
- [Error Handling \(p. 24\)](#)

Using Service Clients

AWS service client classes provide you with an interface to the AWS service that the class represents. Service clients follow the namespace convention `Aws::Service::ServiceClient`.

For example, a client for AWS Identity and Access Management is constructed using the `Aws::IAM::IAMClient` class. For an Amazon Simple Storage Service client, use `Aws::S3::S3Client`.

When you use the client classes to instantiate a service client, you must supply AWS credentials. You can do this by using the default credential provider chain, by manually passing credentials to the client directly, or by using a custom credentials provider.

For more information about setting credentials, see [Providing AWS Credentials \(p. 5\)](#).

Using the Default Credential Provider Chain

The following code shows how to create an Amazon DynamoDB client by using a specialized client configuration, default credential provider chain, and default HTTP client factory.

```
auto limiter =  
    Aws::MakeShared<Aws::Utils::RateLimits::DefaultRateLimiter<>>(ALLOCATION_TAG, 200000);  
  
// Create a client  
ClientConfiguration config;  
config.scheme = Scheme::HTTPS;  
config.connectTimeoutMs = 30000;  
config.requestTimeoutMs = 30000;  
config.readRateLimiter = m_limiter;  
config.writeRateLimiter = m_limiter;  
  
auto client = Aws::MakeShared<DynamoDBClient>(ALLOCATION_TAG, config);
```

Passing Credentials Manually

The following code shows how to use the client constructor that takes three arguments, and use the `Aws::Auth::AWSCredentials` class to pass your credentials manually to the constructor.

```
auto client = Aws::MakeShared<DynamoDBClient>(
    ALLOCATION_TAG, AWSCredentials("access_key_id", "secret_key"), config);
```

Using a Custom Credentials Provider

The following code shows how to pass credentials to the `Aws::MakeShared` function and create a client by using one of the credential providers in the `Aws::Auth` namespace.

```
auto client = Aws::MakeShared<DynamoDBClient>(
    ALLOCATION_TAG,
    Aws::MakeShared<CognitoCachingAnonymousCredentialsProvider>(
        ALLOCATION_TAG, "identityPoolId", "accountId"), config);
```

Utility Modules

The AWS SDK for C++ includes many [utility modules](#) to reduce the complexity of developing AWS applications in C++.

HTTP Stack

An HTTP stack that provides connection pooling, is thread-safe, and can be reused as you need. For more information, see [AWS Client Configuration \(p. 15\)](#).

Headers	/aws/core/http/
API Documentation	Aws::Http

String Utils

Core string functions, such as `trim`, `lowercase`, and numeric conversions.

Header	aws/core/utls/StringUtils.h
API Documentation	Aws::Utils::StringUtils

Hashing Utils

Hashing functions such as `SHA256`, `MD5`, `Base64`, and `SHA256_HMAC`.

Header	/aws/core/utls/HashingUtils.h
API Documentation	Aws::Utils::HashingUtils

JSON Parser

A fully functioning yet lightweight JSON parser (a thin wrapper around `JsonCpp`).

Header	/aws/core/utils/json/JsonSerializer.h
API Documentation	Aws::Utils::Json::JsonValue

XML Parser

A lightweight XML parser (a thin wrapper around *tinyxml2*). The [RAII pattern](#) has been added to the interface.

Header	/aws/core/utils/xml/XmlSerializer.h
API Documentation	Aws::Utils::Xml

Memory Management

The AWS SDK for C++ provides a way to control memory allocation and deallocation in a library.

Note

Custom memory management is available only if you use a version of the library built using the defined compile-time constant `AWS_CUSTOM_MEMORY_MANAGEMENT`.

If you use a version of the library that is built without the compile-time constant, global memory system functions such as `InitializeAWSMemorySystem` won't work; the global `new` and `delete` functions are used instead.

For more information about the compile-time constant, see [STL and AWS Strings and Vectors \(p. 22\)](#).

Allocating and Deallocating Memory

To allocate or deallocate memory

1. Subclass `MemorySystemInterface`: `aws/core/utils/memory/MemorySystemInterface.h`.

```
class MyMemoryManager : public Aws::Utils::Memory::MemorySystemInterface
{
public:
    // ...
    virtual void* AllocateMemory(
        std::size_t blockSize, std::size_t alignment,
        const char *allocationTag = nullptr) override;
    virtual void FreeMemory(void* memoryPtr) override;
};
```

Note

You can change the type signature for `AllocateMemory` as needed.

2. Install a memory manager with an instance of your subclass by calling `InitializeAWSMemorySystem`. This should occur at the beginning of your application. For example, in your `main()` function:

```
int main(void)
{
    MyMemoryManager sdkMemoryManager;
    Aws::Utils::Memory::InitializeAWSMemorySystem(sdkMemoryManager);
    // ... do stuff
}
```

```
Aws::Utils::Memory::ShutdownAWSMemorySystem();  
return 0;  
}
```

3. Just before exit, call `ShutdownAWSMemorySystem` (as shown in the preceding example, but repeated here):

```
Aws::Utils::Memory::ShutdownAWSMemorySystem();
```

STL and AWS Strings and Vectors

When initialized with a memory manager, the AWS SDK for C++ defers all allocation and deallocation to the memory manager. If a memory manager doesn't exist, the SDK uses global `new` and `delete`.

If you use custom STL allocators, you must alter the type signatures for all STL objects to match the allocation policy. Because STL is used prominently in the SDK implementation and interface, a single approach in the SDK would inhibit direct passing of default STL objects into the SDK or control of STL allocation. Alternately, a hybrid approach—using custom allocators internally and allowing standard and custom STL objects on the interface—could potentially make it more difficult to investigate memory issues.

The solution is to use the memory system's compile-time constant `AWS_CUSTOM_MEMORY_MANAGEMENT` to control which STL types the SDK uses.

If the compile-time constant is enabled (on), the types resolve to STL types with a custom allocator connected to the AWS memory system.

If the compile-time constant is disabled (off), all `Aws::*` types resolve to the corresponding default `std::*` type.

Example code from the ```AWSAllocator.h``` file in the SDK

```
#ifndef AWS_CUSTOM_MEMORY_MANAGEMENT  
  
template< typename T >  
class AwsAllocator : public std::allocator< T >  
{  
    ... definition of allocator that uses AWS memory system  
};  
  
#else  
  
template< typename T > using Allocator = std::allocator<T>;  
  
#endif
```

In the example code, the `AwsAllocator` can be a custom allocator or a default allocator, depending on the compile-time constant.

Example code from the ```AWSVector.h``` file in the SDK

```
template<typename T> using Vector = std::vector<T, Aws::Allocator<T>>;
```

In the example code, we define the `Aws::*` types.

If the compile-time constant is enabled (on), the type maps to a vector using custom memory allocation and the AWS memory system.

If the compile-time constant is disabled (off), the type maps to a regular `std::vector` with default type parameters.

Type aliasing is used for all `std::` types in the SDK that perform memory allocation, such as containers, string streams, and string buffers. The AWS SDK for C++ uses these types.

Remaining Issues

You can control memory allocation in the SDK; however, STL types still dominate the public interface through string parameters to the model object `initialize` and `set` methods. If you don't use STL and use strings and containers instead, you have to create a lot of temporaries whenever you want to make a service call.

To remove most of the temporaries and allocation when you make service calls using non-STL, we have implemented the following:

- Every Init/Set function that takes a string has an overload that takes a `const char*`.
- Every Init/Set function that takes a container (map/vector) has an add variant that takes a single entry.
- Every Init/Set function that takes binary data has an overload that takes a pointer to the data and a `length` value.
- (Optional) Every Init/Set function that takes a string has an overload that takes a non-zero terminated `const char*` and a `length` value.

Native SDK Developers and Memory Controls

Follow these rules in the SDK code:

- Don't use `new` and `delete`; use `Aws::New<>` and `Aws::Delete<>` instead.
- Don't use `new[]` and `delete[]`; use `Aws::NewArray<>` and `Aws::DeleteArray<>`.
- Don't use `std::make_shared`; use `Aws::MakeShared`.
- Use `Aws::UniquePtr` for unique pointers to a single object. Use the `Aws::MakeUnique` function to create the unique pointer.
- Use `Aws::UniqueArray` for unique pointers to an array of objects. Use the `Aws::MakeUniqueArray` function to create the unique pointer.
- Don't directly use STL containers; use one of the `Aws::` typedefs or add a typedef for the container you want. For example:

```
Aws::Map<Aws::String, Aws::String> m_kvPairs;
```

- Use `shared_ptr` for any external pointer passed into and managed by the SDK. You must initialize the shared pointer with a destruction policy that matches how the object was allocated. You can use a raw pointer if the SDK is not expected to clean up the pointer.

Logging

The AWS SDK for C++ includes logging support that you can configure. When initializing the logging system, you can control the filter level and the logging target (file with a name that has a configurable prefix or a stream). The log file generated by the prefix option rolls over once per hour to allow for archiving or deleting log files.

```
Aws::Utils::Logging::InitializeAWSLogging(  
    Aws::MakeShared<Aws::Utils::Logging::DefaultLogSystem>(
```

```
"RunUnitTests", Aws::Utils::Logging::LogLevel::Trace, "aws_sdk_"));
```

If you don't call `InitializeAWSLogging` in your program, the SDK will not do any logging. If you do use logging, don't forget to shut it down at the end of your program by calling `ShutdownAWSLogging`:

```
Aws::Utils::Logging::ShutdownAWSLogging();
```

Example integration test with logging

```
#include <aws/external/gtest.h>

#include <aws/core/utils/memory/stl/AWSString.h>
#include <aws/core/utils/logging/DefaultLogSystem.h>
#include <aws/core/utils/logging/AWSLogging.h>

#include <iostream>

int main(int argc, char** argv)
{
    Aws::Utils::Logging::InitializeAWSLogging(
        Aws::MakeShared<Aws::Utils::Logging::DefaultLogSystem>(
            "RunUnitTests", Aws::Utils::Logging::LogLevel::Trace, "aws_sdk_"));
    ::testing::InitGoogleTest(&argc, argv);
    int exitCode = RUN_ALL_TESTS();
    Aws::Utils::Logging::ShutdownAWSLogging();
    return exitCode;
}
```

Error Handling

The AWS SDK for C++ does not use exceptions; however, you can use exceptions in your code. Every service client returns an outcome object that includes the result and an error code.

Example of handling error conditions

```
bool CreateTableAndWaitForItToBeActive()
{
    CreateTableRequest createTableRequest;
    AttributeDefinition hashKey;
    hashKey.SetAttributeName(HASH_KEY_NAME);
    hashKey.SetAttributeType(ScalarAttributeType::S);
    createTableRequest.AddAttributeDefinitions(hashKey);
    KeySchemaElement hashKeySchemaElement;
    hashKeySchemaElement.WithAttributeName(HASH_KEY_NAME).WithKeyType(KeyType::HASH);
    createTableRequest.AddKeySchema(hashKeySchemaElement);
    ProvisionedThroughput provisionedThroughput;
    provisionedThroughput.SetReadCapacityUnits(readCap);
    provisionedThroughput.SetWriteCapacityUnits(writeCap);
    createTableRequest.WithProvisionedThroughput(provisionedThroughput);
    createTableRequest.WithTableName(tableName);

    CreateTableOutcome createTableOutcome = dynamoDbClient->CreateTable(createTableRequest);
    if (createTableOutcome.IsSuccess())
    {
        DescribeTableRequest describeTableRequest;
        describeTableRequest.SetTableName(tableName);
        bool shouldContinue = true;
        DescribeTableOutcome outcome = dynamoDbClient->DescribeTable(describeTableRequest);
    }
}
```

```
while (shouldContinue)
{
    if (outcome.GetResult().GetTable().GetTableStatus() == TableStatus::ACTIVE)
    {
        break;
    }
    else
    {
        std::this_thread::sleep_for(std::chrono::seconds(1));
    }
}
return true;
}
else if(createTableOutcome.GetError().GetErrorType() == DynamoDBErrors::RESOURCE_IN_USE)
{
    return true;
}
return false;
}
```

AWS SDK for C++ Code Examples

This section provides examples, guidance, and tips you can use to work with specific AWS services using the AWS SDK for C++.

Topics

- [Amazon CloudWatch Examples Using the AWS SDK for C++ \(p. 26\)](#)
- [Amazon DynamoDB Examples Using the AWS SDK for C++ \(p. 37\)](#)
- [Amazon EC2 Examples Using the AWS SDK for C++ \(p. 45\)](#)
- [IAM Code Examples Using the AWS SDK for C++ \(p. 62\)](#)
- [Amazon S3 Code Examples Using the AWS SDK for C++ \(p. 81\)](#)
- [Amazon SQS Code Examples Using the AWS SDK for C++ \(p. 95\)](#)

Amazon CloudWatch Examples Using the AWS SDK for C++

Amazon CloudWatch (CloudWatch) is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use the following examples to program [CloudWatch](#) using the [AWS SDK for C++](#).

Amazon CloudWatch monitors your Amazon Web Services (AWS) resources and the applications you run on AWS in real time. You can use CloudWatch to collect and track metrics, which are variables you can measure for your resources and applications. CloudWatch alarms send notifications or automatically make changes to the resources you are monitoring based on rules that you define.

For more information about CloudWatch, see the [Amazon CloudWatch User Guide](#).

Note

Only the code that is necessary to demonstrate each technique is supplied here, but [complete example code is available on GitHub](#), where you can download a single source file or you can clone the repository locally to get all examples, build and run them.

Topics

- [Getting Metrics from CloudWatch \(p. 26\)](#)
- [Publishing Custom Metric Data \(p. 28\)](#)
- [Working with CloudWatch Alarms \(p. 29\)](#)
- [Using Alarm Actions in CloudWatch \(p. 32\)](#)
- [Sending Events to CloudWatch \(p. 34\)](#)

Getting Metrics from CloudWatch

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Listing Metrics

To list CloudWatch metrics, create a [ListMetricsRequest](#) and call the [CloudWatchClient](#)'s `ListMetrics` function. You can use the `ListMetricsRequest` to filter the returned metrics by namespace, metric name, or dimensions.

Note

A list of metrics and dimensions that are posted by AWS services can be found within the [Amazon CloudWatch Metrics and Dimensions Reference](#) in the *Amazon CloudWatch User Guide*.

Includes

```
#include <aws/core/Aws.h>
#include <aws/monitoring/CloudWatchClient.h>
#include <aws/monitoring/model/ListMetricsRequest.h>
#include <aws/monitoring/model/ListMetricsResult.h>
#include <iomanip>
#include <iostream>
```

Code

```
Aws::CloudWatch::CloudWatchClient cw;
Aws::CloudWatch::Model::ListMetricsRequest request;

if (argc > 1)
{
    request.SetMetricName(argv[1]);
}

if (argc > 2)
{
    request.SetNamespace(argv[2]);
}

bool done = false;
bool header = false;
while (!done)
{
    auto outcome = cw.ListMetrics(request);
    if (!outcome.IsSuccess())
    {
        std::cout << "Failed to list cloudwatch metrics:" <<
            outcome.GetError().GetMessage() << std::endl;
        break;
    }

    if (!header)
    {
        std::cout << std::left << std::setw(48) << "MetricName" <<
            std::setw(32) << "Namespace" << "DimensionNameValuePairs" <<
            std::endl;
        header = true;
    }

    const auto &metrics = outcome.GetResult().GetMetrics();
    for (const auto &metric : metrics)
    {
        std::cout << std::left << std::setw(48) <<
            metric.GetMetricName() << std::setw(32) <<
            metric.GetNamespace();
        const auto &dimensions = metric.GetDimensions();
        for (auto iter = dimensions.cbegin();
```

```
        iter != dimensions.cend(); ++iter)
    {
        const auto &dimkv = *iter;
        std::cout << dimkv.GetName() << " = " << dimkv.GetValue();
        if (iter + 1 != dimensions.cend())
        {
            std::cout << ", ";
        }
    }
    std::cout << std::endl;
}

const auto &next_token = outcome.GetResult().GetNextToken();
request.SetNextToken(next_token);
```

The metrics are returned in a [ListMetricsResult](#) by calling its `GetMetrics` function. The results may be *paged*. To retrieve the next batch of results, call `SetNextToken` on the original request object with the return value of the `ListMetricsResult` object's `GetNextToken` function, and pass the modified request object back to another call to `ListMetrics`.

See the [complete example](#).

More Information

- [ListMetrics](#) in the *Amazon CloudWatch API Reference*.

Publishing Custom Metric Data

A number of AWS services publish [their own metrics](#) in namespaces beginning with "AWS/" You can also publish custom metric data using your own namespace (as long as it doesn't begin with "AWS/").

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Publish Custom Metric Data

To publish your own metric data, call the [CloudWatchClient](#)'s `PutMetricData` function with a [PutMetricDataRequest](#). The `PutMetricDataRequest` must include the custom namespace to use for the data, and information about the data point itself in a [MetricDatum](#) object.

Note

You cannot specify a namespace that begins with "AWS/". Namespaces that begin with "AWS/" are reserved for use by Amazon Web Services products. .

Includes

```
#include <aws/core/Aws.h>
#include <aws/monitoring/CloudWatchClient.h>
#include <aws/monitoring/model/PutMetricDataRequest.h>
#include <iostream>
```

Code

```
Aws::CloudWatch::CloudWatchClient cw;

Aws::CloudWatch::Model::Dimension dimension;
```



```
dimension.SetName("UNIQUE_PAGES");
dimension.SetValue("URLS");

Aws::CloudWatch::Model::MetricDatum datum;
datum.SetMetricName("PAGES_VISITED");
datum.SetUnit(Aws::CloudWatch::Model::StandardUnit::None);
datum.SetValue(data_point);
datum.AddDimensions(dimension);

Aws::CloudWatch::Model::PutMetricDataRequest request;
request.SetNamespace("SITE/TRAFFIC");
request.AddMetricData(datum);

auto outcome = cw.PutMetricData(request);
if (!outcome.IsSuccess())
{
    std::cout << "Failed to put sample metric data:" <<
        outcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully put sample metric data" << std::endl;
}
```

See the [complete example](#).

More Information

- [Using Amazon CloudWatch Metrics](#) in the *Amazon CloudWatch User Guide*.
- [AWS Namespaces](#) in the *Amazon CloudWatch User Guide*.
- [PutMetricData](#) in the *Amazon CloudWatch API Reference*.

Working with CloudWatch Alarms

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++](#) (p. 2) and have configured default AWS credentials using the information in [Providing AWS Credentials](#) (p. 5).

Create an Alarm

To create an alarm based on a CloudWatch metric, call the [CloudWatchClient](#)'s `PutMetricAlarm` function with a [PutMetricAlarmRequest](#) filled with the alarm conditions.

Includes

```
#include <aws/core/Aws.h>
#include <aws/monitoring/CloudWatchClient.h>
#include <aws/monitoring/model/PutMetricAlarmRequest.h>
#include <iostream>
```

Code

```
Aws::CloudWatch::CloudWatchClient cw;
Aws::CloudWatch::Model::PutMetricAlarmRequest request;
request.SetAlarmName(alarm_name);
request.SetComparisonOperator(
    Aws::CloudWatch::Model::ComparisonOperator::GreaterThanThreshold);
```

```
request.SetEvaluationPeriods(1);
request.SetMetricName("CPUUtilization");
request.SetNamespace("AWS/EC2");
request.SetPeriod(60);
request.SetStatistic(Aws::CloudWatch::Model::Statistic::Average);
request.SetThreshold(70.0);
request.SetActionsEnabled(false);
request.SetAlarmDescription("Alarm when server CPU exceeds 70%");
request.SetUnit(Aws::CloudWatch::Model::StandardUnit::Seconds);

Aws::CloudWatch::Model::Dimension dimension;
dimension.SetName("InstanceId");
dimension.SetValue(instanceId);

request.AddDimensions(dimension);

auto outcome = cw.PutMetricAlarm(request);
if (!outcome.IsSuccess())
{
    std::cout << "Failed to create cloudwatch alarm:" <<
        outcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully created cloudwatch alarm " << alarm_name
        << std::endl;
}
```

See the [complete example](#).

List Alarms

To list the CloudWatch alarms that you have created, call the [CloudWatchClient's DescribeAlarms](#) function with a [DescribeAlarmsRequest](#) that you can use to set options for the result.

Includes

```
#include <aws/core/Aws.h>
#include <aws/monitoring/CloudWatchClient.h>
#include <aws/monitoring/model/DescribeAlarmsRequest.h>
#include <aws/monitoring/model/DescribeAlarmsResult.h>
#include <iomanip>
#include <iostream>
```

Code

```
Aws::CloudWatch::CloudWatchClient cw;
Aws::CloudWatch::Model::DescribeAlarmsRequest request;
request.SetMaxRecords(1);

bool done = false;
bool header = false;
while (!done)
{
    auto outcome = cw.DescribeAlarms(request);
    if (!outcome.IsSuccess())
    {
        std::cout << "Failed to describe cloudwatch alarms:" <<
            outcome.GetError().GetMessage() << std::endl;
        break;
    }
}
```

```
if (!header)
{
    std::cout << std::left <<
        std::setw(32) << "Name" <<
        std::setw(64) << "Arn" <<
        std::setw(64) << "Description" <<
        std::setw(20) << "LastUpdated" <<
        std::endl;
    header = true;
}

const auto &alarms = outcome.GetResult().GetMetricAlarms();
for (const auto &alarm : alarms)
{
    std::cout << std::left <<
        std::setw(32) << alarm.GetAlarmName() <<
        std::setw(64) << alarm.GetAlarmArn() <<
        std::setw(64) << alarm.GetAlarmDescription() <<
        std::setw(20) <<
        alarm.GetAlarmConfigurationUpdatedTimestamp().ToGmtString(
            SIMPLE_DATE_FORMAT_STR) <<
        std::endl;
}

const auto &next_token = outcome.GetResult().GetNextToken();
request.SetNextToken(next_token);
```

The list of alarms can be obtained by calling `getMetricAlarms` on the [DescribeAlarmsResult](#) that is returned by `DescribeAlarms`.

The results may be *paged*. To retrieve the next batch of results, call `SetNextToken` on the original request object with the return value of the `DescribeAlarmsResult` object's `GetNextToken` function, and pass the modified request object back to another call to `DescribeAlarms`.

Note

You can also retrieve alarms for a specific metric by using the [CloudWatchClient](#)'s `DescribeAlarmsForMetric` function. Its use is similar to `DescribeAlarms`.

See the [complete example](#).

Delete Alarms

To delete CloudWatch alarms, call the [CloudWatchClient](#)'s `DeleteAlarms` function with a [DeleteAlarmsRequest](#) containing one or more names of alarms that you want to delete.

Includes

```
#include <aws/core/Aws.h>
#include <aws/monitoring/CloudWatchClient.h>
#include <aws/monitoring/model/DeleteAlarmsRequest.h>
#include <iostream>
```

Code

```
Aws::CloudWatch::CloudWatchClient cw;
Aws::CloudWatch::Model::DeleteAlarmsRequest request;
request.AddAlarmNames(alarm_name);

auto outcome = cw.DeleteAlarms(request);
if (!outcome.IsSuccess())
```

```
{
    std::cout << "Failed to delete cloudwatch alarm:" <<
        outcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully deleted cloudwatch alarm " << alarm_name
        << std::endl;
}
```

See the [complete example](#).

More Information

- [Creating Amazon CloudWatch Alarms](#) in the *Amazon CloudWatch User Guide*
- [PutMetricAlarm](#) in the *Amazon CloudWatch API Reference*
- [DescribeAlarms](#) in the *Amazon CloudWatch API Reference*
- [DeleteAlarms](#) in the *Amazon CloudWatch API Reference*

Using Alarm Actions in CloudWatch

Using CloudWatch alarm actions, you can create alarms that perform actions such as automatically stopping, terminating, rebooting, or recovering Amazon EC2 instances.

Alarm actions can be added to an alarm by using the [PutMetricAlarmRequest](#)'s `SetAlarmActions` function when [creating an alarm](#) (p. 29).

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++](#) (p. 2) and have configured default AWS credentials using the information in [Providing AWS Credentials](#) (p. 5).

Enable Alarm Actions

To enable alarm actions for a CloudWatch alarm, call the [CloudWatchClient](#)'s `EnableAlarmActions` with a [EnableAlarmActionsRequest](#) containing one or more names of alarms whose actions you want to enable.

Includes

```
#include <aws/core/Aws.h>
#include <aws/monitoring/CloudWatchClient.h>
#include <aws/monitoring/model/EnableAlarmActionsRequest.h>
#include <aws/monitoring/model/PutMetricAlarmRequest.h>
#include <iostream>
```

Code

```
Aws::CloudWatch::CloudWatchClient cw;
Aws::CloudWatch::Model::PutMetricAlarmRequest request;
request.SetAlarmName(alarm_name);
request.SetComparisonOperator(
    Aws::CloudWatch::Model::ComparisonOperator::GreaterThanThreshold);
request.SetEvaluationPeriods(1);
request.SetMetricName("CPUUtilization");
request.SetNamespace("AWS/EC2");
request.SetPeriod(60);
```

```
request.SetStatistic(Aws::CloudWatch::Model::Statistic::Average);
request.SetThreshold(70.0);
request.SetActionsEnabled(false);
request.SetAlarmDescription("Alarm when server CPU exceeds 70%");
request.SetUnit(Aws::CloudWatch::Model::StandardUnit::Seconds);
request.AddAlarmActions(actionArn);

Aws::CloudWatch::Model::Dimension dimension;
dimension.SetName("InstanceId");
dimension.SetValue(instanceId);
request.AddDimensions(dimension);

auto outcome = cw.PutMetricAlarm(request);
if (!outcome.IsSuccess())
{
    std::cout << "Failed to create cloudwatch alarm:" <<
        outcome.GetError().GetMessage() << std::endl;
    return;
}

Aws::CloudWatch::Model::EnableAlarmActionsRequest enable_request;
enable_request.AddAlarmNames(alarm_name);

auto enable_outcome = cw.EnableAlarmActions(enable_request);
if (!enable_outcome.IsSuccess())
{
    std::cout << "Failed to enable alarm actions:" <<
        enable_outcome.GetError().GetMessage() << std::endl;
    return;
}

std::cout << "Successfully created alarm " << alarm_name <<
    " and enabled actions on it." << std::endl;
```

See the [complete example](#).

Disable Alarm Actions

To disable alarm actions for a CloudWatch alarm, call the [CloudWatchClient](#)'s `DisableAlarmActions` with a [DisableAlarmActionsRequest](#) containing one or more names of alarms whose actions you want to disable.

Includes

```
#include <aws/core/Aws.h>
#include <aws/monitoring/CloudWatchClient.h>
#include <aws/monitoring/model/DisableAlarmActionsRequest.h>
#include <iostream>
```

Code

```
Aws::CloudWatch::CloudWatchClient cw;

Aws::CloudWatch::Model::DisableAlarmActionsRequest disableAlarmActionsRequest;
disableAlarmActionsRequest.AddAlarmNames(alarm_name);

auto disableAlarmActionsOutcome = cw.DisableAlarmActions(disableAlarmActionsRequest);
if (!disableAlarmActionsOutcome.IsSuccess())
{
    std::cout << "Failed to disable actions for alarm " << alarm_name <<
        ": " << disableAlarmActionsOutcome.GetError().GetMessage() <<
        std::endl;
```

```
}  
else  
{  
    std::cout << "Successfully disabled actions for alarm " <<  
        alarm_name << std::endl;  
}
```

See the [complete example](#).

More Information

- [Create Alarms to Stop, Terminate, Reboot, or Recover an Instance](#) in the *Amazon CloudWatch User Guide*
- [PutMetricAlarm](#) in the *Amazon CloudWatch API Reference*
- [EnableAlarmActions](#) in the *Amazon CloudWatch API Reference*
- [DisableAlarmActions](#) in the *Amazon CloudWatch API Reference*

Sending Events to CloudWatch

CloudWatch Events delivers a near real-time stream of system events that describe changes in AWS resources to Amazon EC2 instances, Lambda functions, Kinesis streams, Amazon ECS tasks, Step Functions state machines, Amazon SNS topics, Amazon SQS queues, or built-in targets. You can match events and route them to one or more target functions or streams by using simple rules.

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Add Events

To add custom CloudWatch events, call the [CloudWatchEventsClient's PutEvents](#) function with a [PutEventsRequest](#) object that contains one or more [PutEventsRequestEntry](#) objects that provide details about each event. You can specify several parameters for the entry such as the source and type of the event, resources associated with the event, and so on.

Note

You can specify a maximum of 10 events per call to `putEvents`.

Includes

```
#include <aws/core/Aws.h>  
#include <aws/events/CloudWatchEventsClient.h>  
#include <aws/events/model/PutEventsRequest.h>  
#include <aws/events/model/PutEventsResult.h>  
#include <aws/core/utils/Outcome.h>  
#include <iostream>
```

Code

```
:String MakeDetails(const Aws::String &key, const Aws::String& value)  
  
Aws::Utils::Json::JsonValue value_entry;  
value_entry.AsString(value);  
  
Aws::Utils::Json::JsonValue detail_map;
```

```
detail_map.WithObject(key, value_entry);

return detail_map.View().WriteReadable();

    Aws::CloudWatchEvents::CloudWatchEventsClient cwe;

    Aws::CloudWatchEvents::Model::PutEventsRequestEntry event_entry;
    event_entry.SetDetail(MakeDetails(event_key, event_value));
    event_entry.SetDetailType("sampleSubmitted");
    event_entry.AddResources(resource_arn);
    event_entry.SetSource("aws-sdk-cpp-cloudwatch-example");

    Aws::CloudWatchEvents::Model::PutEventsRequest request;
    request.AddEntries(event_entry);

    auto outcome = cwe.PutEvents(request);
    if (!outcome.IsSuccess())
    {
        std::cout << "Failed to post cloudwatch event: " <<
            outcome.GetError().GetMessage() << std::endl;
    }
    else
    {
        std::cout << "Successfully posted cloudwatch event" << std::endl;
    }
}
```

Add Rules

To create or update a rule, call the [CloudWatchEventsClient](#)'s `PutRule` function with a [PutRuleRequest](#) with the name of the rule and optional parameters such as the [event pattern](#), IAM role to associate with the rule, and a [scheduling expression](#) that describes how often the rule is run.

Includes

```
#include <aws/core/Aws.h>
#include <aws/events/CloudWatchEventsClient.h>
#include <aws/events/model/PutRuleRequest.h>
#include <aws/events/model/PutRuleResult.h>
#include <aws/core/Utils/Outcome.h>
#include <iostream>
```

Code

```
Aws::CloudWatchEvents::CloudWatchEventsClient cwe;
Aws::CloudWatchEvents::Model::PutRuleRequest request;
request.SetName(rule_name);
request.SetRoleArn(role_arn);
request.SetScheduleExpression("rate(5 minutes)");
request.SetState(Aws::CloudWatchEvents::Model::RuleState::ENABLED);

auto outcome = cwe.PutRule(request);
if (!outcome.IsSuccess())
{
    std::cout << "Failed to create cloudwatch events rule " <<
        rule_name << ": " << outcome.GetError().GetMessage() <<
        std::endl;
}
else
{
    std::cout << "Successfully created cloudwatch events rule " <<
        rule_name << " with resulting Arn " <<
        outcome.GetResult().GetRuleArn() << std::endl;
}
```

```
}

```

Add Targets

Targets are the resources that are invoked when a rule is triggered. Example targets include Amazon EC2 instances, Lambda functions, Kinesis streams, Amazon ECS tasks, Step Functions state machines, and built-in targets.

To add a target to a rule, call the [CloudWatchEventsClient](#)'s `PutTargets` function with a [PutTargetsRequest](#) containing the rule to update and a list of targets to add to the rule.

Includes

```
#include <aws/core/Aws.h>
#include <aws/events/CloudWatchEventsClient.h>
#include <aws/events/model/PutTargetsRequest.h>
#include <aws/events/model/PutTargetsResult.h>
#include <aws/core/utils/Outcome.h>
#include <iostream>
```

Code

```
Aws::CloudWatchEvents::CloudWatchEventsClient cwe;

Aws::CloudWatchEvents::Model::Target target;
target.SetArn(lambda_arn);
target.SetId(target_id);

Aws::CloudWatchEvents::Model::PutTargetsRequest request;
request.SetRule(rule_name);
request.AddTargets(target);

auto putTargetsOutcome = cwe.PutTargets(request);
if (!putTargetsOutcome.IsSuccess())
{
    std::cout << "Failed to create cloudwatch events target for rule "
              << rule_name << ": " <<
              putTargetsOutcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout <<
      "Successfully created cloudwatch events target for rule "
      << rule_name << std::endl;
}
}
```

See the [complete example](#).

More Information

- [Adding Events with PutEvents](#) in the *Amazon CloudWatch Events User Guide*
- [Schedule Expressions for Rules](#) in the *Amazon CloudWatch Events User Guide*
- [Event Types for CloudWatch Events](#) in the *Amazon CloudWatch Events User Guide*
- [Events and Event Patterns](#) in the *Amazon CloudWatch Events User Guide*
- [PutEvents](#) in the *Amazon CloudWatch Events API Reference*
- [PutTargets](#) in the *Amazon CloudWatch Events API Reference*
- [PutRule](#) in the *Amazon CloudWatch Events API Reference*

Amazon DynamoDB Examples Using the AWS SDK for C++

Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. The following examples show how you can program [DynamoDB](#) using the [AWS SDK for C++](#).

Note

Only the code that is necessary to demonstrate each technique is supplied here, but [complete example code is available on GitHub](#), where you can download a single source file or you can clone the repository locally to get all examples, build and run them.

Topics

- [Working with Tables in DynamoDB \(p. 37\)](#)
- [Working with Items in DynamoDB \(p. 42\)](#)

Working with Tables in DynamoDB

Tables are the containers for all items in a DynamoDB database. Before you can add or remove data from DynamoDB, you must create a table.

For each table, you must define:

- A table *name* that is unique for your account and region.
- A *primary key* for which every value must be unique. No two items in your table can have the same primary key value.

A primary key can be *simple*, consisting of a single partition (HASH) key, or *composite*, consisting of a partition and a sort (RANGE) key.

Each key value has an associated *data type*, enumerated by the [ScalarAttributeType](#) class. The key value can be binary (B), numeric (N), or a string (S). For more information, see [Naming Rules and Data Types](#) in the *Amazon DynamoDB Developer Guide*.

- *Provisioned throughput* values that define the number of reserved read/write capacity units for the table.

Note

[Amazon DynamoDB pricing](#) is based on the provisioned throughput values that you set on your tables, so reserve only as much capacity as you think you'll need for your table. Provisioned throughput for a table can be modified at any time, so you can adjust capacity if your needs change.

Create a Table

Use the [DynamoDB client](#) `CreateTable` method to create a new DynamoDB table. You need to construct table attributes and a table schema, both of which are used to identify the primary key of your table. You must also supply initial provisioned throughput values and a table name. `CreateTable` is an asynchronous operation. `GetTableStatus` will return `CREATING` until the table is `ACTIVE` and ready for use.

Create a Table with a Simple Primary Key

This code creates a table with a simple primary key ("Name").

Includes

```
#include <aws/core/Aws.h>
#include <aws/core/utils/Outcome.h>
#include <aws/dynamodb/DynamoDBClient.h>
#include <aws/dynamodb/model/AttributeDefinition.h>
#include <aws/dynamodb/model/CreateTableRequest.h>
#include <aws/dynamodb/model/KeySchemaElement.h>
#include <aws/dynamodb/model/ProvisionedThroughput.h>
#include <aws/dynamodb/model/ScalarAttributeType.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration clientConfig;
Aws::DynamoDB::DynamoDBClient dynamoClient(clientConfig);
Aws::DynamoDB::Model::CreateTableRequest req;

Aws::DynamoDB::Model::AttributeDefinition haskKey;
haskKey.SetAttributeName("Name");
haskKey.SetAttributeType(Aws::DynamoDB::Model::ScalarAttributeType::S);
req.AddAttributeDefinitions(haskKey);

Aws::DynamoDB::Model::KeySchemaElement keyscelt;
keyscelt.WithAttributeName("Name").WithKeyType(Aws::DynamoDB::Model::KeyType::HASH);
req.AddKeySchema(keyscelt);

Aws::DynamoDB::Model::ProvisionedThroughput thruput;
thruput.WithReadCapacityUnits(5).WithWriteCapacityUnits(5);
req.SetProvisionedThroughput(thruput);

req.SetTableName(table);

const Aws::DynamoDB::Model::CreateTableOutcome& result = dynamoClient.CreateTable(req);
if (result.IsSuccess())
{
    std::cout << "Table \"" << result.GetResult().GetTableDescription().GetTableName() <<
        " was created!" << std::endl;
}
else
{
    std::cout << "Failed to create table: " << result.GetError().GetMessage();
}
```

See the [complete example](#).

Create a Table with a Composite Primary Key

Add another [AttributeDefinition](#) and [KeySchemaElement](#) to [CreateTableRequest](#).

Includes

```
#include <aws/core/Aws.h>
#include <aws/core/utils/Outcome.h>
#include <aws/dynamodb/DynamoDBClient.h>
#include <aws/dynamodb/model/AttributeDefinition.h>
#include <aws/dynamodb/model/CreateTableRequest.h>
#include <aws/dynamodb/model/KeySchemaElement.h>
#include <aws/dynamodb/model/ProvisionedThroughput.h>
#include <aws/dynamodb/model/ScalarAttributeType.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration clientConfig;
Aws::DynamoDB::DynamoDBClient dynamoClient(clientConfig);
Aws::DynamoDB::Model::CreateTableRequest req;

Aws::DynamoDB::Model::AttributeDefinition hashKey1, hashKey2;
hashKey1.WithAttributeName("Language").WithAttributeType(Aws::DynamoDB::Model::ScalarAttributeType::S);
req.AddAttributeDefinitions(hashKey1);
hashKey2.WithAttributeName("Greeting").WithAttributeType(Aws::DynamoDB::Model::ScalarAttributeType::S);
req.AddAttributeDefinitions(hashKey2);

Aws::DynamoDB::Model::KeySchemaElement kse1, kse2;
kse1.WithAttributeName("Language").WithKeyType(Aws::DynamoDB::Model::KeyType::HASH);
req.AddKeySchema(kse1);
kse2.WithAttributeName("Greeting").WithKeyType(Aws::DynamoDB::Model::KeyType::RANGE);
req.AddKeySchema(kse2);

Aws::DynamoDB::Model::ProvisionedThroughput thruput;
thruput.WithReadCapacityUnits(5).WithWriteCapacityUnits(5);
req.SetProvisionedThroughput(thruput);

req.SetTableName(table);

const Aws::DynamoDB::Model::CreateTableOutcome& result = dynamoClient.CreateTable(req);
if (result.IsSuccess())
{
    std::cout << "Table \"" << result.GetResult().GetTableDescription().GetTableName() <<
        "\" was created!\n";
}
else
{
    std::cout << "Failed to create table:" << result.GetError().GetMessage();
}
}
```

See the [complete example](#) on GitHub.

List Tables

You can list the tables in a particular region by calling the [DynamoDB client](#) `ListTables` method.

Includes

```
#include <aws/core/Aws.h>
#include <aws/core/Utils/Outcome.h>
#include <aws/dynamodb/DynamoDBClient.h>
#include <aws/dynamodb/model/ListTablesRequest.h>
#include <aws/dynamodb/model/ListTablesResult.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration clientConfig;
Aws::DynamoDB::DynamoDBClient dynamoClient(clientConfig);

Aws::DynamoDB::Model::ListTablesRequest ltr;
ltr.SetLimit(50);
do
{
    const Aws::DynamoDB::Model::ListTablesOutcome& lto = dynamoClient.ListTables(ltr);
    if (!lto.IsSuccess())
```

```
{
    std::cout << "Error: " << lto.GetError().GetMessage() << std::endl;
    return 1;
}
for (const auto& s : lto.GetResult().GetTableNames())
    std::cout << s << std::endl;
ltr.SetExclusiveStartTableName(lto.GetResult().GetLastEvaluatedTableName());
} while (!ltr.GetExclusiveStartTableName().empty());
```

By default, up to 100 tables are returned per call. Use `GetExclusiveStartTableName` on the returned [ListTablesOutcome](#) object to get the last table that was evaluated. You can use this value to start the listing after the last returned value of the previous listing.

See the [complete example](#).

Describe (Get Information about) a Table

You can find out more about a table by calling the [DynamoDB client](#) `DescribeTable` method.

Includes

```
#include <aws/core/Aws.h>
#include <aws/core/utils/Outcome.h>
#include <aws/dynamodb/DynamoDBClient.h>
#include <aws/dynamodb/model/DescribeTableRequest.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration clientConfig;
Aws::DynamoDB::DynamoDBClient dynamoClient(clientConfig);

Aws::DynamoDB::Model::DescribeTableRequest dtr;
dtr.SetTableName(table);

const Aws::DynamoDB::Model::DescribeTableOutcome& result = dynamoClient.DescribeTable(dtr);

if (result.IsSuccess())
{
    const Aws::DynamoDB::Model::TableDescription& td = result.GetResult().GetTable();
    std::cout << "Table name   : " << td.GetTableName() << std::endl;
    std::cout << "Table ARN   : " << td.GetTableArn() << std::endl;
    std::cout << "Status     : " <<
    Aws::DynamoDB::Model::TableStatusMapper::GetNameForTableStatus(td.GetTableStatus()) <<
    std::endl;
    std::cout << "Item count  : " << td.GetItemCount() << std::endl;
    std::cout << "Size (bytes): " << td.GetTableSizeBytes() << std::endl;

    const Aws::DynamoDB::Model::ProvisionedThroughputDescription& ptd =
    td.GetProvisionedThroughput();
    std::cout << "Throughput" << std::endl;
    std::cout << "  Read Capacity : " << ptd.GetReadCapacityUnits() << std::endl;
    std::cout << "  Write Capacity: " << ptd.GetWriteCapacityUnits() << std::endl;

    const Aws::Vector<Aws::DynamoDB::Model::AttributeDefinition>& ad =
    td.GetAttributeDefinitions();
    std::cout << "Attributes" << std::endl;
    for (const auto& a : ad)
        std::cout << "  " << a.GetAttributeName() << " (" <<
    Aws::DynamoDB::Model::ScalarAttributeTypeMapper::GetNameForScalarAttributeType(a.GetAttributeType())
    <<
```

```
        ")" << std::endl;
    }
    else
    {
        std::cout << "Failed to describe table: " << result.GetError().GetMessage();
    }
}
```

See the [complete example](#) on GitHub.

Modify (Update) a Table

You can modify your table's provisioned throughput values at any time by calling the [DynamoDB client](#) `UpdateTable` method.

Includes

```
#include <aws/core/Aws.h>
#include <aws/core/utils/Outcome.h>
#include <aws/dynamodb/DynamoDBClient.h>
#include <aws/dynamodb/model/ProvisionedThroughput.h>
#include <aws/dynamodb/model/UpdateTableRequest.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration clientConfig;
Aws::DynamoDB::DynamoDBClient dynamoClient(clientConfig);

std::cout << "Updating " << table << " with new provisioned throughput values" <<
    std::endl;
std::cout << "Read capacity : " << rc << std::endl;
std::cout << "Write capacity: " << wc << std::endl;

Aws::DynamoDB::Model::UpdateTableRequest utr;
Aws::DynamoDB::Model::ProvisionedThroughput pt;
pt.WithReadCapacityUnits(rc).WithWriteCapacityUnits(wc);
utr.WithProvisionedThroughput(pt).WithTableName(table);

const Aws::DynamoDB::Model::UpdateTableOutcome& result = dynamoClient.UpdateTable(utr);
if (!result.IsSuccess())
{
    std::cout << result.GetError().GetMessage() << std::endl;
    return 1;
}
std::cout << "Done!" << std::endl;
```

See the [complete example](#).

Delete a Table

Call the [DynamoDB client](#) `DeleteTable` method and pass it the table's name.

Includes

```
#include <aws/core/Aws.h>
#include <aws/core/utils/Outcome.h>
#include <aws/dynamodb/DynamoDBClient.h>
#include <aws/dynamodb/model/DeleteTableRequest.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration clientConfig;
if (!region.empty())
    clientConfig.region = region;
Aws::DynamoDB::DynamoDBClient dynamoClient(clientConfig);

Aws::DynamoDB::Model::DeleteTableRequest dtr;
dtr.SetTableName(table);

const Aws::DynamoDB::Model::DeleteTableOutcome& result = dynamoClient.DeleteTable(dtr);
if (result.IsSuccess())
{
    std::cout << "Table \"" << result.GetResult().GetTableDescription().GetTableName() <<
        " was deleted!\n";
}
else
{
    std::cout << "Failed to delete table: " << result.GetError().GetMessage();
}
```

See the [complete example](#) on GitHub.

More Info

- [Guidelines for Working with Tables](#) in the *Amazon DynamoDB Developer Guide*
- [Working with Tables in DynamoDB](#) in the *Amazon DynamoDB Developer Guide*

Working with Items in DynamoDB

In DynamoDB, an item is a collection of *attributes*, each of which has a *name* and a *value*. An attribute value can be a scalar, set, or document type. For more information, see [Naming Rules and Data Types](#) in the *Amazon DynamoDB Developer Guide*.

Retrieve (Get) an Item from a Table

Call the [DynamoDB client](#) `GetItem` method. Pass it a [GetItemRequest](#) object with the table name and primary key value of the item you want. It returns a [GetItemResult](#) object.

You can use the returned `GetItemResult` object's `GetItem()` method to retrieve an `Aws::Map` of key `Aws::String` and value [AttributeValue](#) pairs associated with the item.

Includes

```
#include <aws/core/Aws.h>
#include <aws/core/utils/Outcome.h>
#include <aws/dynamodb/DynamoDBClient.h>
#include <aws/dynamodb/model/AttributeDefinition.h>
#include <aws/dynamodb/model/GetItemRequest.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration clientConfig;
Aws::DynamoDB::DynamoDBClient dynamoClient(clientConfig);

Aws::DynamoDB::Model::GetItemRequest req;
```

```
if (!projection.empty())
    req.SetProjectionExpression(projection);

Aws::DynamoDB::Model::AttributeValue haskKey;
haskKey.SetS(name);
req.AddKey("Name", haskKey);

req.SetTableName(table);

const Aws::DynamoDB::Model::GetItemOutcome& result = dynamoClient.GetItem(req);
if (result.IsSuccess())
{
    const Aws::Map<Aws::String, Aws::DynamoDB::Model::AttributeValue>& item =
    result.GetResult().GetItem();
    if (item.size() > 0)
    {
        for (const auto& i : item)
            std::cout << i.first << ": " << i.second.GetS() << std::endl;
    }
    else
    {
        std::cout << "No item found with the key " << name << std::endl;
    }
}
else
{
    std::cout << "Failed to get item: " << result.GetError().GetMessage();
}
}
```

See the [complete example](#) on GitHub.

Add a New Item to a Table

Create key `Aws::String` and value `AttributeValue` pairs that represent each item. These must include values for the table's primary key fields. If the item identified by the primary key already exists, its fields are *updated* by the request. Add them to the `PutItemRequest` using the `AddItem` method.

Includes

```
#include <aws/core/Aws.h>
#include <aws/core/utils/Outcome.h>
#include <aws/dynamodb/DynamoDBClient.h>
#include <aws/dynamodb/model/AttributeDefinition.h>
#include <aws/dynamodb/model/PutItemRequest.h>
#include <aws/dynamodb/model/PutItemResult.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration clientConfig;
Aws::DynamoDB::DynamoDBClient dynamoClient(clientConfig);

Aws::DynamoDB::Model::PutItemRequest pir;
pir.SetTableName(table);

Aws::DynamoDB::Model::AttributeValue av;
av.SetS(name);
pir.AddItem("Name", av);

for (int x = 3; x < argc; x++)
```

```
{
    const Aws::String arg(argv[x]);
    const Aws::Vector<Aws::String>& flds = Aws::Utils::StringUtils::Split(arg, ':');
    if (flds.size() == 2)
    {
        Aws::DynamoDB::Model::AttributeValue val;
        val.SetS(flds[1]);
        pir.AddItem(flds[0], val);
    }
    else
    {
        std::cout << "Invalid argument: " << arg << std::endl << USAGE;
        return 1;
    }
}

const Aws::DynamoDB::Model::PutItemOutcome result = dynamoClient.PutItem(pir);
if (!result.IsSuccess())
{
    std::cout << result.GetError().GetMessage() << std::endl;
    return 1;
}
std::cout << "Done!" << std::endl;
```

See the [complete example](#) on GitHub.

Update an Existing Item in a Table

You can update an attribute for an item that already exists in a table by using the [DynamoDBClient's UpdateItem](#) method, providing a table name, primary key value, and fields to update and their corresponding value.

Imports

```
#include <aws/core/Aws.h>
#include <aws/core/Utils/Outcome.h>
#include <aws/dynamodb/DynamoDBClient.h>
#include <aws/dynamodb/model/UpdateItemRequest.h>
#include <aws/dynamodb/model/UpdateItemResult.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration clientConfig;
Aws::DynamoDB::DynamoDBClient dynamoClient(clientConfig);

Aws::DynamoDB::Model::UpdateItemRequest uir;
uir.SetTableName(table);

Aws::DynamoDB::Model::AttributeValue av;
av.SetS(name);
uir.AddKey("Name", av);

for (int x = 3; x < argc; x++)
{
    const Aws::String arg(argv[x]);
    const Aws::Vector<Aws::String>& flds = Aws::Utils::StringUtils::Split(arg, ':');
    if (flds.size() == 2)
    {
        Aws::DynamoDB::Model::AttributeValue val;
        val.SetS(flds[1]);
        Aws::DynamoDB::Model::AttributeValueUpdate avu;
```



```
        avu.SetValue(val);
        uir.AddAttributeUpdates(flds[0], avu);
    }
    else
    {
        std::cout << "Invalid argument: " << arg << std::endl << USAGE;
        return 1;
    }
}

const Aws::DynamoDB::Model::UpdateItemOutcome& result = dynamoClient.UpdateItem(uir);
if (!result.IsSuccess())
{
    std::cout << result.GetError().GetMessage() << std::endl;
    return 1;
}
std::cout << "Done!" << std::endl;
```

See the [complete example](#).

More Info

- [Guidelines for Working with Items](#) in the *Amazon DynamoDB Developer Guide*
- [Working with Items in DynamoDB](#) in the *Amazon DynamoDB Developer Guide*

Amazon EC2 Examples Using the AWS SDK for C++

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable computing capacity—literally servers in Amazon's data centers—that you use to build and host your software systems. You can use the following examples to program [Amazon EC2](#) using the [AWS SDK for C++](#).

Note

Only the code that is necessary to demonstrate each technique is supplied here, but [complete example code is available on GitHub](#), where you can download a single source file or you can clone the repository locally to get all examples, build and run them.

Topics

- [Managing Amazon EC2 Instances \(p. 45\)](#)
- [Using Elastic IP Addresses in Amazon EC2 \(p. 52\)](#)
- [Using Regions and Availability Zones for Amazon EC2 \(p. 55\)](#)
- [Working with Amazon EC2 Key Pairs \(p. 57\)](#)
- [Working with Security Groups in Amazon EC2 \(p. 59\)](#)

Managing Amazon EC2 Instances

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Creating an Instance

Create a new Amazon EC2 instance by calling the [EC2Client's RunInstances](#) function, providing it with a [RunInstancesRequest](#) containing the [Amazon Machine Image \(AMI\)](#) to use and an [instance type](#).

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/CreateTagsRequest.h>
#include <aws/ec2/model/RunInstancesRequest.h>
#include <aws/ec2/model/RunInstancesResponse.h>
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;

Aws::EC2::Model::RunInstancesRequest run_request;
run_request.SetImageId(ami_id);
run_request.SetInstanceType(Aws::EC2::Model::InstanceType::t1_micro);
run_request.SetMinCount(1);
run_request.SetMaxCount(1);

auto run_outcome = ec2.RunInstances(run_request);
if (!run_outcome.IsSuccess())
{
    std::cout << "Failed to start ec2 instance " << instanceName <<
        " based on ami " << ami_id << ":" <<
        run_outcome.GetError().GetMessage() << std::endl;
    return;
}

const auto& instances = run_outcome.GetResult().GetInstances();
if (instances.size() == 0)
{
    std::cout << "Failed to start ec2 instance " << instanceName <<
        " based on ami " << ami_id << ":" <<
        run_outcome.GetError().GetMessage() << std::endl;
    return;
}
```

See the [complete example](#).

Starting an Instance

To start an Amazon EC2 instance, call the [EC2Client's StartInstances](#) function, providing it with a [StartInstancesRequest](#) containing the ID of the instance to start.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/StartInstancesRequest.h>
#include <aws/ec2/model/StartInstancesResponse.h>
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;

Aws::EC2::Model::StartInstancesRequest start_request;
start_request.AddInstanceIds(instance_id);
start_request.SetDryRun(true);

auto dry_run_outcome = ec2.StartInstances(start_request);
```

```
assert(!dry_run_outcome.IsSuccess());
if (dry_run_outcome.GetError().GetErrorType() !=
    Aws::EC2::EC2Errors::DRY_RUN_OPERATION)
{
    std::cout << "Failed dry run to start instance " << instance_id << ": "
              << dry_run_outcome.GetError().GetMessage() << std::endl;
    return;
}

start_request.SetDryRun(false);
auto start_instancesOutcome = ec2.StartInstances(start_request);

if (!start_instancesOutcome.IsSuccess())
{
    std::cout << "Failed to start instance " << instance_id << ": " <<
              start_instancesOutcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully started instance " << instance_id <<
              std::endl;
}
```

See the [complete example](#).

Stopping an Instance

To stop an Amazon EC2 instance, call the [EC2Client's](#) `StopInstances` function, providing it with a [StopInstancesRequest](#) containing the ID of the instance to stop.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/StopInstancesRequest.h>
#include <aws/ec2/model/StopInstancesResponse.h>
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;
Aws::EC2::Model::StopInstancesRequest request;
request.AddInstanceIds(instance_id);
request.SetDryRun(true);

auto dry_run_outcome = ec2.StopInstances(request);
assert(!dry_run_outcome.IsSuccess());

if (dry_run_outcome.GetError().GetErrorType() !=
    Aws::EC2::EC2Errors::DRY_RUN_OPERATION)
{
    std::cout << "Failed dry run to stop instance " << instance_id << ": "
              << dry_run_outcome.GetError().GetMessage() << std::endl;
    return;
}

request.SetDryRun(false);
auto outcome = ec2.StopInstances(request);
if (!outcome.IsSuccess())
{
    std::cout << "Failed to stop instance " << instance_id << ": " <<
              outcome.GetError().GetMessage() << std::endl;
}
```

```
}  
else  
{  
    std::cout << "Successfully stopped instance " << instance_id <<  
        std::endl;  
}
```

See the [complete example](#).

Rebooting an Instance

To reboot an Amazon EC2 instance, call the [EC2Client](#)'s `RebootInstances` function, providing it with a [RebootInstancesRequest](#) containing the ID of the instance to reboot.

Includes

```
#include <aws/core/Aws.h>  
#include <aws/ec2/EC2Client.h>  
#include <aws/ec2/model/RebootInstancesRequest.h>  
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;  
  
Aws::EC2::Model::RebootInstancesRequest request;  
request.AddInstanceIds(instanceId);  
request.SetDryRun(true);  
  
auto dry_run_outcome = ec2.RebootInstances(request);  
assert(!dry_run_outcome.IsSuccess());  
  
if (dry_run_outcome.GetError().GetErrorType()  
    != Aws::EC2::EC2Errors::DRY_RUN_OPERATION)  
{  
    std::cout << "Failed dry run to reboot instance " << instanceId << ": "  
        << dry_run_outcome.GetError().GetMessage() << std::endl;  
    return;  
}  
  
request.SetDryRun(false);  
auto outcome = ec2.RebootInstances(request);  
if (!outcome.IsSuccess())  
{  
    std::cout << "Failed to reboot instance " << instanceId << ": " <<  
        outcome.GetError().GetMessage() << std::endl;  
}  
else  
{  
    std::cout << "Successfully rebooted instance " << instanceId <<  
        std::endl;  
}
```

See the [complete example](#).

Describing Instances

To list your instances, create a [DescribeInstancesRequest](#) and call the [EC2Client](#)'s `DescribeInstances` function. It will return a [DescribeInstancesResponse](#) object that you can use to list the Amazon EC2 instances for your account and region.

Instances are grouped by *reservation*. Each reservation corresponds to the call to `StartInstances` that launched the instance. To list your instances, you must first call the `DescribeInstancesResponse` class' `GetReservations` function, and then call `getInstances` on each returned `Reservation` object.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/DescribeInstancesRequest.h>
#include <aws/ec2/model/DescribeInstancesResponse.h>
#include <iomanip>
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;
Aws::EC2::Model::DescribeInstancesRequest request;
bool header = false;
bool done = false;
while (!done)
{
    auto outcome = ec2.DescribeInstances(request);
    if (outcome.IsSuccess())
    {
        if (!header)
        {
            std::cout << std::left <<
                std::setw(48) << "Name" <<
                std::setw(20) << "ID" <<
                std::setw(15) << "Ami" <<
                std::setw(15) << "Type" <<
                std::setw(15) << "State" <<
                std::setw(15) << "Monitoring" << std::endl;
            header = true;
        }

        const auto &reservations =
            outcome.GetResult().GetReservations();

        for (const auto &reservation : reservations)
        {
            const auto &instances = reservation.GetInstances();
            for (const auto &instance : instances)
            {
                Aws::String instanceStateString =
                    Aws::EC2::Model::InstanceStateNameMapper::GetNameForInstanceStateName(
                        instance.GetState().GetName());

                Aws::String type_string =
                    Aws::EC2::Model::InstanceTypeMapper::GetNameForInstanceType(
                        instance.GetInstanceType());

                Aws::String monitor_str =
                    Aws::EC2::Model::MonitoringStateMapper::GetNameForMonitoringState(
                        instance.GetMonitoring().GetState());
                Aws::String name = "Unknown";

                const auto &tags = instance.GetTags();
                auto nameIter = std::find_if(tags.cbegin(), tags.cend(),
                    [](const Aws::EC2::Model::Tag &tag)
                    {
                        return tag.GetKey() == "Name";
                    });
            }
        }
    }
}
```

```
        if (nameIter != tags.cend())
        {
            name = nameIter->GetValue();
        }
        std::cout <<
            std::setw(48) << name <<
            std::setw(20) << instance.GetInstanceId() <<
            std::setw(15) << instance.GetImageId() <<
            std::setw(15) << type_string <<
            std::setw(15) << instanceStateString <<
            std::setw(15) << monitor_str << std::endl;
    }
}

if (outcome.GetResult().GetNextToken().size() > 0)
{
    request.SetNextToken(outcome.GetResult().GetNextToken());
}
else
{
    done = true;
}
}
else
{
    std::cout << "Failed to describe ec2 instances:" <<
        outcome.GetError().GetMessage() << std::endl;
    done = true;
}
}
```

Results are paged; you can get further results by passing the value returned from the result object's `GetNextToken` function to your original request object's `SetNextToken` function, then using the same request object in your next call to `DescribeInstances`.

See the [complete example](#).

Enable Instance Monitoring

You can monitor various aspects of your Amazon EC2 instances, such as CPU and network utilization, available memory, and disk space remaining. To learn more about instance monitoring, see [Monitoring Amazon EC2](#) in the *Amazon EC2 User Guide for Linux Instances*.

To start monitoring an instance, you must create a `MonitorInstancesRequest` with the ID of the instance to monitor, and pass it to the `EC2Client`'s `MonitorInstances` function.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/MonitorInstancesRequest.h>
#include <aws/ec2/model/MonitorInstancesResponse.h>
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;
Aws::EC2::Model::MonitorInstancesRequest request;
request.AddInstanceIds(instance_id);
request.SetDryRun(true);

auto dry_run_outcome = ec2.MonitorInstances(request);
```

```
assert(!dry_run_outcome.IsSuccess());
if (dry_run_outcome.GetError().GetErrorType()
    != Aws::EC2::EC2Errors::DRY_RUN_OPERATION)
{
    std::cout << "Failed dry run to enable monitoring on instance " <<
        instance_id << ": " << dry_run_outcome.GetError().GetMessage() <<
        std::endl;
    return;
}

request.SetDryRun(false);
auto monitorInstancesOutcome = ec2.MonitorInstances(request);
if (!monitorInstancesOutcome.IsSuccess())
{
    std::cout << "Failed to enable monitoring on instance " <<
        instance_id << ": " <<
        monitorInstancesOutcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully enabled monitoring on instance " <<
        instance_id << std::endl;
}
}
```

See the [complete example](#).

Disable Instance Monitoring

To stop monitoring an instance, create an [UnmonitorInstancesRequest](#) with the ID of the instance to stop monitoring, and pass it to the [EC2Client](#)'s `UnmonitorInstances` function.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/UnmonitorInstancesRequest.h>
#include <aws/ec2/model/UnmonitorInstancesResponse.h>
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;
Aws::EC2::Model::UnmonitorInstancesRequest unrequest;
unrequest.AddInstanceIds(instance_id);
unrequest.SetDryRun(true);

auto undry_run_outcome = ec2.UnmonitorInstances(unrequest);
assert(!undry_run_outcome.IsSuccess());
if (undry_run_outcome.GetError().GetErrorType() !=
    Aws::EC2::EC2Errors::DRY_RUN_OPERATION)
{
    std::cout << "Failed dry run to disable monitoring on instance " <<
        instance_id << ": " << undry_run_outcome.GetError().GetMessage() <<
        std::endl;
    return;
}

unrequest.SetDryRun(false);
auto unmonitorInstancesOutcome = ec2.UnmonitorInstances(unrequest);
if (!unmonitorInstancesOutcome.IsSuccess())
{
    std::cout << "Failed to disable monitoring on instance " << instance_id
```

```
        << " : " << unmonitorInstancesOutcome.GetError().GetMessage() <<
        std::endl;
    }
    else
    {
        std::cout << "Successfully disable monitoring on instance " <<
        instance_id << std::endl;
    }
}
```

See the [complete example](#).

More Information

- [RunInstances](#) in the *Amazon EC2 API Reference*
- [DescribeInstances](#) in the *Amazon EC2 API Reference*
- [StartInstances](#) in the *Amazon EC2 API Reference*
- [StopInstances](#) in the *Amazon EC2 API Reference*
- [RebootInstances](#) in the *Amazon EC2 API Reference*
- [DescribeInstances](#) in the *Amazon EC2 API Reference*
- [MonitorInstances](#) in the *Amazon EC2 API Reference*
- [UnmonitorInstances](#) in the *Amazon EC2 API Reference*

Using Elastic IP Addresses in Amazon EC2

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Allocating an Elastic IP Address

To use an Elastic IP address, you first allocate one to your account, and then associate it with your instance or a network interface.

To allocate an Elastic IP address, call the [EC2Client's](#) `AllocateAddress` function with an [AllocateAddressRequest](#) object containing the network type (classic EC2 or VPC).

The [AllocateAddressResponse](#) class in the response object contains an allocation ID that you can use to associate the address with an instance, by passing the allocation ID and instance ID in a [AssociateAddressRequest](#) to the [EC2Client's](#) `AssociateAddress` function.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/AllocateAddressRequest.h>
#include <aws/ec2/model/AllocateAddressResponse.h>
#include <aws/ec2/model/AssociateAddressRequest.h>
#include <aws/ec2/model/AssociateAddressResponse.h>
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;
```



```
Aws::EC2::Model::AllocateAddressRequest request;
request.SetDomain(Aws::EC2::Model::DomainType::vpc);

auto outcome = ec2.AllocateAddress(request);
if (!outcome.IsSuccess())
{
    std::cout << "Failed to allocate elastic ip address:" <<
        outcome.GetError().GetMessage() << std::endl;
    return;
}

Aws::String allocation_id = outcome.GetResult().GetAllocationId();

Aws::EC2::Model::AssociateAddressRequest associate_request;
associate_request.SetInstanceId(instance_id);
associate_request.SetAllocationId(allocation_id);

auto associate_outcome = ec2.AssociateAddress(associate_request);
if (!associate_outcome.IsSuccess())
{
    std::cout << "Failed to associate elastic ip address" << allocation_id
        << " with instance " << instance_id << ":" <<
        associate_outcome.GetError().GetMessage() << std::endl;
    return;
}

std::cout << "Successfully associated elastic ip address " << allocation_id
    << " with instance " << instance_id << std::endl;
```

See the [complete example](#).

Describing Elastic IP Addresses

To list the Elastic IP addresses assigned to your account, call the `EC2Client`'s `DescribeAddresses` function. It returns an outcome object that contains a `DescribeAddressesResponse` which you can use to get a list of `Address` objects that represent the Elastic IP addresses on your account.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/DescribeAddressesRequest.h>
#include <aws/ec2/model/DescribeAddressesResponse.h>
#include <iomanip>
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;
Aws::EC2::Model::DescribeAddressesRequest request;
auto outcome = ec2.DescribeAddresses(request);
if (outcome.IsSuccess())
{
    std::cout << std::left << std::setw(20) << "InstanceId" <<
        std::setw(15) << "Public IP" << std::setw(10) << "Domain" <<
        std::setw(20) << "Allocation ID" << std::setw(25) <<
        "NIC ID" << std::endl;

    const auto &addresses = outcome.GetResult().GetAddresses();
    for (const auto &address : addresses)
    {
```

```
Aws::String domainString =
    Aws::EC2::Model::DomainTypeMapper::GetNameForDomainType(
        address.GetDomain());

std::cout << std::left << std::setw(20) <<
    address.GetInstanceId() << std::setw(15) <<
    address.GetPublicIp() << std::setw(10) << domainString <<
    std::setw(20) << address.GetAllocationId() << std::setw(25)
    << address.GetNetworkInterfaceId() << std::endl;
}
}
else
{
    std::cout << "Failed to describe elastic ip addresses:" <<
        outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Releasing an Elastic IP Address

To release an Elastic IP address, call the [EC2Client](#)'s `ReleaseAddress` function, passing it a [ReleaseAddressRequest](#) containing the allocation ID of the Elastic IP address you want to release.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/ReleaseAddressRequest.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration config;
config.region = Aws::Region::US_WEST_2;

Aws::EC2::EC2Client ec2(config);

Aws::EC2::Model::ReleaseAddressRequest request;
request.SetAllocationId(allocation_id);

auto outcome = ec2.ReleaseAddress(request);
if (!outcome.IsSuccess())
{
    std::cout << "Failed to release elastic ip address " <<
        allocation_id << ":" << outcome.GetError().GetMessage() <<
        std::endl;
}
else
{
    std::cout << "Successfully released elastic ip address " <<
        allocation_id << std::endl;
}
```

After you release an Elastic IP address, it is released to the AWS IP address pool and might be unavailable to you afterward. Be sure to update your DNS records and any servers or devices that communicate with the address. If you attempt to release an Elastic IP address that you already released, you'll get an *AuthFailure* error if the address is already allocated to another AWS account.

If you are using *EC2-Classic* or a *default VPC*, then releasing an Elastic IP address automatically disassociates it from any instance that it's associated with. To disassociate an Elastic IP address without releasing it, use the [EC2Client](#)'s `DisassociateAddress` function.

If you are using a non-default VPC, you *must* use `DisassociateAddress` to disassociate the Elastic IP address before you try to release it. Otherwise, Amazon EC2 returns an error (`InvalidIPAddress.InUse`).

See the [complete example](#).

More Information

- [Elastic IP Addresses](#) in the *Amazon EC2 User Guide for Linux Instances*
- [AllocateAddress](#) in the *Amazon EC2 API Reference*
- [DescribeAddresses](#) in the *Amazon EC2 API Reference*
- [ReleaseAddress](#) in the *Amazon EC2 API Reference*

Using Regions and Availability Zones for Amazon EC2

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Describing Regions

To list the regions available to your account, call the `EC2Client`'s `DescribeRegions` function with a `DescribeRegionsRequest`.

You will receive a `DescribeRegionsResponse` in the outcome object. Call its `GetRegions` function to get a list of `Region` objects that represent each region.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/DescribeRegionsRequest.h>
#include <aws/ec2/model/DescribeRegionsResponse.h>
#include <iomanip>
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;
Aws::EC2::Model::DescribeRegionsRequest request;
auto outcome = ec2.DescribeRegions(request);
if (outcome.IsSuccess())
{
    std::cout << std::left <<
        std::setw(32) << "RegionName" <<
        std::setw(64) << "Endpoint" << std::endl;

    const auto &regions = outcome.GetResult().GetRegions();
    for (const auto &region : regions)
    {
        std::cout << std::left <<
            std::setw(32) << region.GetRegionName() <<
            std::setw(64) << region.GetEndpoint() << std::endl;
    }
}
else
{
```

```
std::cout << "Failed to describe regions:" <<  
    outcome.GetError().GetMessage() << std::endl;
```

See the [complete example](#).

Describing Availability Zones

To list each availability zone available to your account, call the `EC2Client`'s `DescribeAvailabilityZones` function with a `DescribeAvailabilityZonesRequest`.

You will receive a `DescribeAvailabilityZonesResponse` in the outcome object. Call its `GetAvailabilityZones` function to get a list of `AvailabilityZone` objects that represent each availability zone.

Includes

```
#include <aws/core/Aws.h>  
#include <aws/ec2/EC2Client.h>  
#include <aws/ec2/model/DescribeAvailabilityZonesRequest.h>  
#include <aws/ec2/model/DescribeAvailabilityZonesResponse.h>  
#include <iomanip>  
#include <iostream>
```

Code

```
std::cout << std::endl;  
Aws::EC2::Model::DescribeAvailabilityZonesRequest describe_request;  
auto describe_outcome = ec2.DescribeAvailabilityZones(describe_request);  
  
if (describe_outcome.IsSuccess())  
{  
    std::cout << std::left <<  
        std::setw(32) << "ZoneName" <<  
        std::setw(20) << "State" <<  
        std::setw(32) << "Region" << std::endl;  
  
    const auto &zones =  
        describe_outcome.GetResult().GetAvailabilityZones();  
  
    for (const auto &zone : zones)  
    {  
        Aws::String stateString =  
            Aws::EC2::Model::AvailabilityZoneStateMapper::GetNameForAvailabilityZoneState(  
                zone.GetState());  
        std::cout << std::left <<  
            std::setw(32) << zone.GetZoneName() <<  
            std::setw(20) << stateString <<  
            std::setw(32) << zone.GetRegionName() << std::endl;  
    }  
}  
else  
{  
    std::cout << "Failed to describe availability zones:" <<  
        describe_outcome.GetError().GetMessage() << std::endl;
```

See the [complete example](#).

More Information

- [Regions and Availability Zones](#) in the *Amazon EC2 User Guide for Linux Instances*

- [DescribeRegions](#) in the *Amazon EC2 API Reference*
- [DescribeAvailabilityZones](#) in the *Amazon EC2 API Reference*

Working with Amazon EC2 Key Pairs

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++](#) (p. 2) and have configured default AWS credentials using the information in [Providing AWS Credentials](#) (p. 5).

Creating a Key Pair

To create a key pair, call the [EC2Client](#)'s `CreateKeyPair` function with a [CreateKeyPairRequest](#) that contains the key's name.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/CreateKeyPairRequest.h>
#include <aws/ec2/model/CreateKeyPairResponse.h>
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;
Aws::EC2::Model::CreateKeyPairRequest request;
request.SetKeyName(pair_name);

auto outcome = ec2.CreateKeyPair(request);
if (!outcome.IsSuccess())
{
    std::cout << "Failed to create key pair:" <<
        outcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully created key pair named " <<
        pair_name << std::endl;
}
```

See the [complete example](#).

Describing Key Pairs

To list your key pairs or to get information about them, call the [EC2Client](#)'s `DescribeKeyPairs` function with a [DescribeKeyPairsRequest](#).

You will receive a [DescribeKeyPairsResponse](#) that you can use to access the list of key pairs by calling its `GetKeyPairs` function, which returns a list of [KeyPairInfo](#) objects.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/DescribeKeyPairsRequest.h>
#include <aws/ec2/model/DescribeKeyPairsResponse.h>
#include <iomanip>
```

```
#include <iostream>
```

Code

```
const auto &key_pairs = outcome.GetResult().GetKeyPairs();
for (const auto &key_pair : key_pairs)
{
    std::cout << std::left <<
        std::setw(32) << key_pair.GetKeyName() <<
        std::setw(64) << key_pair.GetKeyFingerprint() << std::endl;
}
else
{
    std::cout << "Failed to describe key pairs:" <<
        outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Deleting a Key Pair

To delete a key pair, call the [EC2Client](#)'s `DeleteKeyPair` function, passing it a [DeleteKeyPairRequest](#) that contains the name of the key pair to delete.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/DeleteKeyPairRequest.h>
#include <iostream>
```

Code

```
Aws::EC2::Model::DeleteKeyPairRequest request;

request.SetKeyName(pair_name);
auto outcome = ec2.DeleteKeyPair(request);

if (!outcome.IsSuccess())
{
    std::cout << "Failed to delete key pair " << pair_name <<
        ":" << outcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully deleted key pair named " << pair_name <<
        std::endl;
}
```

See the [complete example](#).

More Information

- [Amazon EC2 Key Pairs](#) in the *Amazon EC2 User Guide for Linux Instances*
- [CreateKeyPair](#) in the *Amazon EC2 API Reference*
- [DescribeKeyPairs](#) in the *Amazon EC2 API Reference*
- [DeleteKeyPair](#) in the *Amazon EC2 API Reference*

Working with Security Groups in Amazon EC2

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Creating a Security Group

To create a security group, call the `EC2Client`'s `CreateSecurityGroup` function with a `CreateSecurityGroupRequest` that contains the key's name.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/CreateSecurityGroupRequest.h>
#include <aws/ec2/model/CreateSecurityGroupResponse.h>
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;
Aws::EC2::Model::CreateSecurityGroupRequest request;

request.SetGroupName(group_name);
request.SetDescription(description);
request.SetVpcId(vpc_id);

auto outcome = ec2.CreateSecurityGroup(request);

if (!outcome.IsSuccess())
{
    std::cout << "Failed to create security group:" <<
        outcome.GetError().GetMessage() << std::endl;
    return;
}

std::cout << "Successfully created security group named " << group_name <<
    std::endl;
```

See the [complete example](#).

Configuring a Security Group

A security group can control both inbound (ingress) and outbound (egress) traffic to your Amazon EC2 instances.

To add ingress rules to your security group, use the `EC2Client`'s `AuthorizeSecurityGroupIngress` function, providing the name of the security group and the access rules ([IpPermission](#)) you want to assign to it within an `AuthorizeSecurityGroupIngressRequest` object. The following example shows how to add IP permissions to a security group.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/CreateSecurityGroupRequest.h>
#include <aws/ec2/model/CreateSecurityGroupResponse.h>
```

```
#include <aws/ec2/model/AuthorizeSecurityGroupIngressRequest.h>
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;
authorize_request.SetGroupName(group_name);

BuildSampleIngressRule(authorize_request);
Aws::EC2::Model::IpRange ip_range;
ip_range.SetCidrIp("0.0.0.0/0");

Aws::EC2::Model::IpPermission permission1;
permission1.SetIpProtocol("tcp");
permission1.SetToPort(80);
permission1.SetFromPort(80);
permission1.AddIpRanges(ip_range);

authorize_request.AddIpPermissions(permission1);

Aws::EC2::Model::IpPermission permission2;
permission2.SetIpProtocol("tcp");
permission2.SetToPort(22);
permission2.SetFromPort(22);
permission2.AddIpRanges(ip_range);

authorize_request.AddIpPermissions(permission2);
auto ingress_request = ec2.AuthorizeSecurityGroupIngress(
    authorize_request);

if (!ingress_request.IsSuccess())
{
    std::cout << "Failed to set ingress policy for security group " <<
        group_name << ":" << ingress_request.GetError().GetMessage() <<
        std::endl;
    return;
}

std::cout << "Successfully added ingress policy to security group " <<
    group_name << std::endl;
```

To add an egress rule to the security group, provide similar data in an [AuthorizeSecurityGroupEgressRequest](#) to the [EC2Client](#)'s `AuthorizeSecurityGroupEgress` function.

See the [complete example](#).

Describing Security Groups

To describe your security groups or get information about them, call the [EC2Client](#)'s `DescribeSecurityGroups` function with a [DescribeSecurityGroupsRequest](#).

You will receive a [DescribeSecurityGroupsResponse](#) in the outcome object that you can use to access the list of security groups by calling its `GetSecurityGroups` function, which returns a list of [SecurityGroup](#) objects.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/DescribeSecurityGroupsRequest.h>
#include <aws/ec2/model/DescribeSecurityGroupsResponse.h>
#include <iomanip>
```



```
#include <iostream>
```

Code

```
Aws::EC2::EC2Client ec2;
Aws::EC2::Model::DescribeSecurityGroupsRequest request;

if (argc == 2)
{
    request.AddGroupIds(argv[1]);
}

auto outcome = ec2.DescribeSecurityGroups(request);

if (outcome.IsSuccess())
{
    std::cout << std::left <<
        std::setw(32) << "Name" <<
        std::setw(20) << "GroupId" <<
        std::setw(20) << "VpcId" <<
        std::setw(64) << "Description" << std::endl;

    const auto &securityGroups =
        outcome.GetResult().GetSecurityGroups();

    for (const auto &securityGroup : securityGroups)
    {
        std::cout << std::left <<
            std::setw(32) << securityGroup.GetGroupName() <<
            std::setw(20) << securityGroup.GetGroupId() <<
            std::setw(20) << securityGroup.GetVpcId() <<
            std::setw(64) << securityGroup.GetDescription() <<
            std::endl;
    }
}
else
{
    std::cout << "Failed to describe security groups:" <<
        outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Deleting a Security Group

To delete a security group, call the [EC2Client](#)'s `DeleteSecurityGroup` function, passing it a [DeleteSecurityGroupRequest](#) that contains the ID of the security group to delete.

Includes

```
#include <aws/core/Aws.h>
#include <aws/ec2/EC2Client.h>
#include <aws/ec2/model/DeleteSecurityGroupRequest.h>
#include <iomanip>
```

Code

```
Aws::EC2::EC2Client ec2;
Aws::EC2::Model::DeleteSecurityGroupRequest request;

request.SetGroupId(groupId);
```

```
auto outcome = ec2.DeleteSecurityGroup(request);

if (!outcome.IsSuccess())
{
    std::cout << "Failed to delete security group " << groupId <<
        ":" << outcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully deleted security group " << groupId <<
        std::endl;
}
```

See the [complete example](#).

More Information

- [Amazon EC2 Security Groups](#) in the *Amazon EC2 User Guide for Linux Instances*
- [Authorizing Inbound Traffic for Your Linux Instances](#) in the *Amazon EC2 User Guide for Linux Instances*
- [CreateSecurityGroup](#) in the *Amazon EC2 API Reference*
- [DescribeSecurityGroups](#) in the *Amazon EC2 API Reference*
- [DeleteSecurityGroup](#) in the *Amazon EC2 API Reference*
- [AuthorizeSecurityGroupIngress](#) in the *Amazon EC2 API Reference*

IAM Code Examples Using the AWS SDK for C++

AWS Identity and Access Management (IAM) is a web service for securely controlling access to AWS services. You can use the following examples to program IAM using the [AWS SDK for C++](#).

Note

Only the code that is necessary to demonstrate each technique is supplied here, but [complete example code is available on GitHub](#), where you can download a single source file or you can clone the repository locally to get all examples, build and run them.

Topics

- [Managing IAM Access Keys](#) (p. 62)
- [Managing IAM Users](#) (p. 66)
- [Using IAM Account Aliases](#) (p. 70)
- [Working with IAM Policies](#) (p. 72)
- [Working with IAM Server Certificates](#) (p. 78)

Managing IAM Access Keys

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++](#) (p. 2) and have configured default AWS credentials using the information in [Providing AWS Credentials](#) (p. 5).

Creating an Access Key

To create an IAM access key, call the [IAMClient](#)'s `CreateAccessKey` function with an `CreateAccessKeyRequest` object.

You must set the user name using the `CreateAccessKeyRequest`'s `WithUserName` setter function before passing it to the `CreateAccessKey` function.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/CreateAccessKeyRequest.h>
#include <aws/iam/model/CreateAccessKeyResult.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;

Aws::IAM::Model::CreateAccessKeyRequest request;
request.SetUserName(user_name);

auto outcome = iam.CreateAccessKey(request);
if (!outcome.IsSuccess())
{
    std::cout << "Error creating access key for IAM user " << user_name
              << ":" << outcome.GetError().GetMessage() << std::endl;
}
else
{
    const auto &accessKey = outcome.GetResult().GetAccessKey();
    std::cout << "Successfully created access key for IAM user " <<
              user_name << std::endl << "  aws_access_key_id = " <<
              accessKey.GetAccessKeyId() << std::endl <<
              "  aws_secret_access_key = " << accessKey.GetSecretAccessKey() <<
              std::endl;
}
}
```

See the [complete example](#).

Listing Access Keys

To list the access keys for a given user, create a [ListAccessKeysRequest](#) object that contains the user name to list keys for, and pass it to the [IAMClient](#)'s `ListAccessKeys` function.

Note

If you do not supply a user name to `ListAccessKeys`, it will attempt to list access keys associated with the AWS account that signed the request.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/ListAccessKeysRequest.h>
#include <aws/iam/model/ListAccessKeysResult.h>
#include <iomanip>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::ListAccessKeysRequest request;
request.SetUserName(userName);

bool done = false;
```

```
bool header = false;
while (!done)
{
    auto outcome = iam.ListAccessKeys(request);
    if (!outcome.IsSuccess())
    {
        std::cout << "Failed to list access keys for user " << userName
            << ": " << outcome.GetError().GetMessage() << std::endl;
        break;
    }

    if (!header)
    {
        std::cout << std::left << std::setw(32) << "UserName" <<
            std::setw(30) << "KeyID" << std::setw(20) << "Status" <<
            std::setw(20) << "CreateDate" << std::endl;
        header = true;
    }

    const auto &keys = outcome.GetResult().GetAccessKeyMetadata();
    for (const auto &key : keys)
    {
        Aws::String statusString =
            Aws::IAM::Model::StatusTypeMapper::GetNameForStatusType(
                key.GetStatus());
        std::cout << std::left << std::setw(32) << key.GetUserName() <<
            std::setw(30) << key.GetAccessKeyId() << std::setw(20) <<
            statusString << std::setw(20) <<
            key.GetCreateDate().ToGmtString(DATE_FORMAT) << std::endl;
    }

    if (outcome.GetResult().GetIsTruncated())
    {
        request.SetMarker(outcome.GetResult().GetMarker());
    }
    else
    {
        done = true;
    }
}
```

The results of `ListAccessKeys` are paged (with a default maximum of 100 records per call). You can call `GetIsTruncated` on the returned [ListAccessKeysResult](#) object to see if the query returned fewer results than are available. If so, then call `SetMarker` on the `ListAccessKeysRequest` and pass it back to the next invocation of `ListAccessKeys`.

See the [complete example](#).

Retrieving an Access Key's Last Used Time

To get the time an access key was last used, call the `IAMClient`'s `GetAccessKeyLastUsed` function with the access key's ID (which can be passed in using a [GetAccessKeyLastUsedRequest](#) object, or directly to the overload that takes the access key ID directly).

You can then use the returned [GetAccessKeyLastUsedResult](#) object to retrieve the key's last used time.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/GetAccessKeyLastUsedRequest.h>
#include <aws/iam/model/GetAccessKeyLastUsedResult.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::GetAccessKeyLastUsedRequest request;

request.SetAccessKeyId(key_id);

auto outcome = iam.GetAccessKeyLastUsed(request);

if (!outcome.IsSuccess())
{
    std::cout << "Error querying last used time for access key " <<
        key_id << ":" << outcome.GetError().GetMessage() << std::endl;
}
else
{
    auto lastUsedTimeString =
        outcome.GetResult()
            .GetAccessKeyLastUsed()
            .GetLastUsedDate()
            .ToGmtString(Aws::Utils::DateFormat::ISO_8601);
    std::cout << "Access key " << key_id << " last used at time " <<
        lastUsedTimeString << std::endl;
}
```

See the [complete example](#).

Activating or Deactivating Access Keys

You can activate or deactivate an access key by creating an [UpdateAccessKeyRequest](#) object, providing the access key ID, optionally the user name, and the desired `Status` Type, then passing the request object to the [IAMClient](#)'s `UpdateAccessKey` function.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/UpdateAccessKeyRequest.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::UpdateAccessKeyRequest request;
request.SetUserName(user_name);
request.SetAccessKeyId(accessKeyId);
request.SetStatus(status);

auto outcome = iam.UpdateAccessKey(request);
if (outcome.IsSuccess())
{
    std::cout << "Successfully updated status of access key " <<
        accessKeyId << " for user " << user_name << std::endl;
}
else
{
    std::cout << "Error updated status of access key " << accessKeyId <<
        " for user " << user_name << ": " <<
        outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Deleting an Access Key

To permanently delete an access key, call the `IAMClient`'s `DeleteKey` function, providing it with a `DeleteAccessKeyRequest` containing the access key's ID and username.

Note

Once deleted, a key can no longer be retrieved or used. To temporarily deactivate a key so that it can be activated again later, use [updateAccessKey \(p. 65\)](#) function instead.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/DeleteAccessKeyRequest.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;

Aws::IAM::Model::DeleteAccessKeyRequest request;
request.SetUserName(user_name);
request.SetAccessKeyId(key_id);

auto outcome = iam.DeleteAccessKey(request);

if (!outcome.IsSuccess())
{
    std::cout << "Error deleting access key " << key_id << " from user "
              << user_name << ": " << outcome.GetError().GetMessage() <<
              std::endl;
}
else
{
    std::cout << "Successfully deleted access key " << key_id
              << " for IAM user " << user_name << std::endl;
}
}
```

See the [complete example](#).

More Information

- [CreateAccessKey](#) in the *IAM API Reference*
- [ListAccessKeys](#) in the *IAM API Reference*
- [GetAccessKeyLastUsed](#) in the *IAM API Reference*
- [UpdateAccessKey](#) in the *IAM API Reference*
- [DeleteAccessKey](#) in the *IAM API Reference*

Managing IAM Users

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Create a User

Use the `IAMClient::CreateUser` function, passing it a `CreateUserRequest` with the name of the user to create.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/CreateUserRequest.h>
#include <aws/iam/model/CreateUserResult.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::CreateUserRequest create_request;
create_request.SetUserName(user_name);

auto create_outcome = iam.CreateUser(create_request);
if (!create_outcome.IsSuccess())
{
    std::cout << "Error creating IAM user " << user_name << " " <<
        create_outcome.GetError().GetMessage() << std::endl;
    return;
}
std::cout << "Successfully created IAM user " << user_name << std::endl;
```

Get Information About a User

To get information about a particular user, such as the user's creation date, path, ID or ARN, call the `IAMClient::GetUser` function with a `GetUserRequest` containing the user name. If successful, you can get the `User` from the returned `GetUserResult` outcome.

If the user doesn't already exist, `GetUser` will fail with `Aws::IAM::IAMErrors::NO_SUCH_ENTITY`.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/GetUserRequest.h>
#include <aws/iam/model/GetUserResult.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::GetUserRequest get_request;
get_request.SetUserName(user_name);

auto get_outcome = iam.GetUser(get_request);
if (get_outcome.IsSuccess())
{
    std::cout << "IAM user " << user_name << " already exists" << std::endl;
    return;
}
else if (get_outcome.GetError().GetErrorType() !=
    Aws::IAM::IAMErrors::NO_SUCH_ENTITY)
{
```

```
std::cout << "Error checking existence of IAM user " << user_name << ":"  
    << get_outcome.GetError().GetMessage() << std::endl;  
return;  
}
```

See the [complete example](#).

Listing Users

List the existing IAM users for your account by calling the `IAMClient::ListUsers` function, passing it a `ListUsersRequest` object. The list of users is returned in a `ListUsersResult` object that you can use to get information about the users.

The result may be paginated; to check to see if there are more results available, check the value of `GetResult().GetIsTruncated()`. If `true`, then set a marker on the request and call `ListUsers` again to get the next batch of users. This code demonstrates the technique.

Includes:

```
#include <aws/core/Aws.h>  
#include <aws/iam/IAMClient.h>  
#include <aws/iam/model/ListUsersRequest.h>  
#include <aws/iam/model/ListUsersResult.h>  
#include <iomanip>  
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;  
Aws::IAM::Model::ListUsersRequest request;  
  
bool done = false;  
bool header = false;  
while (!done)  
{  
    auto outcome = iam.ListUsers(request);  
    if (!outcome.IsSuccess())  
    {  
        std::cout << "Failed to list iam users:" <<  
            outcome.GetError().GetMessage() << std::endl;  
        break;  
    }  
  
    if (!header)  
    {  
        std::cout << std::left << std::setw(32) << "Name" <<  
            std::setw(30) << "ID" << std::setw(64) << "Arn" <<  
            std::setw(20) << "CreateDate" << std::endl;  
        header = true;  
    }  
  
    const auto &users = outcome.GetResult().GetUsers();  
    for (const auto &user : users)  
    {  
        std::cout << std::left << std::setw(32) << user.GetUserName() <<  
            std::setw(30) << user.GetUserId() << std::setw(64) <<  
            user.GetArn() << std::setw(20) <<  
            user.GetCreateDate().ToGmtString(DATE_FORMAT) << std::endl;  
    }  
  
    if (outcome.GetResult().GetIsTruncated())
```



```
{
    request.SetMarker(outcome.GetResult().GetMarker());
}
else
{
    done = true;
}
```

See the [complete example](#).

Update a User

To update an existing user, create an [UpdateUserRequest](#) and pass it to the `IAMClientUpdateUser` member function.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/UpdateUserRequest.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;

Aws::IAM::Model::UpdateUserRequest request;
request.SetUserName(old_name);
request.SetNewUserName(new_name);

auto outcome = iam.UpdateUser(request);
if (outcome.IsSuccess())
{
    std::cout << "IAM user " << old_name <<
        " successfully updated with new user name " << new_name <<
        std::endl;
}
else
{
    std::cout << "Error updating user name for IAM user " << old_name <<
        ":" << outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Delete a User

To delete an existing user, call the `IAMClientDeleteUser` function, passing it a [DeleteUserRequest](#) object containing the name of the user to delete.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/DeleteUserRequest.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::DeleteUserRequest request;
request.SetUserName(user_name);
auto outcome = iam.DeleteUser(request);
if (!outcome.IsSuccess())
{
    std::cout << "Error deleting IAM user " << user_name << ": " <<
        outcome.GetError().GetMessage() << std::endl;
    return;
}
std::cout << "Successfully deleted IAM user " << user_name << std::endl;
```

See the [complete example](#).

Using IAM Account Aliases

If you want the URL for your sign-in page to contain your company name or other friendly identifier instead of your AWS account ID, you can create an alias for your AWS account.

Note

AWS supports exactly one account alias per account.

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Creating an Account Alias

To create an account alias, call the `IAMClient`'s `CreateAccountAlias` function with a `CreateAccountAliasRequest` object that contains the alias name.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/CreateAccountAliasRequest.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::CreateAccountAliasRequest request;
request.SetAccountAlias(alias_name);

auto outcome = iam.CreateAccountAlias(request);
if (!outcome.IsSuccess())
{
    std::cout << "Error creating account alias " << alias_name << ": " <<
        outcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully created account alias " << alias_name <<
        std::endl;
}
```

See the [complete example](#).

Listing Account Aliases

To list your account's alias, if any, call the `IAMClient`'s `ListAccountAliases` function. It takes a `ListAccountAliasesRequest` object.

Note

The returned `ListAccountAliasesResult` supports the same `GetIsTruncated` and `GetMarker` functions as other AWS SDK for C++ *list* functions, but an AWS account can have only *one* account alias.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/ListAccountAliasesRequest.h>
#include <aws/iam/model/ListAccountAliasesResult.h>
#include <iomanip>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::ListAccountAliasesRequest request;

bool done = false;
bool header = false;
while (!done)
{
    auto outcome = iam.ListAccountAliases(request);
    if (!outcome.IsSuccess())
    {
        std::cout << "Failed to list account aliases: " <<
            outcome.GetError().GetMessage() << std::endl;
        break;
    }

    const auto &aliases = outcome.GetResult().GetAccountAliases();
    if (!header)
    {
        if (aliases.size() == 0)
        {
            std::cout << "Account has no aliases" << std::endl;
            break;
        }
        std::cout << std::left << std::setw(32) << "Alias" << std::endl;
        header = true;
    }

    for (const auto &alias : aliases)
    {
        std::cout << std::left << std::setw(32) << alias << std::endl;
    }

    if (outcome.GetResult().GetIsTruncated())
    {
        request.SetMarker(outcome.GetResult().GetMarker());
    }
    else
    {
        done = true;
    }
}
```

see the [complete example](#).

Deleting an account alias

To delete your account's alias, call the `IAMClient`'s `DeleteAccountAlias` function. When deleting an account alias, you must supply its name using a `DeleteAccountAliasRequest` object.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/DeleteAccountAliasRequest.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;

Aws::IAM::Model::DeleteAccountAliasRequest request;
request.SetAccountAlias(alias_name);

const auto outcome = iam.DeleteAccountAlias(request);
if (!outcome.IsSuccess())
{
    std::cout << "Error deleting account alias " << alias_name << ": "
              << outcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully deleted account alias " << alias_name <<
              << std::endl;
}
```

See the [complete example](#).

More Information

- [Your AWS Account ID and Its Alias](#) in the *IAM User Guide*
- [CreateAccountAlias](#) in the *IAM API Reference*
- [ListAccountAliases](#) in the *IAM API Reference*
- [DeleteAccountAlias](#) in the *IAM API Reference*

Working with IAM Policies

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Creating a Policy

To create a new policy, provide the policy's name and a JSON-formatted policy document in a `CreatePolicyRequest` to the `IAMClient`'s `CreatePolicy` function.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/CreatePolicyRequest.h>
#include <aws/iam/model/CreatePolicyResult.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;

Aws::IAM::Model::CreatePolicyRequest request;
request.SetPolicyName(policy_name);
request.SetPolicyDocument(BuildSamplePolicyDocument(rsrc_arn));

auto outcome = iam.CreatePolicy(request);
if (!outcome.IsSuccess())
{
    std::cout << "Error creating policy " << policy_name << ": " <<
        outcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully created policy " << policy_name <<
        std::endl;
}
```

IAM policy documents are JSON strings with a [well-documented syntax](#). Here is an example that provides access to make particular requests to DynamoDB. It takes the policy ARN as a passed-in variable.

```
static const char* const POLICY_TEMPLATE =
"{
  \"Version\": \"2012-10-17\",
  \"Statement\": [
    {
      \"Effect\": \"Allow\",
      \"Action\": \"logs:CreateLogGroup\",
      \"Resource\": \"%s\"
    },
    {
      \"Effect\": \"Allow\",
      \"Action\": [
        \"dynamodb:DeleteItem\",
        \"dynamodb:GetItem\",
        \"dynamodb:PutItem\",
        \"dynamodb:Scan\",
        \"dynamodb:UpdateItem\"
      ],
      \"Resource\": \"%s\"
    }
  ]
}";

Aws::String BuildSamplePolicyDocument(const Aws::String& rsrc_arn)
{
    char policyBuffer[512];
#ifdef WIN32
    sprintf_s(policyBuffer, POLICY_TEMPLATE, rsrc_arn.c_str(), rsrc_arn.c_str());
#else
    sprintf(policyBuffer, POLICY_TEMPLATE, rsrc_arn.c_str(), rsrc_arn.c_str());
#endif // WIN32
    return Aws::String(policyBuffer);
}
```

See the [complete example](#).

Getting a Policy

To retrieve an existing policy, call the `IAMClient`'s `GetPolicy` function, providing the policy's ARN within a `GetPolicyRequest` object.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/GetPolicyRequest.h>
#include <aws/iam/model/GetPolicyResult.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::GetPolicyRequest request;
request.SetPolicyArn(policy_arn);

auto outcome = iam.GetPolicy(request);
if (!outcome.IsSuccess())
{
    std::cout << "Error getting policy " << policy_arn << ": " <<
        outcome.GetError().GetMessage() << std::endl;
}
else
{
    const auto &policy = outcome.GetResult().GetPolicy();
    std::cout << "Name: " << policy.GetPolicyName() << std::endl <<
        "ID: " << policy.GetPolicyId() << std::endl << "Arn: " <<
        policy.GetArn() << std::endl << "Description: " <<
        policy.GetDescription() << std::endl << "CreateDate: " <<
        policy.GetCreateDate().ToGmtString(Aws::Utils::DateFormat::ISO_8601)
        << std::endl;
}
}
```

See the [complete example](#).

Deleting a Policy

To delete a policy, provide the policy's ARN in a `DeletePolicyRequest` to the `IAMClient`'s `DeletePolicy` function.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/DeletePolicyRequest.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::DeletePolicyRequest request;
request.SetPolicyArn(policy_arn);

auto outcome = iam.DeletePolicy(request);
```

```
if (!outcome.IsSuccess())
{
    std::cout << "Error deleting policy with arn " << policy_arn << ": "
                << outcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully deleted policy with arn " << policy_arn
                << std::endl;
}
```

See the [complete example](#).

Attaching a Policy

You can attach a policy to an IAMrole by calling the IAMClient's AttachRolePolicy function, providing it with the role name and policy ARN in an AttachRolePolicyRequest.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/AttachRolePolicyRequest.h>
#include <aws/iam/model/ListAttachedRolePoliciesRequest.h>
#include <aws/iam/model/ListAttachedRolePoliciesResult.h>
#include <iostream>
#include <iomanip>
```

Code:

```
Aws::IAM::IAMClient iam;

Aws::IAM::Model::ListAttachedRolePoliciesRequest list_request;
list_request.SetRoleName(role_name);

bool done = false;
while (!done)
{
    auto list_outcome = iam.ListAttachedRolePolicies(list_request);
    if (!list_outcome.IsSuccess())
    {
        std::cout << "Failed to list attached policies of role " <<
                    role_name << ": " << list_outcome.GetError().GetMessage() <<
                    std::endl;
        return;
    }

    const auto& policies = list_outcome.GetResult().GetAttachedPolicies();
    if (std::any_of(policies.cbegin(), policies.cend(),
                   [=](const Aws::IAM::Model::AttachedPolicy& policy)
                   {
                       return policy.GetPolicyArn() == policy_arn;
                   })))
    {
        std::cout << "Policy " << policy_arn <<
                    " is already attached to role " << role_name << std::endl;
        return;
    }

    done = !list_outcome.GetResult().GetIsTruncated();
    list_request.SetMarker(list_outcome.GetResult().GetMarker());
}
```

```
Aws::IAM::Model::AttachRolePolicyRequest request;
request.SetRoleName(role_name);
request.SetPolicyArn(policy_arn);

auto outcome = iam.AttachRolePolicy(request);
if (!outcome.IsSuccess())
{
    std::cout << "Failed to attach policy " << policy_arn << " to role " <<
        role_name << ": " << outcome.GetError().GetMessage() << std::endl;
    return;
}

std::cout << "Successfully attached policy " << policy_arn << " to role " <<
    role_name << std::endl;
```

See the [complete example](#).

Listing Attached Policies

List attached policies on a role by calling the `IAMClient`'s `ListAttachedRolePolicies` function. It takes a `ListAttachedRolePoliciesRequest` object that contains the role name to list the policies for.

Call `GetAttachedPolicies` on the returned `ListAttachedRolePoliciesResult` object to get the list of attached policies. Results may be truncated; if the `ListAttachedRolePoliciesResult` object's `GetIsTruncated` function returns `true`, call the `ListAttachedRolePoliciesRequest` object's `SetMarker` function and use it to call `ListAttachedRolePolicies` again to get the next batch of results.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/ListPoliciesRequest.h>
#include <aws/iam/model/ListPoliciesResult.h>
#include <iomanip>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::ListPoliciesRequest request;

bool done = false;
bool header = false;
while (!done)
{
    auto outcome = iam.ListPolicies(request);
    if (!outcome.IsSuccess())
    {
        std::cout << "Failed to list iam policies: " <<
            outcome.GetError().GetMessage() << std::endl;
        break;
    }

    if (!header)
    {
        std::cout << std::left << std::setw(55) << "Name" <<
            std::setw(30) << "ID" << std::setw(80) << "Arn" <<
            std::setw(64) << "Description" << std::setw(12) <<
            "CreateDate" << std::endl;
    }
}
```



```
        header = true;
    }

    const auto &policies = outcome.GetResult().GetPolicies();
    for (const auto &policy : policies)
    {
        std::cout << std::left << std::setw(55) <<
            policy.GetPolicyName() << std::setw(30) <<
            policy.GetPolicyId() << std::setw(80) << policy.GetArn() <<
            std::setw(64) << policy.GetDescription() << std::setw(12) <<
            policy.GetCreateDate().ToGmtString(DATE_FORMAT) <<
            std::endl;
    }

    if (outcome.GetResult().GetIsTruncated())
    {
        request.SetMarker(outcome.GetResult().GetMarker());
    }
    else
    {
        done = true;
    }
}
```

See the [complete example](#).

Detaching a Policy

To detach a policy from a role, call the `IAMClient`'s `DetachRolePolicy` function, providing it with the role name and policy ARN in a [DetachRolePolicyRequest](#).

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/DetachRolePolicyRequest.h>
#include <aws/iam/model/ListAttachedRolePoliciesRequest.h>
#include <aws/iam/model/ListAttachedRolePoliciesResult.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::DetachRolePolicyRequest detach_request;
detach_request.SetRoleName(role_name);
detach_request.SetPolicyArn(policy_arn);

auto detach_outcome = iam.DetachRolePolicy(detach_request);
if (!detach_outcome.IsSuccess())
{
    std::cout << "Failed to detach policy " << policy_arn << " from role "
        << role_name << ": " << detach_outcome.GetError().GetMessage() <<
        std::endl;
    return;
}
```

See the [complete example](#).

More Information

- [Overview of IAM Policies](#) in the *IAM User Guide*.

- [AWS IAM Policy Reference](#) in the *IAM User Guide*.
- [CreatePolicy](#) in the *IAM API Reference*
- [GetPolicy](#) in the *IAM API Reference*
- [DeletePolicy](#) in the *IAM API Reference*
- [AttachGroupPolicy](#), [AttachRolePolicy](#) and [AttachUserPolicy](#) in the *IAM API Reference*
- [DetachGroupPolicy](#), [DetachRolePolicy](#) and [DetachUserPolicy](#) in the *IAM API Reference*
- [ListAttachedGroupPolicies](#), [ListAttachedRolePolicies](#) and [ListAttachedUserPolicies](#) in the *IAM API Reference*

Working with IAM Server Certificates

To enable HTTPS connections to your website or application on AWS, you need an SSL/TLS *server certificate*. You can use a server certificate provided by AWS Certificate Manager or one that you obtained from an external provider.

We recommend that you use ACM to provision, manage, and deploy your server certificates. With ACM you can request a certificate, deploy it to your AWS resources, and let ACM handle certificate renewals for you. Certificates provided by ACM are free. For more information about ACM, see the [ACM User Guide](#).

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Getting a Server Certificate

You can retrieve a server certificate by calling the `IAMClient`'s `GetServerCertificate` function, passing it a `GetServerCertificateRequest` with the certificate's name.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/GetServerCertificateRequest.h>
#include <aws/iam/model/GetServerCertificateResult.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::GetServerCertificateRequest request;
request.SetServerCertificateName(cert_name);

auto outcome = iam.GetServerCertificate(request);
if (!outcome.IsSuccess())
{
    std::cout << "Error getting server certificate " << cert_name <<
        ": " << outcome.GetError().GetMessage() << std::endl;
}
else
{
    const auto &certificate = outcome.GetResult().GetServerCertificate();
    std::cout << "Name: " <<
        certificate.GetServerCertificateMetadata().GetServerCertificateName()
        << std::endl << "Body: " << certificate.GetCertificateBody() <<
        std::endl << "Chain: " << certificate.GetCertificateChain() <<
        std::endl;
```

```
}
```

See the [complete example](#).

Listing Server Certificates

To list your server certificates, call the `IAMClient`'s `ListServerCertificates` function with a `ListServerCertificatesRequest`. It returns a `ListServerCertificatesResult`.

Call the returned `ListServerCertificateResult` object's `GetServerCertificateMetadataList` function to get a list of `ServerCertificateMetadata` objects that you can use to get information about each certificate.

Results may be truncated; if the `ListServerCertificateResult` object's `GetIsTruncated` function returns `true`, call the `ListServerCertificatesRequest` object's `SetMarker` function and use it to call `listServerCertificates` again to get the next batch of results.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/ListServerCertificatesRequest.h>
#include <aws/iam/model/ListServerCertificatesResult.h>
#include <iomanip>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::ListServerCertificatesRequest request;

bool done = false;
bool header = false;
while (!done)
{
    auto outcome = iam.ListServerCertificates(request);
    if (!outcome.IsSuccess())
    {
        std::cout << "Failed to list server certificates: " <<
            outcome.GetError().GetMessage() << std::endl;
        break;
    }

    if (!header)
    {
        std::cout << std::left << std::setw(55) << "Name" <<
            std::setw(30) << "ID" << std::setw(80) << "Arn" <<
            std::setw(14) << "UploadDate" << std::setw(14) <<
            "ExpirationDate" << std::endl;
        header = true;
    }

    const auto &certificates =
        outcome.GetResult().GetServerCertificateMetadataList();

    for (const auto &certificate : certificates)
    {
        std::cout << std::left << std::setw(55) <<
            certificate.GetServerCertificateName() << std::setw(30) <<
            certificate.GetServerCertificateId() << std::setw(80) <<
```

```
certificate.GetArn() << std::setw(14) <<
certificate.GetUploadDate().ToGmtString(DATE_FORMAT) <<
std::setw(14) <<
certificate.GetExpiration().ToGmtString(DATE_FORMAT) <<
std::endl;
}

if (outcome.GetResult().GetIsTruncated())
{
    request.SetMarker(outcome.GetResult().GetMarker());
}
else
{
    done = true;
}
}
```

See the [complete example](#).

Updating a Server Certificate

You can update a server certificate's name or path by calling the `IAMClient`'s `UpdateServerCertificate` function. It takes a `UpdateServerCertificateRequest` object set with the server certificate's current name and either a new name or new path to use.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/UpdateServerCertificateRequest.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::UpdateServerCertificateRequest request;
request.SetServerCertificateName(old_name);
request.SetNewServerCertificateName(new_name);

auto outcome = iam.UpdateServerCertificate(request);
if (outcome.IsSuccess())
{
    std::cout << "Server certificate " << old_name
              << " successfully renamed as " << new_name
              << std::endl;
}
else
{
    std::cout << "Error changing name of server certificate " <<
              old_name << " to " << new_name << ":" <<
              outcome.GetError().GetMessage() << std::endl;
}
}
```

See the [complete example](#).

Deleting a Server Certificate

To delete a server certificate, call the `IAMClient`'s `DeleteServerCertificate` function with a `DeleteServerCertificateRequest` containing the certificate's name.

Includes:

```
#include <aws/core/Aws.h>
#include <aws/iam/IAMClient.h>
#include <aws/iam/model/DeleteServerCertificateRequest.h>
#include <iostream>
```

Code:

```
Aws::IAM::IAMClient iam;
Aws::IAM::Model::DeleteServerCertificateRequest request;
request.SetServerCertificateName(cert_name);

const auto outcome = iam.DeleteServerCertificate(request);
if (!outcome.IsSuccess())
{
    std::cout << "Error deleting server certificate " << cert_name <<
        ": " << outcome.GetError().GetMessage() << std::endl;
}
else
{
    std::cout << "Successfully deleted server certificate " << cert_name
        << std::endl;
}
```

See the [complete example](#).

More Information

- [Working with Server Certificates](#) in the *IAM User Guide*
- [GetServerCertificate](#) in the *IAM API Reference*
- [ListServerCertificates](#) in the *IAM API Reference*
- [UpdateServerCertificate](#) in the *IAM API Reference*
- [DeleteServerCertificate](#) in the *IAM API Reference*
- [ACM User Guide](#)

Amazon S3 Code Examples Using the AWS SDK for C++

Amazon Simple Storage Service (Amazon S3) is storage for the internet. You can use the following examples to program [Amazon S3](#) using the [AWS SDK for C++](#).

Note

Only the code that is necessary to demonstrate each technique is supplied here, but [complete example code is available on GitHub](#), where you can download a single source file or you can clone the repository locally to get all examples, build and run them.

Topics

- [Creating, Listing, and Deleting Buckets](#) (p. 82)
- [Operations on Objects](#) (p. 84)
- [Managing Amazon S3 Access Permissions for Buckets and Objects](#) (p. 87)
- [Managing Access to Amazon S3 Buckets Using Bucket Policies](#) (p. 90)
- [Configuring an Amazon S3 Bucket as a Website](#) (p. 92)

Creating, Listing, and Deleting Buckets

Every object (file) in Amazon Simple Storage Service must reside within a *bucket*, which represents a collection (container) of objects. Each bucket is known by a *key* (name), which must be unique. For detailed information about buckets and their configuration, see [Working with Amazon S3 Buckets](#) in the *Amazon S3 Developer Guide*.

Note

Best Practice

We recommend that you enable the [AbortIncompleteMultipartUpload](#) lifecycle rule on your Amazon S3 buckets.

This rule directs Amazon S3 to abort multipart uploads that don't complete within a specified number of days after being initiated. When the set time limit is exceeded, Amazon S3 aborts the upload and then deletes the incomplete upload data.

For more information, see [Lifecycle Configuration for a Bucket with Versioning](#) in the *Amazon S3 User Guide*.

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++](#) (p. 2) and have configured default AWS credentials using the information in [Providing AWS Credentials](#) (p. 5).

Create a Bucket

Use the [S3Client](#) object `CreateBucket` method, passing it a `CreateBucketRequest` with the bucket's name.

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/CreateBucketRequest.h>
```

Code

```
Aws::S3::S3Client s3_client;

Aws::S3::Model::CreateBucketRequest request;
request.SetBucket(bucket_name);

auto outcome = s3_client.CreateBucket(request);

if (outcome.IsSuccess())
{
    std::cout << "Done!" << std::endl;
}
else
{
    std::cout << "CreateBucket error: "
              << outcome.GetError().GetExceptionName() << std::endl
              << outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

List Buckets

Use the [S3Client](#) object `ListBucket` method. If successful, the method returns a `ListBucketOutcome` object, which contains a `ListBucketResult` object.

Use the `ListBucketResult` object `GetBuckets` method to get a list of `Bucket` objects that contain information about each Amazon S3 bucket in your account.

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/Bucket.h>
```

Code

```
Aws::S3::S3Client s3_client;
auto outcome = s3_client.ListBuckets();

if (outcome.IsSuccess())
{
    std::cout << "Your Amazon S3 buckets:" << std::endl;

    Aws::Vector<Aws::S3::Model::Bucket> bucket_list =
        outcome.GetResult().GetBuckets();

    for (auto const &bucket : bucket_list)
    {
        std::cout << " * " << bucket.GetName() << std::endl;
    }
}
else
{
    std::cout << "ListBuckets error: "
        << outcome.GetError().GetExceptionName() << " - "
        << outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Delete a Bucket

Use the `S3Client` object `DeleteBucket` method, passing it a `DeleteBucketRequest` object that is set with the name of the bucket to delete. *The bucket must be empty or an error will result.*

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/DeleteBucketRequest.h>
```

Code

```
Aws::Client::ClientConfiguration config;
config.region = user_region;
Aws::S3::S3Client s3_client(config);

Aws::S3::Model::DeleteBucketRequest bucket_request;
bucket_request.SetBucket(bucket_name);

auto outcome = s3_client.DeleteBucket(bucket_request);

if (outcome.IsSuccess())
{
    std::cout << "Done!" << std::endl;
}
```

```
}  
else  
{  
    std::cout << "DeleteBucket error: "  
        << outcome.GetError().GetExceptionName() << " - "  
        << outcome.GetError().GetMessage() << std::endl;  
}
```

See the [complete example](#).

Operations on Objects

An Amazon S3 object represents a *file*, which is a collection of data. Every object must reside within a [bucket](#) (p. 82).

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++](#) (p. 2) and have configured default AWS credentials using the information in [Providing AWS Credentials](#) (p. 5).

Upload an Object

Use the [S3Client](#) object `PutObject` function, supplying it with a bucket name, key name, and file to upload. *The bucket must exist or an error will result.*

Includes

```
#include <aws/core/Aws.h>  
#include <aws/s3/S3Client.h>  
#include <aws/s3/model/PutObjectRequest.h>  
#include <iostream>  
#include <fstream>
```

Code

```
Aws::Client::ClientConfiguration clientConfig;  
if (!region.empty())  
    clientConfig.region = region;  
Aws::S3::S3Client s3_client(clientConfig);  
  
Aws::S3::Model::PutObjectRequest object_request;  
object_request.WithBucket(bucket_name).WithKey(key_name);  
  
// Binary files must also have the std::ios_base::bin flag or'ed in  
auto input_data = Aws::MakeShared<Aws::FStream>("PutObjectInputStream",  
    file_name.c_str(), std::ios_base::in | std::ios_base::binary);  
  
object_request.SetBody(input_data);  
  
auto put_object_outcome = s3_client.PutObject(object_request);  
  
if (put_object_outcome.IsSuccess())  
{  
    std::cout << "Done!" << std::endl;  
}  
else  
{  
    std::cout << "PutObject error: " <<  
        put_object_outcome.GetError().GetExceptionName() << " " <<  
        put_object_outcome.GetError().GetMessage() << std::endl;  
}
```



```
}

```

See the [complete example](#).

List Objects

To get a list of objects within a bucket, use the [S3Client](#) object `ListObjects` function. Supply it with a `ListObjectsRequest` that you set with the name of a bucket to list the contents of.

The `ListObjects` function returns a `ListObjectsOutcome` object that you can use to get a list of objects in the form of `Object` instances.

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/ListObjectsRequest.h>
#include <aws/s3/model/Object.h>

```

Code

```
Aws::S3::S3Client s3_client;

Aws::S3::Model::ListObjectsRequest objects_request;
objects_request.WithBucket(bucket_name);

auto list_objects_outcome = s3_client.ListObjects(objects_request);

if (list_objects_outcome.IsSuccess())
{
    Aws::Vector<Aws::S3::Model::Object> object_list =
        list_objects_outcome.GetResult().GetContents();

    for (auto const &s3_object : object_list)
    {
        std::cout << "*" << s3_object.GetKey() << std::endl;
    }
}
else
{
    std::cout << "ListObjects error: " <<
        list_objects_outcome.GetError().GetExceptionName() << " " <<
        list_objects_outcome.GetError().GetMessage() << std::endl;
}

```

See the [complete example](#).

Download an Object

Use the [S3Client](#) object `GetObject` function, passing it a `GetObjectRequest` that you set with the name of a bucket and the object key to download. `GetObject` returns a `GetObjectOutcome` object that you can use to access the S3 object's data.

The following example downloads an object from Amazon S3 and saves its contents to a file (using the same name as the object's key).

Includes

```
#include <aws/core/Aws.h>

```

```
#include <aws/s3/S3Client.h>
#include <aws/s3/model/GetObjectRequest.h>
#include <fstream>
```

Code

```
Aws::S3::S3Client s3_client;

Aws::S3::Model::GetObjectRequest object_request;
object_request.WithBucket(bucket_name).WithKey(key_name);

auto get_object_outcome = s3_client.GetObject(object_request);

if (get_object_outcome.IsSuccess())
{
    Aws::OFStream local_file;
    local_file.open(key_name.c_str(), std::ios::out | std::ios::binary);
    local_file << get_object_outcome.GetResult().GetBody().rdbuf();
    std::cout << "Done!" << std::endl;
}
else
{
    std::cout << "GetObject error: " <<
        get_object_outcome.GetError().GetExceptionName() << " " <<
        get_object_outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Delete an Object

Use the [S3Client](#) object's `DeleteObject` function, passing it a `DeleteObjectRequest` that you set with the name of a bucket and object to download. *The specified bucket and object key must exist or an error will result.*

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/DeleteObjectRequest.h>
#include <fstream>
```

Code

```
Aws::S3::S3Client s3_client;

Aws::S3::Model::DeleteObjectRequest object_request;
object_request.WithBucket(bucket_name).WithKey(key_name);

auto delete_object_outcome = s3_client.DeleteObject(object_request);

if (delete_object_outcome.IsSuccess())
{
    std::cout << "Done!" << std::endl;
}
else
{
    std::cout << "DeleteObject error: " <<
        delete_object_outcome.GetError().GetExceptionName() << " " <<
        delete_object_outcome.GetError().GetMessage() << std::endl;
}
```

```
}
```

See the [complete example](#).

Managing Amazon S3 Access Permissions for Buckets and Objects

You can use access control lists (ACLs) for Amazon S3 buckets and objects for fine-grained control over your Amazon S3 resources.

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Get the Access Control List for a Bucket

To get the ACL for an Amazon S3 bucket, call the `S3Client`'s `GetBucketAcl` function with a `GetBucketAclRequest`, providing it with the *bucket name*.

Results are returned in an `GetBucketAclResult` that you can use to get the list of `Grants` by calling its `GetGrants` function.

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/GetBucketAclRequest.h>
#include <aws/s3/model/Permission.h>
#include <aws/s3/model/Grant.h>
```

Code

```
Aws::Client::ClientConfiguration config;
config.region = user_region;
Aws::S3::S3Client s3_client(config);

Aws::S3::Model::GetBucketAclRequest request;
request.SetBucket(bucket_name);

auto outcome = s3_client.GetBucketAcl(request);

if (outcome.IsSuccess())
{
    Aws::Vector<Aws::S3::Model::Grant> grants =
        outcome.GetResult().GetGrants();
    for (auto it = grants.begin(); it != grants.end(); it++)
    {
        Aws::S3::Model::Grant grant = *it;
        std::cout << grant.GetGrantee().GetDisplayName() << ": "
            << GetPermissionString(grant.GetPermission())
            << std::endl;
    }
}
else
{
    std::cout << "GetBucketAcl error: "
        << outcome.GetError().GetExceptionName() << " - "
        << outcome.GetError().GetMessage() << std::endl;
}
```

```
}
```

See the [complete example](#).

Set the Access Control List for a Bucket

To set the ACL for a bucket, call the `S3Client`'s `PutBucketAcl` function, passing it a `PutBucketAclRequest` object with the bucket name and list of grantees and permissions within an `AccessControlPolicy` object.

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/AccessControlPolicy.h>
#include <aws/s3/model/GetBucketAclRequest.h>
#include <aws/s3/model/PutBucketAclRequest.h>
#include <aws/s3/model/Grantee.h>
#include <aws/s3/model/Permission.h>
```

Code

```
Aws::Client::ClientConfiguration config;
config.region = user_region;
Aws::S3::S3Client s3_client(config);

Aws::S3::Model::GetBucketAclRequest get_request;
get_request.SetBucket(bucket_name);

auto get_outcome = s3_client.GetBucketAcl(get_request);

if (get_outcome.IsSuccess())
{
    Aws::S3::Model::Grantee grantee;
    grantee.SetEmailAddress(email);
    Aws::S3::Model::PutBucketAclRequest put_request;
    put_request.SetBucket(bucket_name);
    s3_client.PutBucketAcl(put_request);
}
else
{
    std::cout << "GetBucketAcl error: "
                << get_outcome.GetError().GetExceptionName() << " - "
                << get_outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Get the Access Control List for an Object

To get the ACL for an Amazon S3 object, call the `S3Client`'s `GetObjectAcl` function with a `GetObjectAclRequest`, providing it with the *bucket name* and *object key*.

Results are returned in an `GetObjectAclResult` that you can use to get the list of `Grants` by calling its `GetGrants` function.

Includes

```
#include <aws/core/Aws.h>
```

```
#include <aws/s3/S3Client.h>
#include <aws/s3/model/GetObjectAclRequest.h>
#include <aws/s3/model/Permission.h>
#include <aws/s3/model/Grant.h>
```

Code

```
Aws::Client::ClientConfiguration config;
config.region = user_region;
Aws::S3::S3Client s3_client(config);

Aws::S3::Model::GetObjectAclRequest request;
request.SetBucket(bucket_name);
request.SetKey(object_key);

auto outcome = s3_client.GetObjectAcl(request);

if (outcome.IsSuccess())
{
    Aws::Vector<Aws::S3::Model::Grant> grants =
        outcome.GetResult().GetGrants();
    for (auto it = grants.begin(); it != grants.end(); it++)
    {
        Aws::S3::Model::Grant grant = *it;
        std::cout << grant.GetGrantee().GetDisplayName() << ": "
            << GetPermissionString(grant.GetPermission())
            << std::endl;
    }
}
else
{
    std::cout << "GetObjectAcl error: "
        << outcome.GetError().GetExceptionName() << " - "
        << outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Set the Access Control List for an Object

To set the ACL for an object, call the [S3Client's PutObjectAcl](#) function, passing it a [PutObjectAclRequest](#) object with the object name and list of grantees and permissions within an [AccessControlPolicy](#) object.

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/AccessControlPolicy.h>
#include <aws/s3/model/GetObjectAclRequest.h>
#include <aws/s3/model/PutObjectAclRequest.h>
#include <aws/s3/model/Grantee.h>
#include <aws/s3/model/Permission.h>
```

Code

```
Aws::Client::ClientConfiguration config;
config.region = user_region;
Aws::S3::S3Client s3_client(config);
```

```
Aws::S3::Model::GetBucketAclRequest get_request;
get_request.SetBucket(bucket_name);

auto get_outcome = s3_client.GetBucketAcl(get_request);

if (get_outcome.IsSuccess())
{
    Aws::S3::Model::Grantee grantee;
    grantee.SetEmailAddress(email);
    Aws::S3::Model::PutBucketAclRequest put_request;
    put_request.SetBucket(bucket_name);
    s3_client.PutBucketAcl(put_request);
}
else
{
    std::cout << "GetBucketAcl error: "
                << get_outcome.GetError().GetExceptionName() << " - "
                << get_outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

More Information

- [GET Bucket acl](#) in the *Amazon S3 API Reference*
- [PUT Bucket acl](#) in the *Amazon S3 API Reference*
- [GET Object acl](#) in the *Amazon S3 API Reference*
- [PUT Object acl](#) in the *Amazon S3 API Reference*

Managing Access to Amazon S3 Buckets Using Bucket Policies

You can set, get, or delete a *bucket policy* to manage access to your Amazon S3 buckets.

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Set a Bucket Policy

You can set the bucket policy for a particular S3 bucket by calling the [S3Client's PutBucketPolicy](#) function and providing it with the bucket name and policy's JSON representation in a [PutBucketPolicyRequest](#).

Includes

```
#include <cstdio>
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/PutBucketPolicyRequest.h>
```

Code

```
Aws::Client::ClientConfiguration config;
```

```
config.region = user_region;
Aws::S3::S3Client s3_client(config);

auto request_body = Aws::MakeShared<Aws::StringStream>("");
st_body << policy_string;

Aws::S3::Model::PutBucketPolicyRequest request;
request.SetBucket(bucket_name);
request.SetBody(request_body);

auto outcome = s3_client.PutBucketPolicy(request);

if (outcome.IsSuccess()) {
    std::cout << "Done!" << std::endl;
} else {
    std::cout << "SetBucketPolicy error: "
                << outcome.GetError().GetExceptionName() << std::endl
                << outcome.GetError().GetMessage() << std::endl;
}
}
```

Note

The [Aws::Utils::Json::JsonValue](#) utility class can be used to help you construct valid JSON objects to pass to `PutBucketPolicy`.

See the [complete example](#).

Get a Bucket Policy

To retrieve the policy for an Amazon S3 bucket, call the [S3Client's](#) `GetBucketPolicy` function, passing it the name of the bucket in a [GetBucketPolicyRequest](#).

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/GetBucketPolicyRequest.h>
```

Code

```
Aws::Client::ClientConfiguration config;
config.region = user_region;
Aws::S3::S3Client s3_client(config);

Aws::S3::Model::GetBucketPolicyRequest request;
request.SetBucket(bucket_name);

auto outcome = s3_client.GetBucketPolicy(request);

if (outcome.IsSuccess())
{
    Aws::StringStream policyStream;
    Aws::String line;
    while (outcome.GetResult().GetPolicy())
    {
        outcome.GetResult().GetPolicy() >> line;
        policyStream << line;
    }
    std::cout << "Policy: " << std::endl << policyStream.str() << std::endl;
}
else
{
}
```

```
std::cout << "GetBucketPolicy error: " <<
outcome.GetError().GetExceptionName() << std::endl <<
outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Delete a Bucket Policy

To delete a bucket policy, call the [S3Client](#)'s `DeleteBucketPolicy` function, providing it with the bucket name in a [DeleteBucketPolicyRequest](#).

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/DeleteBucketPolicyRequest.h>
```

Code

```
Aws::S3::S3Client s3_client(config);

Aws::S3::Model::DeleteBucketPolicyRequest request;
request.SetBucket(bucket_name);

auto outcome = s3_client.DeleteBucketPolicy(request);

if (outcome.IsSuccess())
{
    std::cout << "Done!" << std::endl;
}
else
{
    std::cout << "DeleteBucketPolicy error: "
        << outcome.GetError().GetExceptionName() << " - "
        << outcome.GetError().GetMessage() << std::endl;
}
```

This function succeeds even if the bucket doesn't already have a policy. If you specify a bucket name that doesn't exist or if you don't have access to the bucket, an `AmazonServiceException` is thrown.

See the [complete example](#).

More Info

- [PutBucketPolicy](#) in the *Amazon S3 API Reference*
- [GetBucketPolicy](#) in the *Amazon S3 API Reference*
- [DeleteBucketPolicy](#) in the *Amazon S3 API Reference*
- [Access Policy Language Overview](#) in the *Amazon S3 Developer Guide*
- [Bucket Policy Examples](#) in the *Amazon S3 Developer Guide*

Configuring an Amazon S3 Bucket as a Website

You can configure an Amazon S3 bucket to behave as a website. To do this, you need to set its website configuration.

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Set a Bucket's Website Configuration

To set an Amazon S3 bucket's website configuration, call the [S3Client](#)'s `PutBucketWebsite` function with a [PutBucketWebsiteRequest](#) object containing the bucket name and its website configuration, provided in a [WebsiteConfiguration](#) object.

Setting an index document is *required*; all other parameters are optional.

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/IndexDocument.h>
#include <aws/s3/model/ErrorDocument.h>
#include <aws/s3/model/WebsiteConfiguration.h>
#include <aws/s3/model/PutBucketWebsiteRequest.h>
```

Code

```
Aws::Client::ClientConfiguration config;
config.region = user_region;
Aws::S3::S3Client s3_client(config);

Aws::S3::Model::IndexDocument index_doc;
index_doc.SetSuffix(index_suffix);

Aws::S3::Model::ErrorDocument error_doc;
error_doc.SetKey(error_key);

Aws::S3::Model::WebsiteConfiguration website_config;
website_config.SetIndexDocument(index_doc);
website_config.SetErrorDocument(error_doc);

Aws::S3::Model::PutBucketWebsiteRequest request;
request.SetBucket(bucket_name);
request.SetWebsiteConfiguration(website_config);

auto outcome = s3_client.PutBucketWebsite(request);

if (outcome.IsSuccess())
{
    std::cout << "Done!" << std::endl;
}
else
{
    std::cout << "PutBucketWebsite error: "
              << outcome.GetError().GetExceptionName() << std::endl
              << outcome.GetError().GetMessage() << std::endl;
}
```

Note

Setting a website configuration does not modify the access permissions for your bucket. To make your files visible on the web, you will also need to set a *bucket policy* that allows public read access to the files in the bucket. For more information, see [Managing Access to Amazon S3 Buckets Using Bucket Policies \(p. 90\)](#).

See the [complete example](#).

Get a Bucket's Website Configuration

To get an Amazon S3 bucket's website configuration, call the `S3Client`'s `GetBucketWebsite` function with a `GetBucketWebsiteRequest` containing the name of the bucket to retrieve the configuration for.

The configuration will be returned as a `GetBucketWebsiteResult` object within the outcome object. If there is no website configuration for the bucket, then `null` will be returned.

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/GetBucketWebsiteRequest.h>
```

Code

```
Aws::Client::ClientConfiguration config;
config.region = user_region;
Aws::S3::S3Client s3_client(config);

Aws::S3::Model::GetBucketWebsiteRequest request;
request.SetBucket(bucket_name);

auto outcome = s3_client.GetBucketWebsite(request);

if (outcome.IsSuccess())
{
    std::cout << "  Index page: "
              << outcome.GetResult().GetIndexDocument().GetSuffix()
              << std::endl
              << "  Error page: "
              << outcome.GetResult().GetErrorDocument().GetKey()
              << std::endl;
}
else
{
    std::cout << "GetBucketWebsite error: "
              << outcome.GetError().GetExceptionName() << " - "
              << outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Delete a Bucket's Website Configuration

To delete an Amazon S3 bucket's website configuration, call the `S3Client`'s `DeleteBucketWebsite` function with a `DeleteBucketWebsiteRequest`: containing the name of the bucket to delete the configuration from.

Includes

```
#include <aws/core/Aws.h>
#include <aws/s3/S3Client.h>
#include <aws/s3/model/DeleteBucketWebsiteRequest.h>
```

Code

```
Aws::Client::ClientConfiguration config;
```

```
config.region = user_region;
Aws::S3::S3Client s3_client(config);

Aws::S3::Model::DeleteBucketWebsiteRequest request;
request.SetBucket(bucket_name);

auto outcome = s3_client.DeleteBucketWebsite(request);

if (outcome.IsSuccess())
{
    std::cout << "Done!" << std::endl;
}
else
{
    std::cout << "DeleteBucketWebsite error: "
        << outcome.GetError().GetExceptionName() << std::endl
        << outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

More Information

- [PUT Bucket website](#) in the *Amazon S3 API Reference*
- [GET Bucket website](#) in the *Amazon S3 API Reference*
- [DELETE Bucket website](#) in the *Amazon S3 API Reference*

Amazon SQS Code Examples Using the AWS SDK for C++

Amazon Simple Queue Service (Amazon SQS) is a fully managed message queuing service that makes it easy to decouple and scale microservices, distributed systems, and serverless applications. You can use the following examples to program [Amazon SQS](#) using the [AWS SDK for C++](#).

Note

Only the code that is necessary to demonstrate each technique is supplied here, but [complete example code is available on GitHub](#), where you can download a single source file or you can clone the repository locally to get all examples, build and run them.

Topics

- [Working with Amazon SQS Message Queues](#) (p. 95)
- [Sending, Receiving, and Deleting Amazon SQS Messages](#) (p. 98)
- [Enabling Long Polling for Amazon SQS Message Queues](#) (p. 100)
- [Setting Visibility Timeout in Amazon SQS](#) (p. 103)
- [Using Dead Letter Queues in Amazon SQS](#) (p. 104)

Working with Amazon SQS Message Queues

A *message queue* is the logical container you use to send messages reliably in Amazon SQS. There are two types of queues: *standard* and *first-in, first-out* (FIFO). To learn more about queues and the differences between these types, see the [Amazon SQS Developer Guide](#).

These C++ examples show you how to use the AWS SDK for C++ to create, list, delete, and get the URL of an Amazon SQS queue.

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Create a Queue

Use the `SQSClient` class `CreateQueue` member function, and provide it with a `CreateQueueRequest` object that describes the queue parameters.

Includes

```
#include <aws/core/Aws.h>
#include <aws/sqs/SQSClient.h>
#include <aws/sqs/model/CreateQueueRequest.h>
#include <aws/sqs/model/CreateQueueResult.h>
#include <iostream>
```

Code

```
Aws::SQS::SQSClient sqs;

Aws::SQS::Model::CreateQueueRequest cq_req;
cq_req.SetQueueName(queue_name);

auto cq_out = sqs.CreateQueue(cq_req);
if (cq_out.IsSuccess())
{
    std::cout << "Successfully created queue " << queue_name << std::endl;
}
else
{
    std::cout << "Error creating queue " << queue_name << ": " <<
        cq_out.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

List Queues

To list Amazon SQS queues for your account, call the `SQSClient` class `ListQueues` member function, and pass it a `ListQueuesRequest` object.

Includes

```
#include <aws/core/Aws.h>
#include <aws/sqs/SQSClient.h>
#include <aws/sqs/model/ListQueuesRequest.h>
#include <aws/sqs/model/ListQueuesResult.h>
#include <iostream>
```

Code

```
Aws::SQS::SQSClient sqs;

Aws::SQS::Model::ListQueuesRequest lq_req;

auto lq_out = sqs.ListQueues(lq_req);
if (lq_out.IsSuccess())
```

```
{
    std::cout << "Queue Urls:" << std::endl << std::endl;
    const auto &queue_urls = lq_out.GetResult().GetQueueUrls();
    for (const auto &iter : queue_urls)
    {
        std::cout << " " << iter << std::endl;
    }
}
else
{
    std::cout << "Error listing queues: " <<
        lq_out.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Get the Queue's URL

To get the URL for an existing Amazon SQS queue, call the [SQSClient](#) class `GetQueueUrl` member function.

Includes

```
#include <aws/core/Aws.h>
#include <aws/sqs/SQSClient.h>
#include <aws/sqs/model/GetQueueUrlRequest.h>
#include <aws/sqs/model/GetQueueUrlResult.h>
#include <iostream>
```

Code

```
Aws::SQS::SQSClient sqs;

Aws::SQS::Model::GetQueueUrlRequest gqu_req;
gqu_req.SetQueueName(queue_name);

auto gqu_out = sqs.GetQueueUrl(gqu_req);
if (gqu_out.IsSuccess()) {
    std::cout << "Queue " << queue_name << " has url " <<
        gqu_out.GetResult().GetQueueUrl() << std::endl;
} else {
    std::cout << "Error getting url for queue " << queue_name << ": " <<
        gqu_out.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Delete a Queue

Provide the [URL \(p. 97\)](#) to the [SQSClient](#) class `DeleteQueue` member function.

Includes

```
#include <aws/core/Aws.h>
#include <aws/core/client/DefaultRetryStrategy.h>
#include <aws/sqs/SQSClient.h>
#include <aws/sqs/model/DeleteQueueRequest.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration client_cfg;
client_cfg.retryStrategy =
    Aws::MakeShared<Aws::Client::DefaultRetryStrategy>(
        "sqs_delete_queue", 0);
Aws::SQS::SQSClient sqs(client_cfg);

Aws::SQS::Model::DeleteQueueRequest dq_req;
dq_req.SetQueueUrl(queue_url);

auto dq_out = sqs.DeleteQueue(dq_req);
if (dq_out.IsSuccess())
{
    std::cout << "Successfully deleted queue with url " << queue_url <<
        std::endl;
}
else
{
    std::cout << "Error deleting queue " << queue_url << ": " <<
        dq_out.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

More Info

- [How Amazon SQS Queues Work](#) in the *Amazon SQS Developer Guide*
- [CreateQueue](#) in the *Amazon SQS API Reference*
- [GetQueueUrl](#) in the *Amazon SQS API Reference*
- [ListQueues](#) in the *Amazon SQS API Reference*
- [DeleteQueues](#) in the *Amazon SQS API Reference*

Sending, Receiving, and Deleting Amazon SQS Messages

Messages are always delivered using an [SQS queue \(p. 95\)](#). These C++ examples show you how to use the AWS SDK for C++ to send, receive, and delete Amazon SQS messages from SQS queues.

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Send a Message

You can add a single message to an Amazon SQS queue by calling the [SQSClient](#) class `SendMessage` member function. You provide `SendMessage` with a [SendMessageRequest](#) object containing the queue's [URL \(p. 97\)](#), the message body, and an optional delay value (in seconds).

Includes

```
#include <aws/core/Aws.h>
#include <aws/sqs/SQSClient.h>
#include <aws/sqs/model/SendMessageRequest.h>
#include <aws/sqs/model/SendMessageResult.h>
#include <iostream>
```

Code

```
Aws::SQS::SQSClient sqs;

Aws::SQS::Model::SendMessageRequest sm_req;
sm_req.SetQueueUrl(queue_url);
sm_req.SetMessageBody(msg_body);

auto sm_out = sqs.SendMessage(sm_req);
if (sm_out.IsSuccess())
{
    std::cout << "Successfully sent message to " << queue_url <<
        std::endl;
}
else
{
    std::cout << "Error sending message to " << queue_url << ": " <<
        sm_out.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Receive Messages

Retrieve any messages that are currently in the queue by calling the `SQSClient` class `ReceiveMessage` member function, passing it the queue's URL. Messages are returned as a list of `Message` objects.

Includes

```
#include <aws/core/Aws.h>
#include <aws/sqs/SQSClient.h>
#include <aws/sqs/model/ReceiveMessageRequest.h>
#include <aws/sqs/model/ReceiveMessageResult.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration client_cfg;
client_cfg.requestTimeoutMs = 30000;

Aws::SQS::SQSClient sqs(client_cfg);

Aws::SQS::Model::ReceiveMessageRequest rm_req;
rm_req.SetQueueUrl(queue_url);
rm_req.SetMaxNumberOfMessages(1);

auto rm_out = sqs.ReceiveMessage(rm_req);
if (!rm_out.IsSuccess())
{
    std::cout << "Error receiving message from queue " << queue_url << ": "
        << rm_out.GetError().GetMessage() << std::endl;
    return;
}

const auto& messages = rm_out.GetResult().GetMessages();
if (messages.size() == 0)
{
    std::cout << "No messages received from queue " << queue_url <<
        std::endl;
    return;
}
```

```
const auto& message = messages[0];
std::cout << "Received message:" << std::endl;
std::cout << "  MessageId: " << message.GetMessageId() << std::endl;
std::cout << "  ReceiptHandle: " << message.GetReceiptHandle() << std::endl;
std::cout << "  Body: " << message.GetBody() << std::endl << std::endl;
```

See the [complete example](#).

Delete Messages after Receipt

After receiving a message and processing its contents, delete the message from the queue by sending the message's receipt handle and the queue URL to the [SQSClient](#) class `DeleteMessage` member function.

Includes

```
#include <aws/core/Aws.h>
#include <aws/sqs/SQSClient.h>
#include <aws/sqs/model/DeleteMessageRequest.h>
#include <iostream>
```

Code

```
Aws::SQS::Model::DeleteMessageRequest dm_req;
dm_req.SetQueueUrl(queue_url);
dm_req.SetReceiptHandle(message.GetReceiptHandle());

auto dm_out = sqs.DeleteMessage(dm_req);
if (dm_out.IsSuccess())
{
    std::cout << "Successfully deleted message " << message.GetMessageId()
              << " from queue " << queue_url << std::endl;
}
else
{
    std::cout << "Error deleting message " << message.GetMessageId() <<
              " from queue " << queue_url << ": " <<
              dm_out.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

More Info

- [How Amazon SQS Queues Work](#) in the *Amazon SQS Developer Guide*
- [SendMessage](#) in the *Amazon SQS API Reference*
- [SendMessageBatch](#) in the *Amazon SQS API Reference*
- [ReceiveMessage](#) in the *Amazon SQS API Reference*
- [DeleteMessage](#) in the *Amazon SQS API Reference*

Enabling Long Polling for Amazon SQS Message Queues

Amazon SQS uses *short polling* by default, querying only a subset of the servers—based on a weighted random distribution—to determine whether any messages are available for inclusion in the response.

Long polling helps reduce your cost of using Amazon SQS by reducing the number of empty responses when there are no messages available to return in reply to a `ReceiveMessage` request sent to an Amazon SQS queue and eliminating false empty responses. You can set a long polling frequency from *1–20 seconds*.

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Enabling Long Polling when Creating a Queue

To enable long polling when creating an Amazon SQS queue, set the `ReceiveMessageWaitTimeSeconds` attribute on the `CreateQueueRequest` object before calling the `SQSClient` class' `CreateQueue` member function.

Includes

```
#include <aws/core/Aws.h>
#include <aws/sqs/SQSClient.h>
#include <aws/sqs/model/CreateQueueRequest.h>
#include <aws/sqs/model/CreateQueueResult.h>
#include <iostream>
```

Code

```
Aws::SQS::SQSClient sqs;

Aws::SQS::Model::CreateQueueRequest request;
request.SetQueueName(queue_name);
request.AddAttributes(
    Aws::SQS::Model::QueueAttributeName::ReceiveMessageWaitTimeSeconds,
    poll_time);

auto outcome = sqs.CreateQueue(request);
if (outcome.IsSuccess())
{
    std::cout << "Successfully created queue " << queue_name <<
        std::endl;
}
else
{
    std::cout << "Error creating queue " << queue_name << ": " <<
        outcome.GetError().GetMessage() << std::endl;
}
```

See the [complete example](#).

Enabling Long Polling on an Existing Queue

In addition to enabling long polling when creating a queue, you can also enable it on an existing queue by setting `ReceiveMessageWaitTimeSeconds` on the `SetQueueAttributesRequest` before calling the `SQSClient` class' `SetQueueAttributes` member function.

Includes

```
#include <aws/core/Aws.h>
#include <aws/sqs/SQSClient.h>
#include <aws/sqs/model/SetQueueAttributesRequest.h>
#include <iostream>
```

Code

```
Aws::SQS::SQSClient sqs;

Aws::SQS::Model::SetQueueAttributesRequest request;
request.SetQueueUrl(queue_url);
request.AddAttributes(
    Aws::SQS::Model::QueueAttributeName::ReceiveMessageWaitTimeSeconds,
    poll_time);

auto outcome = sqs.SetQueueAttributes(request);
if (outcome.IsSuccess())
{
    std::cout << "Successfully updated long polling time for queue " <<
        queue_url << " to " << poll_time << std::endl;
}
else
{
    std::cout << "Error updating long polling time for queue " <<
        queue_url << ": " << outcome.GetError().GetMessage() <<
        std::endl;
}
```

See the [complete example](#).

Enabling Long Polling on Message Receipt

You can enable long polling when receiving a message by setting the wait time in seconds on the [ReceiveMessageRequest](#) that you supply to the [SQSClient](#) class' `ReceiveMessage` member function.

Note

You should make sure that the AWS client's request timeout is larger than the maximum long poll time (20s) so that your `ReceiveMessage` requests don't time out while waiting for the next poll event!

Includes

```
#include <aws/core/Aws.h>
#include <aws/sqs/SQSClient.h>
#include <aws/sqs/model/ReceiveMessageRequest.h>
#include <aws/sqs/model/ReceiveMessageResult.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration client_cfg;
client_cfg.requestTimeoutMs = 30000;

Aws::SQS::SQSClient sqs(client_cfg);

Aws::SQS::Model::ReceiveMessageRequest request;
request.SetQueueUrl(queue_url);
request.SetMaxNumberOfMessages(1);
request.SetWaitTimeSeconds(wait_time);

auto outcome = sqs.ReceiveMessage(request);
if (!outcome.IsSuccess())
{
    std::cout << "Error receiving message from queue " << queue_url << ": "
        << outcome.GetError().GetMessage() << std::endl;
    return;
}
```

```
}
```

See the [complete example](#).

More Info

- [Amazon SQS Long Polling](#) in the *Amazon SQS Developer Guide*
- [CreateQueue](#) in the *Amazon SQS API Reference*
- [ReceiveMessage](#) in the *Amazon SQS API Reference*
- [SetQueueAttributes](#) in the *Amazon SQS API Reference*

Setting Visibility Timeout in Amazon SQS

When a message is received in Amazon SQS, it remains on the queue until it's deleted in order to ensure receipt. A message that was received, but not deleted, will be available in subsequent requests after a given *visibility timeout* to help prevent the message from being received more than once before it can be processed and deleted.

When using [standard queues](#), visibility timeout isn't a guarantee against receiving a message twice. If you are using a standard queue, be sure that your code can handle the case where the same message has been delivered more than once.

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Setting the Message Visibility Timeout upon Message Receipt

When you have received a message, you can modify its visibility timeout by passing its receipt handle in a [ChangeMessageVisibilityRequest](#) that you pass to the [SQSClient](#) class' `ChangeMessageVisibility` member function.

Includes

```
#include <aws/core/Aws.h>
#include <aws/sqs/SQSClient.h>
#include <aws/sqs/model/ChangeMessageVisibilityRequest.h>
#include <aws/sqs/model/ReceiveMessageRequest.h>
#include <aws/sqs/model/ReceiveMessageResult.h>
#include <iostream>
```

Code

```
Aws::Client::ClientConfiguration client_config;
client_config.requestTimeoutMs = 30000;

Aws::SQS::SQSClient sqs(client_config);

Aws::SQS::Model::ReceiveMessageRequest receive_request;
receive_request.SetQueueUrl(queue_url);
receive_request.SetMaxNumberOfMessages(1);

auto receive_outcome = sqs.ReceiveMessage(receive_request);
if (!receive_outcome.IsSuccess())
{
```

```
std::cout << "Error receiving message from queue " << queue_url << ": "
    << receive_outcome.GetError().GetMessage() << std::endl;
return;
}

const auto& messages = receive_outcome.GetResult().GetMessages();
if (messages.size() == 0)
{
    std::cout << "No messages received from queue " << queue_url <<
        std::endl;
    return;
}

const auto& message = messages[0];
std::cout << "Received message:" << std::endl;
std::cout << "  MessageId: " << message.GetMessageId() << std::endl;
std::cout << "  ReceiptHandle: " << message.GetReceiptHandle() << std::endl;
std::cout << "  Body: " << message.GetBody() << std::endl << std::endl;

Aws::SQS::Model::ChangeMessageVisibilityRequest request;
request.SetQueueUrl(queue_url);
request.SetReceiptHandle(message.GetReceiptHandle());
request.SetVisibilityTimeout(visibility_timeout);
auto outcome = sqs.ChangeMessageVisibility(request);
if (outcome.IsSuccess())
{
    std::cout << "Successfully changed visibility of message " <<
        message.GetMessageId() << " from queue " << queue_url << std::endl;
}
else
{
    std::cout << "Error changing visibility of message " <<
        message.GetMessageId() << " from queue " << queue_url << ": " <<
        outcome.GetError().GetMessage() << std::endl;
}
}
```

See the [complete example](#).

More Info

- [Visibility Timeout](#) in the *Amazon SQS Developer Guide*
- [SetQueueAttributes](#) in the *Amazon SQS API Reference*
- [GetQueueAttributes](#) in the *Amazon SQS API Reference*
- [ReceiveMessage](#) in the *Amazon SQS API Reference*
- [ChangeMessageVisibility](#) in the *Amazon SQS API Reference*
- [ChangeMessageVisibilityBatch](#) in the *Amazon SQS API Reference*

Using Dead Letter Queues in Amazon SQS

Amazon SQS provides support for *dead letter queues*. A dead letter queue is a queue that other (source) queues can target for messages that can't be processed successfully. You can set aside and isolate these messages in the dead letter queue to determine why their processing did not succeed.

To create a dead letter queue, you must first create a *redrive policy*, and then set the policy in the queue's attributes.

Important

A dead letter queue must be the same type of queue (FIFO or standard) that the source queue is. It must also be created using the same AWS account and region as the source queue.

Note

These code snippets assume that you understand the material in [Getting Started Using the AWS SDK for C++ \(p. 2\)](#) and have configured default AWS credentials using the information in [Providing AWS Credentials \(p. 5\)](#).

Creating a Redrive Policy

A redrive policy is specified in JSON. To create it, you can use the JSON utility class provided with the AWS SDK for C++.

Here is an example function that creates a redrive policy by providing it with the ARN of your dead letter queue and the maximum number of times the message can be received and not processed before it's sent to the dead letter queue.

Includes

```
#include <aws/core/Aws.h>
#include <aws/core/utils/json/JsonSerializer.h>
```

Code

```
Aws::String MakeRedrivePolicy(const Aws::String& queue_arn, int max_msg)
{
    Aws::Utils::Json::JsonValue redrive_arn_entry;
    redrive_arn_entry.AsString(queue_arn);

    Aws::Utils::Json::JsonValue max_msg_entry;
    max_msg_entry.AsInteger(max_msg);

    Aws::Utils::Json::JsonValue policy_map;
    policy_map.WithObject("deadLetterTargetArn", redrive_arn_entry);
    policy_map.WithObject("maxReceiveCount", max_msg_entry);

    return policy_map.View().WriteReadable();
}
```

See the [complete example](#).

Setting the Redrive Policy on your Source Queue

To finish setting up your dead letter queue, call the `SQSClient` class' `SetQueueAttributes` member function with a `SetQueueAttributesRequest` object for which you've set the `RedrivePolicy` attribute with your JSON redrive policy.

Includes

```
#include <aws/core/Aws.h>
#include <aws/sqs/SQSClient.h>
#include <aws/sqs/model/SetQueueAttributesRequest.h>
```

Code

```
Aws::SQS::SQSClient sqs;

Aws::String redrivePolicy = MakeRedrivePolicy(queue_arn, max_msg);

Aws::SQS::Model::SetQueueAttributesRequest request;
request.SetQueueUrl(src_queue_url);
```

```
request.AddAttributes(  
    Aws::SQS::Model::QueueAttributeName::RedrivePolicy,  
    redrivePolicy);  
  
auto outcome = sqs.SetQueueAttributes(request);  
if (outcome.IsSuccess())  
{  
    std::cout << "Successfully set dead letter queue for queue " <<  
        src_queue_url << " to " << queue_arn << std::endl;  
}  
else  
{  
    std::cout << "Error setting dead letter queue for queue " <<  
        src_queue_url << ": " << outcome.GetError().GetMessage() <<  
        std::endl;  
}
```

See the [complete example](#).

More Info

- [Using Amazon SQS Dead Letter Queues](#) in the *Amazon SQS Developer Guide*
- [SetQueueAttributes](#) in the *Amazon SQS API Reference*

Document History for the *AWS SDK for C++ Developer Guide*

This topic lists major changes to the *AWS SDK for C++ Developer Guide* over the course of its history.

- **Latest documentation update:** Oct 18, 2018

April 13, 2017

Added new examples: [Managing Access to Amazon S3 Buckets Using Bucket Policies](#) (p. 90) and [Configuring an Amazon S3 Bucket as a Website](#) (p. 92), as well as refreshing the content in [Creating, Listing, and Deleting Buckets](#) (p. 82) and [Operations on Objects](#) (p. 84).

Made the [Utility Modules](#) (p. 20) topic more useful!

April 03, 2017

Added new topics to the [Amazon CloudWatch Examples Using the AWS SDK for C++](#) (p. 26) section: [Getting Metrics from CloudWatch](#) (p. 26), [Publishing Custom Metric Data](#) (p. 28) [Working with CloudWatch Alarms](#) (p. 29), [Using Alarm Actions in CloudWatch](#) (p. 32) and [Sending Events to CloudWatch](#) (p. 34)

March 28, 2017

Added new topics to the [Amazon EC2 Examples Using the AWS SDK for C++](#) (p. 45) section: [Managing Amazon EC2 Instances](#) (p. 45), [Using Elastic IP Addresses in Amazon EC2](#) (p. 52), [Using Regions and Availability Zones for Amazon EC2](#) (p. 55), [Working with Amazon EC2 Key Pairs](#) (p. 57), and [Working with Security Groups in Amazon EC2](#) (p. 59)

March 23, 2017

Added new topics to the [IAM Code Examples Using the AWS SDK for C++](#) (p. 62) section: [Managing IAM Access Keys](#) (p. 62), [Using IAM Account Aliases](#) (p. 70), [Working with IAM Policies](#) (p. 72), and [Working with IAM Server Certificates](#) (p. 78)

March 10, 2017

Added new topics to [Amazon SQS Code Examples Using the AWS SDK for C++](#) (p. 95): [Using Dead Letter Queues in Amazon SQS](#) (p. 104), [doc:examples-sqs-long-polling](#), and [Setting Visibility Timeout in Amazon SQS](#) (p. 103)

February 27, 2017

A new topic in the **Getting Started** section, [Using the AWS SDK for C++](#) (p. 6), has been added to show how to properly initialize and shutdown the SDK.

In addition to the existing [Amazon S3 Code Examples Using the AWS SDK for C++](#) (p. 81) examples, new code examples have been added for [Amazon SQS Code Examples Using the AWS SDK for C++](#) (p. 95) and [IAM Code Examples Using the AWS SDK for C++](#) (p. 62).

February 02, 2016

Documentation first created.