
Amazon Simple Storage Service

Getting Started Guide

API Version 2006-03-01



Amazon Simple Storage Service: Getting Started Guide

Copyright © 2016 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.


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

Table of Contents

Getting Started	1
Amazon S3 Basics	1
Sign Up for Amazon S3	2
Create a Bucket	3
Add an Object to a Bucket	5
View an Object	7
Move an Object	8
Delete an Object and Bucket	10
Where Do I Go From Here?	13
Common Use Scenarios	13
Considerations Going Forward	13
AWS Account and Security Credentials	14
Security	14
AWS Integration	14
Naming Strategy	14
Pricing	14
Advanced Amazon S3 Features	15
Development Resources	15
Reference Resources	15
About This Guide	17

Getting Started with Amazon Simple Storage Service

Amazon Simple Storage Service (Amazon S3) is storage for the Internet. You can use Amazon S3 to store and retrieve any amount of data at any time, from anywhere on the web. You can accomplish these tasks using the AWS Management Console, which is a simple and intuitive web interface. This guide introduces you to Amazon S3 and how to use the AWS Management Console to complete the tasks shown in the following figure.



For information on Amazon S3 features, pricing, and to see the FAQ, go to the [Amazon S3 product page](#).

Amazon S3 Basics


To get the most out of Amazon S3, you need to understand a few simple concepts. Amazon S3 stores data as objects within buckets. An object consists of a file and optionally any metadata that describes that file.

To store an object in Amazon S3, you upload the file you want to store to a bucket. When you upload a file, you can set permissions on the object as well as any metadata.

Buckets are the containers for objects. You can have one or more buckets. For each bucket, you can control access to it (who can create, delete, and list objects in the bucket), view access logs for it and its objects, and choose the geographical region where Amazon S3 will store the bucket and its contents.

When using the AWS Management Console, you can create folders to group objects, and you can nest folders (create folders within folders). If you have used the Amazon S3 API or other utilities, you can learn some important aspects about how folders work with other grouping conventions in the [Amazon Simple Storage Service Console User Guide](#).

Sign Up for Amazon S3



To use Amazon S3, you need an AWS account. If you don't already have one, you'll be prompted to create one when you sign up for Amazon S3. You will not be charged for Amazon S3 until you use it. For information on Amazon S3 features, pricing, and to see the FAQ, go to the [Amazon S3 product page](#).

To sign up for Amazon S3

1. Go to <https://aws.amazon.com/s3/> and click **Sign Up**.
2. Follow the on-screen instructions.

AWS will notify you by email when your account is active and available for you to use.

Create a Bucket



Now that you've signed up for Amazon S3, you're ready to create a bucket using the AWS Management Console. Every object in Amazon S3 is stored in a bucket. Before you can store data in Amazon S3, you must create a bucket.

Note

You are not charged for creating a bucket; you are charged only for storing objects in the bucket and for transferring objects in and out of the bucket. The charges you will incur through following the examples in this guide are minimal (less than \$1). For more information, go to [Amazon S3 Pricing](#).

To create a bucket

1. Sign into the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3>.
2. Click **Create Bucket**.

Create a Bucket - Select a Bucket Name and Region Cancel x

A bucket is a container for objects stored in Amazon S3. When creating a bucket, you can choose a Region to optimize for latency, minimize costs, or address regulatory requirements. For more information regarding bucket naming conventions, please visit the [Amazon S3 documentation](#).

Bucket Name:

Region:

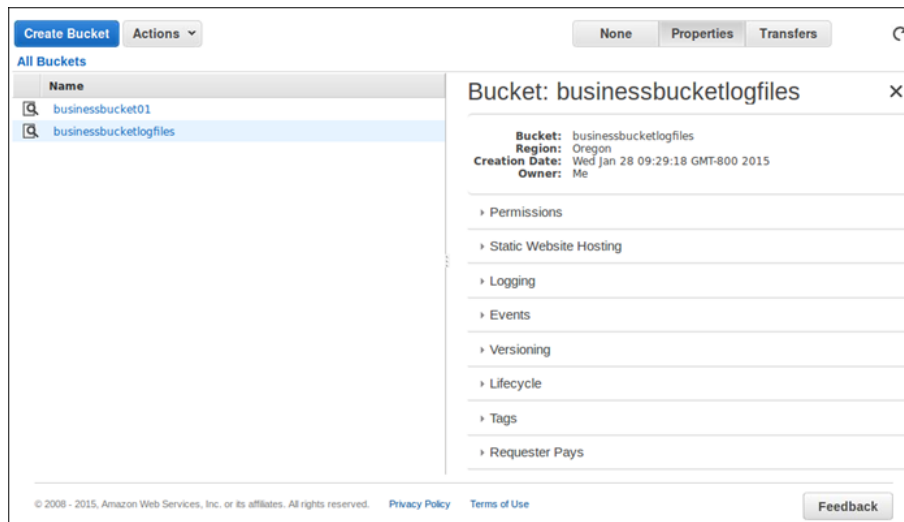
3. In the **Create a Bucket** dialog box, in the **Bucket Name** box, enter a bucket name. The bucket name you choose must be unique across all existing bucket names in Amazon S3. One way to help ensure uniqueness is to prefix your bucket names with the name of your organization. Bucket names must comply with certain rules. For more information, see [Bucket Restrictions and Limitations](#) in the *Amazon Simple Storage Service Developer Guide*.

Note

After you create a bucket, you cannot change its name. In addition, the bucket name is visible in the URL that points to the objects stored in the bucket. Ensure that the bucket name you choose is appropriate.

4. In the **Region** box, select a region. For this exercise, select **Oregon** from the drop-down list. You can choose a region to optimize latency, minimize costs, or address regulatory requirements. Objects stored in a region never leave that region unless you explicitly transfer them to another region. For more information about regions, see [Accessing a Bucket](#) in the *Amazon Simple Storage Service Developer Guide*.
5. Click **Create**.

When Amazon S3 successfully creates your bucket, the console displays your empty bucket in the **Buckets** panel.



You've created a bucket in Amazon S3.

Add an Object to a Bucket




Now that you've created a bucket, you're ready to add an object to it. An object can be any kind of file: a text file, a photo, a video and so forth. When you add a file to Amazon S3, you have the option of including metadata with the file and setting permissions to control access to the file.

Important


For information about storage charges and related information, go to the [Amazon S3 Pricing page](#).

To upload an object

1. In the [Amazon S3 console](#), click the name of bucket that you want to upload an object to and then click **Upload**.



2. In the **Upload - Select Files** wizard, if you want to upload an entire folder, you must click **Enable Enhanced Uploader** to install the necessary Java applet. You only need to do this once per console session.




Note

If you are behind a firewall you will need to install your organization's supported proxy client in order for the Java applet to work.

3. Click **Add Files**.

A file selection dialog box opens:

- If you enabled the advanced uploader in step 2, you see a Java dialog box titled **Select files and folders to upload**, as shown.
- If not, you see the **File Upload** dialog box associated with your operating system.



4. Select the file that you want to upload and then click **Open**.

5. Click **Start Upload**.

You can watch the progress of the upload from within the **Transfer** panel.

Tip

To hide the **Transfer** dialog box, click the **Close** button at top right in the **Transfers** panel. To open it again, click **Transfers**.

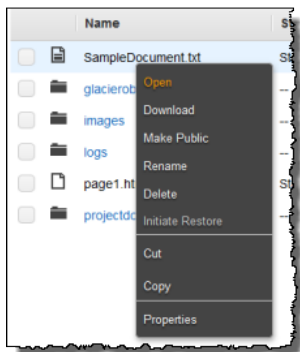
View an Object



Now that you've added an object to a bucket, you can open and view it in a browser. You can also download the object to your local computer.

To open or download an object

1. In the [Amazon S3 console](#), in the **Objects and Folders** list, right-click the object or objects that you want to open or download, then click **Open** or **Download** as appropriate.



2. If you are downloading the object, specify where you want to save it. The procedure for saving the object depends on the browser and operating system that you are using.

Note

By default, your Amazon S3 buckets and objects are private. To make an object viewable by using a URL, for example, <https://s3.amazonaws.com/Bucket/Object>, you must make the object publicly readable. Otherwise, you will need to create a signed URL that includes a signature with authentication information. For information on making objects publicly accessible, see [Editing Object Permissions](#). For more information about pre-signed URLs, go to [Share an Object with Others](#) in the *Amazon Simple Storage Service Developer Guide*.

Move an Object



Now that you've added an object to a bucket and viewed it, you can move the object to a different bucket or folder.

Note

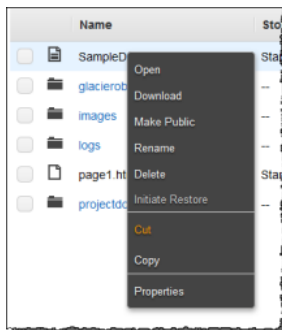
In this exercise, you work with one object; however, you can also follow these steps for moving a folder.

To move an object

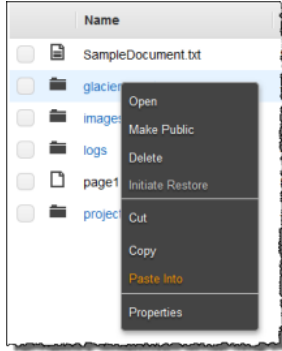
1. In the [Amazon S3 console](#), right-click the object that you want to move, and then click **Cut**.

Tip

You can use the **SHIFT** and **CTRL** keys to select multiple objects and perform the same action on them simultaneously.




2. Navigate to the bucket or folder where you want to move the object. Right-click the folder or bucket and then click **Paste Into**.



Tip

You can monitor the progress of the move using the **Transfers** dialog box. To open it, click the **Transfers** button. To hide it, click the **Close** button at top right in the **Transfers** dialog box.

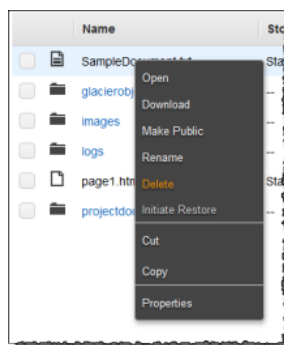
Delete an Object and Bucket



If you no longer need to store the objects that you uploaded and moved while going through this guide, you should delete them to prevent further charges.

To delete an object

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Objects and Folders** panel, right-click the object that you want to delete, and then click **Delete**.



Tip

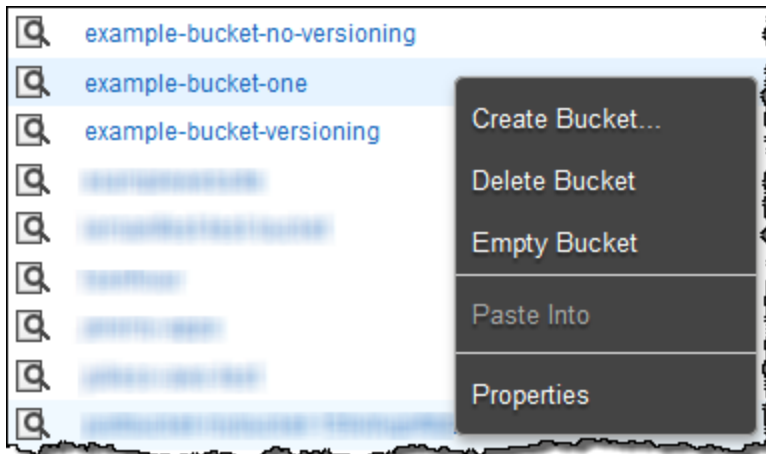
You can use the **SHIFT** and **CRTL** keys to select multiple objects and perform the same action on them simultaneously.

3. When a confirmation message appears, click **OK**.

You can empty a bucket, which deletes all the objects in the bucket without deleting the bucket. For information on the limitations for emptying a bucket, see [Deleting/Emptying a Bucket](#) in the *Amazon Simple Storage Service Developer Guide*.

To empty a bucket

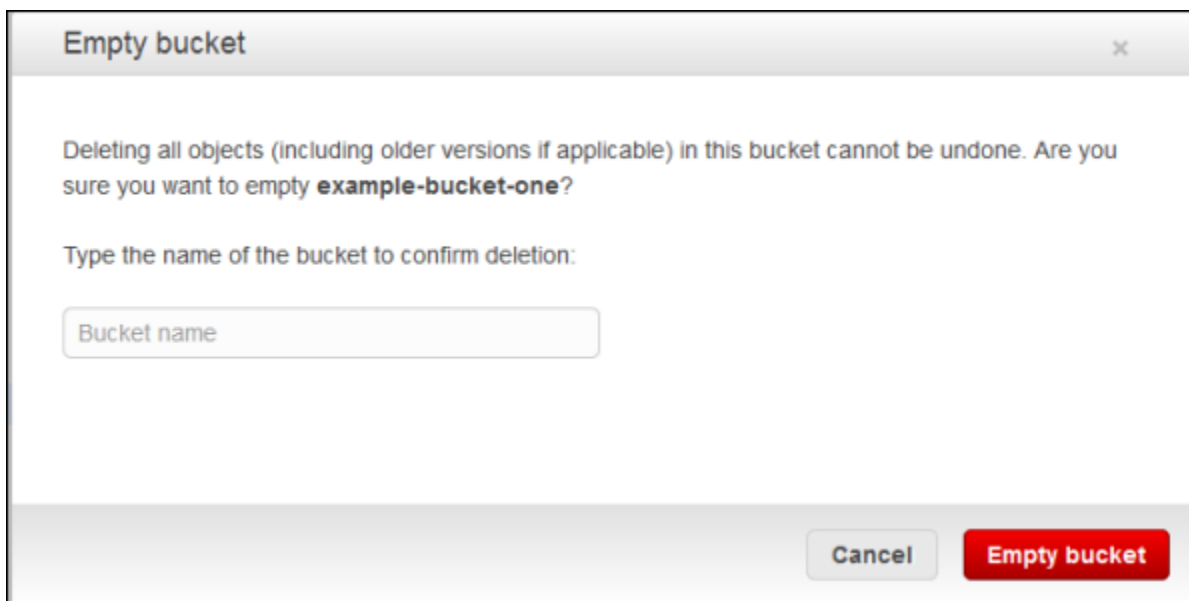
1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Right-click the bucket that you want to empty, and then click **Empty Bucket**.



Tip

Optionally, to get this menu you can click the bucket and then click **Actions**, which is near the top of the console window next to **Create Bucket**.

3. When a confirmation message appears, enter the bucket name and then click **Empty bucket**.



You can delete a bucket and all the objects contained in the bucket. For information on the limitations for deleting a bucket, see [Deleting/Emptying a Bucket](#) in the *Amazon Simple Storage Service Developer Guide*.

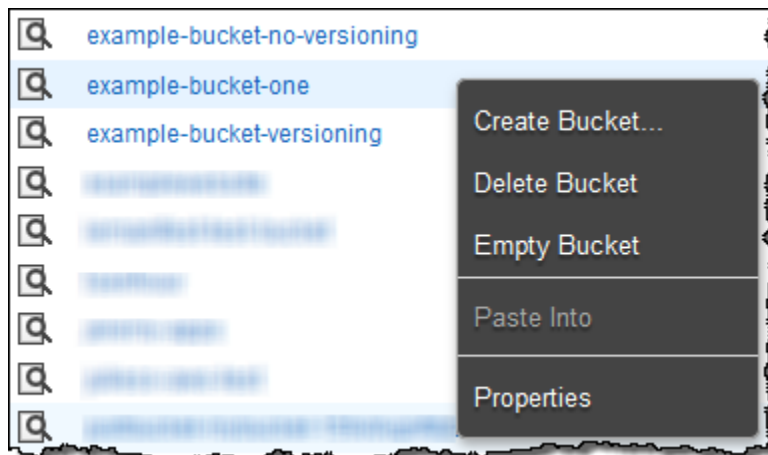
Important

If you want to continue to use the same bucket name, don't delete the bucket. We recommend that you empty the bucket and keep it. After a bucket is deleted, the name becomes available to reuse, but the name might not be available for you to reuse for various reasons. For

example, it might take some time before the name can be reused and some other account could create a bucket with that name before you do.

To delete a bucket

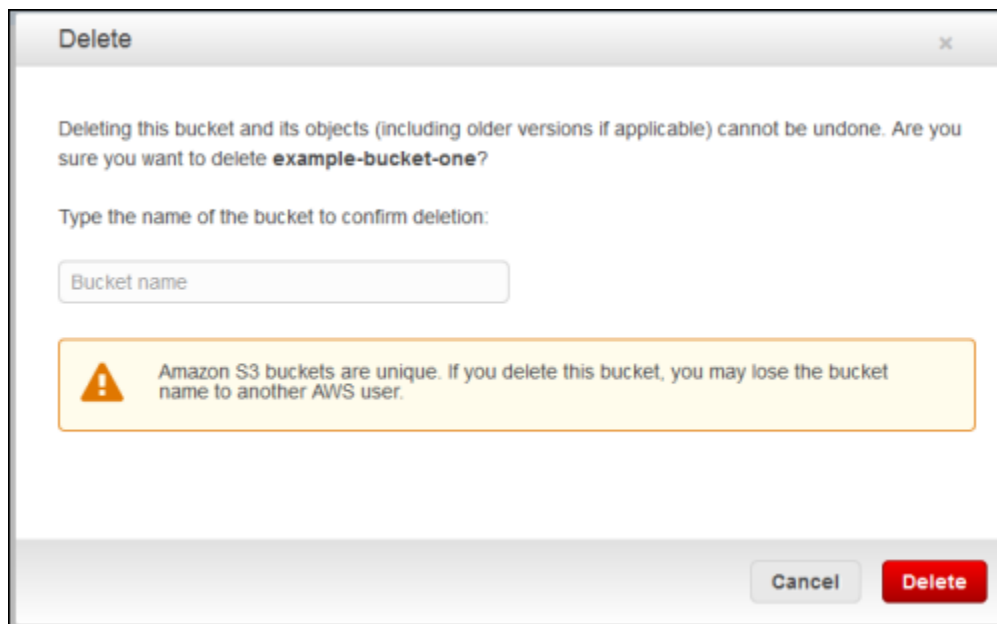
1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Right-click the bucket that you want to delete, and then click **Delete Bucket**.



Tip

Optionally, to get this menu you can click the bucket and then click **Actions**, which is near the top of the console window next to **Create Bucket**.

3. When a confirmation message appears, enter the bucket name and click **Delete**.



Where Do I Go From Here?

Topics

- [Common Use Scenarios](#) (p. 13)
- [Considerations Going Forward](#) (p. 13)
- [Advanced Amazon S3 Features](#) (p. 15)
- [Development Resources](#) (p. 15)
- [Reference Resources](#) (p. 15)

In the preceding examples, you learned how to perform some basic Amazon S3 tasks. You can also learn to use our APIs and code samples to perform basic and advanced Amazon S3 tasks by reading the *Amazon Simple Storage Service Developer Guide* and by looking at code samples. For more information, go to the [Amazon Simple Storage Service Developer Guide](#) and <https://aws.amazon.com/code/Amazon-S3>, respectively.

This section explains various ways in which you can gain a deeper understanding of Amazon S3 so that you can implement it in your applications.

Common Use Scenarios

The AWS Solutions web page lists many of the ways you can use Amazon S3. The following list summarizes some of those ways.

- **Backup and Storage** – Provide data backup and storage services for others.
- **Application Hosting** – Provide services that deploy, install, and manage web applications.
- **Media Hosting** – Build a redundant, scalable, and highly available infrastructure that hosts video, photo, or music uploads and downloads.
- **Software Delivery** – Host your software applications that customers can download.

For more information, go to [AWS Solutions](#).

Considerations Going Forward

Topics

- [AWS Account and Security Credentials](#) (p. 14)

- [Security \(p. 14\)](#)
- [AWS Integration \(p. 14\)](#)
- [Naming Strategy \(p. 14\)](#)
- [Pricing \(p. 14\)](#)

This section introduces you to topics you should consider before launching your own Amazon S3 product.

AWS Account and Security Credentials

So far you signed up for the service, got an AWS account and security credentials, and then completed a short exercise covering the essential product functions. Now that you're finished with the exercise, we recommend that you check with an administrator or coworker in your organization to determine if he or she already has an AWS account and security credentials for you to use in future interactions with AWS.

If you're an account owner or administrator and want to know more about AWS Identity and Access Management, see the product description at <https://aws.amazon.com/iam> or the technical documentation in the [IAM User Guide](#).

Security

Amazon S3 provides authentication mechanisms to secure data stored in Amazon S3 against unauthorized access. Unless you specify otherwise, only the AWS account owner can access data uploaded to Amazon S3. For more information about how to manage access to buckets and objects, go to [Managing Access Permissions to Your Amazon S3 Resources](#) in the *Amazon Simple Storage Service Developer Guide*.

You can also encrypt your data before uploading it to Amazon S3.

AWS Integration

You can use Amazon S3 alone or in concert with one or more other Amazon products. The most common products used with Amazon S3 are:

- [Amazon EC2](#)
- [Amazon Elastic MapReduce](#)
- [Amazon SQS](#)
- [Amazon CloudFront](#)
- [Amazon DevPay](#)

Naming Strategy

Plan your bucket names in advance. The location of your data in Amazon S3 is a URL, generally of the form: `http://[bucket-name].s3.amazonaws.com/[key]`. The bucket and key names should be descriptive of the objects. Each bucket is a namespace. Within one bucket, key names cannot clash. Before naming objects in buckets, you should develop a naming strategy.

Pricing

Learn the pricing structure for storing and transferring data on Amazon S3. For more information, go to [Amazon S3 Pricing](#).

Advanced Amazon S3 Features

The examples in this guide show how to accomplish the basic tasks of creating a bucket, uploading and downloading data to and from it, and moving and deleting the data. The following table summarizes some of the most common advanced functionality offered by Amazon S3. Note that some advanced functionality is not available in the AWS Management Console and requires that you use the Amazon S3 API. All advanced functionality and how to use it is described in the [Amazon Simple Storage Service Developer Guide](#).

Link	Functionality
Using Amazon DevPay with Amazon S3	Learn how you can use Amazon DevPay to charge customers that access the data you store on Amazon S3.
Requester Pays Buckets	Learn how to configure a bucket so that a customer pays for the downloads they make.
Using BitTorrent With Amazon S3	Use BitTorrent, which is an open, peer-to-peer protocol for distributing files.
Versioning	Learn about Amazon S3's Versioning capabilities.
Hosting Static Websites	Learn how to host a static website on Amazon S3.
Object Lifecycle Management	Learn how to manage the lifecycle of objects in your bucket. Lifecycle management includes expiring objects and archiving objects (transitioning objects to the <code>GLACIER</code> storage class).

Development Resources

To help you build applications using the language of your choice, we provide the following resources:

- **Sample Code and Libraries** – The AWS Resource Center has sample code and libraries written especially for Amazon S3. You can use these code samples as a means of understanding how to implement the Amazon S3 API. For more information, go to <https://aws.amazon.com/code/Amazon-S3>.
- **Tutorials** – Our Resource Center also offers more Amazon S3 tutorials. These tutorials provide a hands-on approach for learning Amazon S3 functionality. For more information, go to <https://aws.amazon.com/articles/Amazon-S3>.
- **Customer Forum** – We recommend you review the Amazon S3 forum to get an idea of what other users are doing and to benefit from the questions they've asked. The forum can help you understand what you can and can't do with Amazon S3. The forum also serves as a place for you to ask questions that other users or Amazon representatives might answer. You can use the forum to report issues with the service or the API. For more information, go to [Discussion Forum](#).
- **Case Studies** – To see code for applications that combine multiple Amazon products, including Amazon S3, go to [Customer Apps: Amazon Simple Storage Service](#).

Reference Resources

The following list shows additional resources you can use to further your understanding of Amazon S3.

- The [Amazon Simple Storage Service Console User Guide](#) describes all of the AWS Management Console functions related to Amazon S3.

- The [Amazon Simple Storage Service Developer Guide](#) provides a detailed discussion of the service. It includes an architectural overview, detailed concept descriptions, and procedures for using the API.
- The [Amazon Simple Storage Service API Reference](#) provides a detailed discussion of the actions and parameters in Amazon S3.
- The [Amazon Simple Storage Service Quick Reference Card](#) provides quick access to important Amazon S3 API operations and functionality.
- The AWS Developer Resource Center is the landing page for Amazon S3 code samples, tutorials, documentation, links to customer forums, pricing policies, and other information to help you build innovative applications with Amazon S3. For more information, go to [the AWS home page](#).
- The Service Health Dashboard shows you the status of the Amazon S3 web service. The dashboard shows you whether Amazon S3 (and all other AWS products) are functioning properly. For more information, go to [Service Health Dashboard](#).

About This Guide

This is the *Amazon Simple Storage Service Getting Started Guide*.

Amazon Simple Storage Service is frequently referred to within this guide as "Amazon S3." All copyrights and legal protections still apply.