
AWS Storage Gateway

User Guide

API Version 2012-06-30



AWS Storage Gateway: User Guide

Copyright © 2013 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

The following are trademarks of Amazon Web Services, Inc.: Amazon, Amazon Web Services Design, AWS, Amazon CloudFront, Cloudfront, Amazon DevPay, DynamoDB, ElastiCache, Amazon EC2, Amazon Elastic Compute Cloud, Amazon Glacier, Kindle, Kindle Fire, AWS Marketplace Design, Mechanical Turk, Amazon Redshift, Amazon Route 53, Amazon S3, Amazon VPC. In addition, Amazon.com graphics, logos, page headers, button icons, scripts, and service names are trademarks, or trade dress of Amazon in the U.S. and/or other countries. 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.

What Is AWS Storage Gateway?	1
How AWS Storage Gateway Works	3
Requirements	6
Pricing	6
Getting Started	7
Step 1: Sign Up	8
Step 2: Try an Example Setup	8
Step 2.1: Set Up and Activate a Gateway	9
Set Up and Activate (VMware Host)	9
Provision Host	9
Download and Deploy the VM	10
Provision Local Disk Storage	18
Provision Local Disk Storage (Gateway-Cached)	19
Provision Local Disk Storage (Gateway-Stored)	24
Configure the VM to Use Paravirtualized Disk Controllers	29
Activate the Gateway	30
Set Up and Activate (Hyper-V Host)	33
Provision Host	34
Download and Deploy the VM	34
Provision Local Disk Storage	43
Provision Local Disk Storage (Gateway-Cached)	44
Provision Local Disk Storage (Gateway-Stored)	50
Activate the Gateway	56
Step 2.2: Create Volumes	61
Create Volumes (Gateway-Cached)	61
Create Volumes (Gateway-Stored)	67
Step 2.3: Access Your Volumes	72
Step 2.4: Test the Setup	79
Test the Setup (Gateway-Cached)	80
Test the Setup (Gateway-Stored)	83
Where Do I Go from Here?	87
Setting Up AWS Storage Gateway	90
Deploying and Activating Up a Gateway On-Premises	90
Deploying and Activating a Gateway on a VMware ESXi Host	91
Downloading and Deploying the VM	91
Provisioning Local Disk Storage	92
Provisioning Local Disks (Gateway-Cached)	93
Adding Local Disks for Cache Storage	94
Adding Local Disks for Upload Buffer	98
Provisioning Local Disks (Gateway-Stored)	102
Adding Local Disks for Volume Storage	102
Adding Local Disks for Upload Buffer	105
Configure VM to Use Paravirtualization	110
Activating a Gateway	111
Deploying and Activating a Gateway on a Microsoft Hyper-V Host	115
Downloading and Deploying the VM	115
Provisioning Local Disk Storage	115
Provisioning Local Disks (Gateway-Cached)	117
Adding Local Disks for Cache Storage	117
Adding Local Disks for Upload Buffer	120
Provisioning Local Disks (Gateway-Stored)	124
Adding Local Disks for Volume Storage	125
Adding Local Disks for Upload Buffer	127
Activating a Gateway	131
Deploying and Activating a Gateway on Amazon EC2	137
Launching and Activating an Amazon EC2 Gateway AMI	139
Managing Your Amazon EC2 Gateway	146
Adding and Removing Amazon EBS Volumes from Your Instance	147

Enabling and Disabling AWS Support Access	149
Configuring Security Groups	149
Cleaning Up Resources After Deleting a Gateway	150
Configuring Upload Buffer and Cache Storage	150
Configuring Upload Buffer (Gateway-Cached)	150
Configuring Cache Storage (Gateway-Cached)	152
Configuring Upload Buffer (Gateway-Stored)	154
Creating Storage Volumes	157
Creating a Storage Volume (Gateway-Cached)	157
Creating a Storage Volume (Gateway-Stored)	159
Configuring Your Application Access to Storage Volumes	161
Connecting from a Windows Client to Your Storage Volume	163
Connecting from a Red Hat Client to Your Storage Volume	165
Configuring CHAP Authentication for Your Storage Volume	167
Managing Your Activated Gateway	176
Managing Storage Volumes	176
Managing Storage Volumes (Gateway-Cached)	181
Managing Storage Volumes (Gateway-Stored)	183
Managing Upload Buffer and Cache Storage (Gateway-Cached)	187
Adding and Removing Upload Buffer Capacity	190
Adding Cache Storage	193
Configuring the Upload Buffer (Gateway-Stored)	193
Adding and Removing Upload Buffer Capacity	195
Working with Snapshots	199
Finding a Snapshot	200
Finding Snapshots Using the AWS SDK for Java	202
Finding Snapshots Using the AWS SDK for .NET	204
Finding Snapshots Using the AWS Tools for Windows PowerShell	206
Editing a Snapshot Schedule	208
Creating an Ad-Hoc Snapshot	209
Deleting Snapshots	209
Deleting Snapshots Using the AWS SDK for Java	211
Deleting Snapshots Using the AWS SDK for .NET	214
Deleting Snapshots Using the AWS Tools for Windows PowerShell	217
Restoring a Snapshot	219
Restoring a Snapshot	219
Restoring a Snapshot to an Amazon EBS Volume	222
Performing Maintenance Tasks	222
Shutting Down and Turning On a Gateway	224
Managing Gateway Updates	226
Updating Gateway Rate Limits	227
Updating Gateway Rate Limits Using the AWS SDK for Java	228
Updating Gateway Rate Limits Using the AWS SDK for .NET	230
Updating Gateway Rate Limits Using the AWS Tools for Windows PowerShell	232
Deleting a Gateway	233
Logging into Your Gateway Local Console	234
Routing Your Gateway Through a Proxy	238
Configuring Your Gateway to Use a Static IP Address	239
Testing Your Gateway Connectivity to the Internet	242
Synchronizing Your Gateway VM Time	243
Configuring Your Gateway for Multiple Network Adapters (NICs)	245
Creating a Storage Volume on a Gateway with Multiple Network Adapters	251
Troubleshooting	252
Optimizing Gateway Performance	260
Monitoring Your AWS Storage Gateway	261
Using the Amazon CloudWatch Console	261
Measuring Performance Between Your Application and Gateway	262
Measuring Performance Between Your Gateway and AWS	264

Monitoring the Upload Buffer	267
Monitoring Cache Storage	271
Understanding AWS Storage Gateway Metrics	272
Access Control	277
API Reference	283
Required Request Headers	283
Signing Requests	285
Error Responses	287
Operations in AWS Storage Gateway	305
ActivateGateway	307
AddCache	310
AddUploadBuffer	313
AddWorkingStorage	315
CreateCachediSCSIVolume	318
CreateSnapshot	321
CreateSnapshotFromVolumeRecoveryPoint	324
CreateStorediSCSIVolume	327
DeleteBandwidthRateLimit	331
DeleteChapCredentials	333
DeleteGateway	336
DeleteSnapshotSchedule	338
DeleteVolume	340
DescribeBandwidthRateLimit	343
DescribeCache	345
DescribeCachediSCSIVolumes	348
DescribeChapCredentials	352
DescribeGatewayInformation	354
DescribeMaintenanceStartTime	358
DescribeSnapshotSchedule	360
DescribeStorediSCSIVolumes	363
DescribeUploadBuffer	367
DescribeWorkingStorage	370
ListGateways	372
ListLocalDisks	375
ListVolumeRecoveryPoints	378
ListVolumes	381
ShutdownGateway	384
StartGateway	387
UpdateBandwidthRateLimit	389
UpdateChapCredentials	391
UpdateGatewayInformation	394
UpdateGatewaySoftwareNow	396
UpdateMaintenanceStartTime	399
UpdateSnapshotSchedule	401
Data Types	404
CachediSCSIVolume	404
ChapInfo	406
Disk	407
GatewayInfo	407
NetworkInterface	408
StorediSCSIVolume	408
VolumeInfo	409
VolumeiSCSIAttributes	410
VolumeRecoveryPointInfo	410
Enumeration Types	411
BandwidthType	411
DiskAllocationType	411
GatewayState	411

GatewayTimezone	411
GatewayType	412
Regions	412
VolumeStatus	412
VolumeType	412
Document History	413
Appendices	415
Appendix A: The Components in Your vSphere Environment	415
Appendix B: Configuring a VMware ESXi Host	417
Appendix C: The Components in Your Hyper-V Environment	421
Appendix D: Configuring a Microsoft Hyper-V Host	422
Appendix E: About AWS Storage Gateway	433

What Is AWS Storage Gateway?

Topics

- [How AWS Storage Gateway Works \(p. 3\)](#)
- [Requirements \(p. 6\)](#)
- [Pricing \(p. 6\)](#)
- [AWS Storage Gateway API \(p. 6\)](#)

Welcome to the *AWS Storage Gateway User Guide*. AWS Storage Gateway is a service that connects an on-premises software appliance with cloud-based storage to provide seamless and secure integration between your on-premises IT environment and AWS's storage infrastructure. The service offers you the following storage solutions:

- **Gateway-Cached Volume Solution**—In this storage architecture, you create your storage volumes and mount them as iSCSI devices from your on-premises application servers. The gateway stores data you write to your gateway-cached volume in Amazon Simple Storage Service (Amazon S3), and stores only a cache of frequently accessed data on your on-premises storage hardware. Storing your volume data in Amazon S3 minimizes the need for you to scale your on-premises storage infrastructure, since Amazon S3 scales on demand.
- **Gateway-Stored Volume Solution**—In this storage architecture, you store all your data locally in storage volumes on your on-premises storage hardware. The gateway periodically takes snapshots as incremental backups and stores them in Amazon S3.

Note that both of these storage solutions enable you to schedule snapshots that the gateway stores in Amazon S3 in the form of Amazon Elastic Block Store (Amazon EBS) snapshots. For more information, see [How AWS Storage Gateway Works \(p. 3\)](#).

AWS Storage Gateway enables a wide range of use cases, including the following:

- **Corporate File Sharing** – Managing on-premises storage for departmental file shares and home directories typically results in high capital and maintenance costs, under-utilized hardware, and restrictive user quotas. AWS Storage Gateway addresses these on-premises scaling and maintenance issues by enabling you to seamlessly store your corporate file shares on Amazon S3, while keeping a copy of your frequently accessed files on-premises. This minimizes the need to scale your on-premises file storage infrastructure, while still providing low-latency access to your frequently accessed data.
- **Backup** – Both the storage solutions AWS Storage Gateway offers enable your existing on-premises applications to store data backups off-site in Amazon S3. All data is securely transferred to AWS over SSL and stored encrypted in Amazon S3 using AES 256-bit encryption. AWS Storage Gateway provides

an attractive alternative to the traditional choice of either maintaining costly hardware in multiple data centers, or dealing with the longer recovery times and operational burden of managing off-site tape storage.

- **Disaster Recovery and Resilience** – AWS Storage Gateway addresses the data replication challenges of disaster recovery (DR) by enabling you to create Gateway-Stored volumes that maintain your primary data on-premises, while storing point-in-time backup snapshots of this data in Amazon S3 as Amazon EBS snapshots. Amazon S3 redundantly stores these snapshots in multiple facilities and on multiple devices within each facility, quickly detecting and repairing any lost redundancy. Using [AWS CLI](#), you can configure virtual machine images of your DR application servers in AWS, and pay for these servers only when you need them. If your on-premises infrastructure goes down, you simply launch the Amazon EC2 compute instances that you need, restore your snapshots to new Amazon EBS volumes, attach the volumes to your running Amazon EC2 instances, and your DR environment is up and running.
- **Data Mirroring to Cloud-Based Compute Resources** – If you want to leverage Amazon EC2's on-demand compute capacity for additional capacity during peak periods, whether for new projects or as a more cost-effective way to run your normal workloads, you can use AWS Storage Gateway to mirror your on-premises data to Amazon EC2 instances.

For more information about use cases, go to the [Common Use Cases](#) section, and for service highlights, go to [Service Highlights](#) section on the *AWS Storage Gateway product detail page*.

If you are a first-time user of AWS Storage Gateway, we recommend that you begin by reading the following sections:

- **What is AWS Storage Gateway**—The rest of this section provides service highlights, a deployment overview, and the requirements for deploying the AWS Storage Gateway virtual machine (VM).
- **Getting Started with AWS Storage Gateway (p. 7)**—The Getting Started section provides you with instructions to set up an AWS Storage Gateway virtual machine (VM), activate it, and configure it so that you have a working gateway. You also test the setup in which you save sample data locally, take a backup snapshot that the gateway uploads to AWS, and restore the snapshot to your local storage volume, showing you how AWS Storage Gateway enables you to recover your data.

Beyond the Getting Started exercise, you'll learn more about how to use AWS Storage Gateway. The following sections cover the fundamentals of setting up, managing, troubleshooting, and monitoring your gateway.

- **Setting Up AWS Storage Gateway (p. 90)** – The Getting Started section provides the minimum required steps to set up and test a gateway. This section provides additional information, such as how to estimate the amount of working storage that your gateway requires. Additionally, if you follow the AWS Storage Gateway console wizard to set up your gateway, the wizard steps provide help links to the topics in this section.
- **Managing Your Activated Gateway (p. 176)** – After you deploy and activate your gateway, this section provides you with information about how to manage your gateway. The ongoing management tasks include adding storage volumes and working storage, working with snapshots, general maintenance, troubleshooting, and monitoring your gateway.

When working with snapshots, you want to know the difference between default and ad-hoc snapshots, how to find information about a snapshot, and how to schedule a snapshot. For more information, see [Working with Snapshots \(p. 199\)](#). This section also describes how to restore a snapshot locally to a new AWS Storage Gateway volume, or use a snapshot to create an Amazon EBS volume and attach it to an Amazon EC2 instance. For more information, see [Restoring a Snapshot \(p. 219\)](#).

You can monitor your gateway using Amazon CloudWatch metrics. AWS Storage Gateway displays key operational metrics for your gateway, storage volumes, and working storage in the AWS Management Console. In Amazon CloudWatch, you can measure the performance between your application and your gateway and between the gateway and AWS. You can also view metrics for throughput, latency,

and a number of input/output operations. For more information, see [Monitoring Your AWS Storage Gateway](#) (p. 261).

How AWS Storage Gateway Works

Topics

- [AWS Storage Gateway: Gateway-Cached Volume Architecture](#) (p. 3)
- [AWS Storage Gateway: Gateway-Stored Volume Architecture](#) (p. 5)

AWS Storage Gateway service architecture enables integration between your organization's on-premises IT environment and AWS's storage infrastructure. AWS Storage Gateway provides the following two storage options to enable this integration.

- Gateway-cached volumes enable you to utilize Amazon S3 as your primary data storage while retaining frequently accessed data local in your AWS Storage Gateway. Gateway-cached volumes minimize the need to scale your on-premises storage infrastructure, while still providing your applications with low-latency access to their frequently accessed data. You can create storage volumes up to 32 TiB in size and attach to them as iSCSI devices from your on-premises application servers. Data written to these volumes is stored in Amazon S3 and retained along with recently read data in your on-premises AWS Storage Gateway's cache and upload buffer storage.

Gateway-cached volumes can range from 1 GiB to 32 TiB in size and must be rounded to the nearest GiB. Each gateway configured for gateway-cached volumes can support up to 20 volumes and a total volume storage of 150 TiB.

- Gateway-stored volumes enable you to store your primary data locally, while asynchronously backing up that data to AWS. Gateway-stored volumes provide your on-premises applications with low-latency access to their entire data sets, while providing durable, off-site backups. You can create storage volumes up to 1 TiB in size and mount them as iSCSI devices from your on-premises application servers. Data written to your gateway-stored volumes is stored on your on-premises storage hardware, and asynchronously backed up to Amazon S3 in the form of Amazon EBS snapshots.

Gateway-stored volumes can range from 1 GiB to 1 TiB in size and must be rounded to the nearest GiB. Each gateway configured for gateway-stored volumes can support up to 12 volumes and a total volume storage of 12 TiB.

In both cases, AWS Storage Gateway takes snapshots, makes incremental backups, and stores them in AWS.

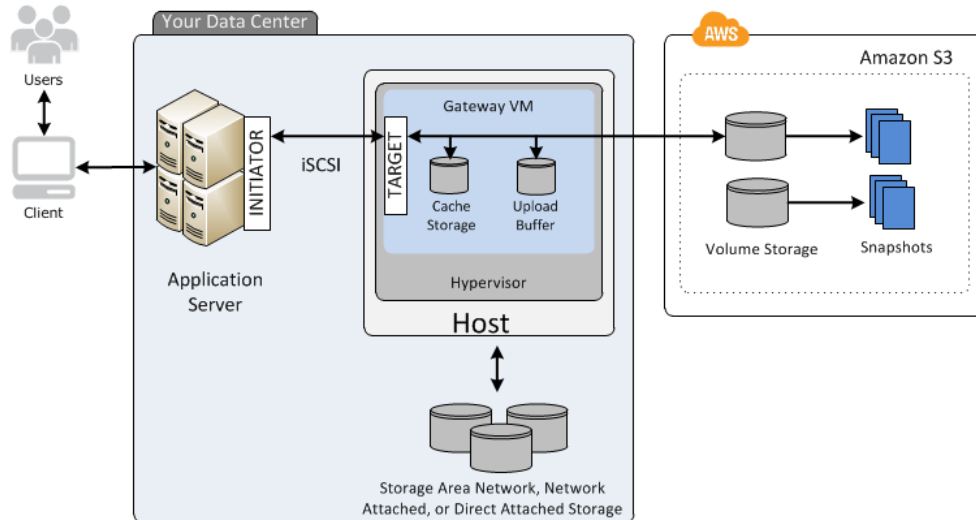
AWS Storage Gateway: Gateway-Cached Volume Architecture

In the gateway-cached volume solution, AWS Storage Gateway stores all your on-premises application data in a storage volume in Amazon S3.

The following diagram provides an overview of the AWS Storage Gateway's cached volume deployment.

AWS Storage Gateway User Guide

AWS Storage Gateway: Gateway-Cached Volume Architecture



Once you've installed AWS Storage Gateway's software appliance (the virtual machine (VM)) on a host in your data center and activated it, you can use the AWS Management Console to provision storage volumes backed by Amazon S3. You can also provision storage volumes programmatically using the AWS Storage Gateway API or the AWS SDK libraries. You then mount these storage volumes to your on-premises application servers as iSCSI devices.

You also allocate disks on-premises for the VM. These on-premises disks serve the following purposes:

- **Disks for use by the gateway as cache storage**—As your applications write data to the storage volumes in AWS, the gateway initially stores the data on the on-premises disks referred to as cache storage before uploading it to Amazon S3. The cache storage acts as the on-premises durable store for data that is pending upload to Amazon S3 from the upload buffer.

The cache storage also enables the gateway to store your application's recently accessed data on-premises for low-latency access. If your application requests data, the gateway first checks the cache storage for the data before checking Amazon S3.

There are some rules to the amount of disk space you can allocate for the cache storage. As a general rule, you should allocate at least 20 percent of your existing file store size; however, cache storage should be larger than the upload buffer. This ensures cache storage is large enough to be able to persistently hold all data that is in the upload buffer that has not yet been uploaded to Amazon S3.

- **Disks for use by the gateway as the upload buffer**—To prepare for upload to Amazon S3, your gateway also stores incoming data in a staging area, referred to as an upload buffer. Your gateway uploads this buffer data over an encrypted SSL connection to AWS where it is stored encrypted in Amazon S3.

You can take incremental backups, called *snapshots*, of your storage volumes in Amazon S3. These point-in-time snapshots are also stored in Amazon S3 as Amazon EBS snapshots. When you take a new snapshot, only the data that has changed since your last snapshot is stored. You can initiate snapshots on a scheduled or ad-hoc basis. When you delete a snapshot, only the data not needed for any other snapshots is removed.

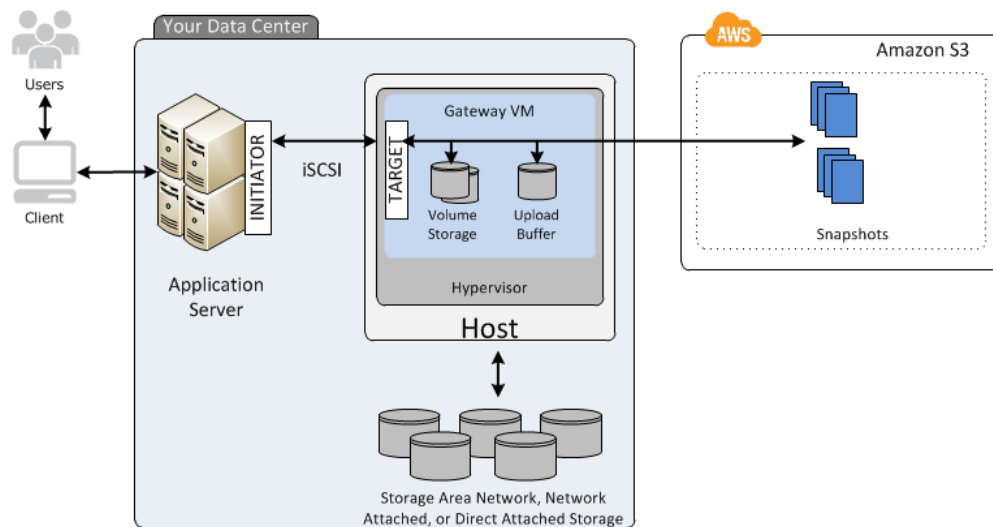
You can restore an Amazon EBS snapshot to a gateway storage volume in the event you need to recover a backup of your data. We plan to add support for Amazon EC2 deployable gateways in the near future, enabling you to restore your snapshot to an Amazon EC2 gateway storage volume. Alternatively, for snapshots up to 1 TiB in size, you can use the snapshot as a starting point for a new Amazon EBS volume, which you can then attach to an Amazon EC2 instance.

All gateway-cached volume data and snapshot data is stored in Amazon S3 encrypted at rest using Server Side Encryption (SSE). However, you cannot access this data using Amazon S3 APIs or with other tools such as the Amazon S3 console.

AWS Storage Gateway: Gateway-Stored Volume Architecture

In the gateway-stored volume solution, you maintain your volume storage on-premises in your data center. That is, you store all your application data on your on-premises storage hardware. The gateway then securely uploads data to the AWS cloud for cost-effective backup and rapid disaster recovery. This is an ideal solution if you want to keep data locally on-premises because you need low-latency access to all your data and maintain backups in AWS.

The following diagram provides an overview of the AWS Storage Gateway's stored volume deployment



Once you've installed AWS Storage Gateway's software appliance (the virtual machine (VM)) on a host in your data center and activated it, you can create gateway *storage volumes* and map them to on-premises Direct Attached Storage (DAS) or Storage Area Network (SAN) disks. You can start with either new disks or disks already holding data. You can then mount these storage volumes to your on-premises application servers as iSCSI devices. As your on-premises applications write data to and read data from a gateway's storage volume, this data is stored and retrieved from the volume's assigned disk.

To prepare data for upload to Amazon S3, your gateway also stores incoming data in a staging area, referred to as an *upload buffer*. You can use on-premises DAS or SAN disks for working storage. Your gateway uploads data from the upload buffer over an encrypted SSL connection to the AWS Storage Gateway service running in the AWS cloud. The service then stores the data encrypted in Amazon S3.

You can take incremental backups, called *snapshots*, of your storage volumes. The gateway stores these snapshots in Amazon S3 as Amazon EBS snapshots. When taking a new snapshot, only the data that has changed since your last snapshot is stored. You can initiate snapshots on a scheduled or ad-hoc basis. When you delete a snapshot, only the data not needed for any other snapshot is removed.

You can restore an Amazon EBS snapshot to an on-premises gateway storage volume in the event that you need to recover a backup of your data. You can also use the snapshot as a starting point for a new Amazon EBS volume, which you can then attach to an Amazon EC2 instance.

Requirements

The AWS Storage Gateway runs as a virtual machine (VM) that you deploy on a host in your data center. The host must be running VMware ESXi Hypervisor (v 4.1 or v 5) or Microsoft Hyper-V 2008 R2. A free version of VMware is available on the [VMware website](#). A free, stand-alone version of Hyper-V is available at the [Microsoft Download Center](#).

Once deployed, the VM will have the following configuration:

- 4 virtual processors assigned to the VM
- 7.5 GB of RAM assigned to the VM
- 75 GB of disk space for installation of VM image and system data

Ensure that your host provides the required hardware for the VM footprint. You also need to provide additional disk space for your application data and disk space for the gateway to use as working storage.

AWS Storage Gateway allows you to create iSCSI storage volumes for your on-premises applications to connect to and store data. AWS Storage Gateway supports the mounting of its storage volumes using the following iSCSI initiators:

- Windows Server 2008 and Windows 7
- Red Hat Enterprise Linux 5

The following list describes the ports required in your AWS Storage Gateway deployment:

- Ports 80 and 443 are used by the vSphere client and the Hyper-V manager to communicate to the host.
- Port 80 is used when you activate your gateway from the AWS Storage Gateway console.
- Port 3260 is the default port that your application server uses to connect to iSCSI targets.

To deploy the VM, provision virtual disks and perform other VM functions that you must connect to your on-premises host's hypervisor. The instructions in this documentation show you how to use the VMware vSphere client and the Microsoft Hyper-V Manager on a Windows client computer to connect to the host and perform these tasks.

Pricing

For current information about pricing, go to the [AWS Storage Gateway Detail Page](#).

AWS Storage Gateway API

All the preceding sections use the AWS Storage Gateway console to perform various gateway configuration and management tasks. Additionally, you can use AWS Storage Gateway API to programmatically configure and manage your gateways. For more information about the API, see [API Reference for AWS Storage Gateway \(p. 283\)](#). You can also use the AWS SDKs when developing applications with AWS Storage Gateway. The AWS SDKs for Java, .NET, and PHP wrap the underlying AWS Storage Gateway API, simplifying your programming tasks. For information about downloading the SDK libraries, go to [Sample Code Libraries](#).

Getting Started with AWS Storage Gateway

Topics

- [Getting Started Requirements for AWS Storage Gateway \(p. 8\)](#)
- [Getting Started Video for AWS Storage Gateway \(p. 8\)](#)
- [Step 1: Sign Up for AWS Storage Gateway \(p. 8\)](#)
- [Step 2: Try an Example Setup \(p. 8\)](#)
- [Where Do I Go from Here? \(p. 87\)](#)

The Getting Started section provides instructions for setting up an AWS Storage Gateway virtual machine (VM), activate it, and configure it so that you have a working gateway. You test the setup by saving sample data locally to your storage volume over an iSCSI connection, and taking a point-in-time backup snapshot. The gateway uploads the snapshot to AWS. To complete the getting started exercise, you then restore the snapshot to a new volume and see how AWS Storage Gateway enables you to recover your data.

At the end of the Getting Started exercise, you will have a working gateway with the following sample configuration:

- An AWS Storage Gateway VM deployed on your VMware ESXi hypervisor host or a Microsoft Hyper-V host
- A gateway that is activated for either cached-volumes or stored-volumes
- Your Windows client connected to one of your local storage volumes over iSCSI

Note

As you follow the steps in this Getting Started section, you will be using the **Setup and Activate Gateway** wizard in the AWS Storage Gateway console. At several steps in the wizard, you perform tasks outside of the console and then return. If your session times out or the browser closes, you can always return to the console to continue from your last step.

Getting Started Requirements for AWS Storage Gateway

To deploy, configure, and test your AWS Storage Gateway setup as described here, you need a host to deploy the AWS Storage Gateway VM. You also need a client to deploy the gateway VM on the host and test the setup. For more information, see [Requirements \(p. 6\)](#).

The Getting Started exercise assumes that Dynamic Host Configuration Protocol (DHCP) is used for the automatic configuration of the gateway IP address. If the environment in which you are deploying the AWS Storage Gateway requires that you specify a static IP address for the gateway, you can do so. For more information about configuring your gateway to use static IP addresses, see [Configuring Your AWS Storage Gateway to Use a Static IP Address \(p. 239\)](#).

Getting Started Video for AWS Storage Gateway

Before you begin this tutorial, you can review this getting started video for the end-to-end setup experience: [Getting Started with AWS Storage Gateway](#)

Step 1: Sign Up for AWS Storage Gateway

When you sign up for an AWS Storage Gateway account, you create an Amazon Web Service (AWS) account that gives you access to all Amazon Web Services, resources, forums, support, and usage reports. You are not charged for any of the services unless you use them. If you already have an account, you can skip this step.

To sign up for AWS Storage Gateway

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

Part of the sign-up procedure involves receiving a phone call and entering a PIN using the phone keypad.

Step 2: Try an Example Setup

Topics

- [Step 2.1: Set Up and Activate AWS Storage Gateway \(p. 9\)](#)
- [Step 2.2: Create Volumes Using the AWS Storage Gateway Console \(p. 61\)](#)
- [Step 2.3: Access Your AWS Storage Gateway Volumes \(p. 72\)](#)
- [Step 2.4: Test the Setup \(p. 79\)](#)

In this getting started exercise, you have two decisions to make that determine the path you will follow. You must decide:

- The on-premises virtualization platform you will use to host the AWS Storage Gateway.
- The type of AWS Storage Gateway you are going to configure, a cached or stored gateway. For more information about these gateway setups, see [How AWS Storage Gateway Works \(p. 3\)](#).

This section provides instructions both supported virtualization platforms (VMware ESXi and Microsoft Hyper-V) and both types of gateway configurations (cached or stored gateway).

Step 2.1: Set Up and Activate AWS Storage Gateway

Topics

- [Set Up and Activate \(VMware Host\) \(p. 9\)](#)
- [Set Up and Activate \(Hyper-V Host\) \(p. 33\)](#)

The getting started exercise requires you to use the AWS Storage Gateway console to download the latest gateway VM and activate your gateway. Go to the console at <http://console.aws.amazon.com/storagegateway>. If you signed up for the service and have not yet activated a gateway, the console shows the following page where you begin deploying the gateway. If you have already activated a gateway, click **Deploy a New Gateway** in the navigation pane to start the **Setup and Activate Gateway** wizard.



The wizard walks you through a series of steps required to deploy and configure your gateway. You first choose a subsection to follow based on the hypervisor you plan to use, either VMware or Hyper-V. The gateway deployment process for the two host types is conceptually similar. After setting up and activating your gateway, the remaining steps (creating and accessing volumes) are the same for both host types.

Set Up and Activate (VMware Host)

In this section, you will provision an on-premises VMware host, download and deploy the gateway VM to the host, configure the gateway, and then activate it.

Provision a VMware Host to Deploy the AWS Storage Gateway VM

In this procedure, you create a VMware host in your data center on which you deploy the gateway virtual machine (VM).

To provision a host

1. Review the minimum host requirements. For more information, see the [Requirements \(p. 6\)](#).
2. Set up a host in your data center with the VMware ESXi hypervisor.

An appendix in this guide provides the minimum instructions to install the hypervisor OS. For more information, see [Appendix B: Configuring a VMware ESXi Host for AWS Storage Gateway](#) (p. 417).

Note

If you plan to deploy AWS Storage Gateway using VMware High Availability (HA) for failover protection, see [Using AWS Storage Gateway with VMware High Availability](#) (p. 92). In this tutorial, you deploy your AWS Storage Gateway VM on a single host with no clustering or failover provision.

Download and Deploy the AWS Storage Gateway VM on Your Host

The AWS Storage Gateway virtual machine is available as a VMware ESX .ova package. This section explains how to download the .ova file locally, deploy it to your host, and synchronize the VM time with the host time.

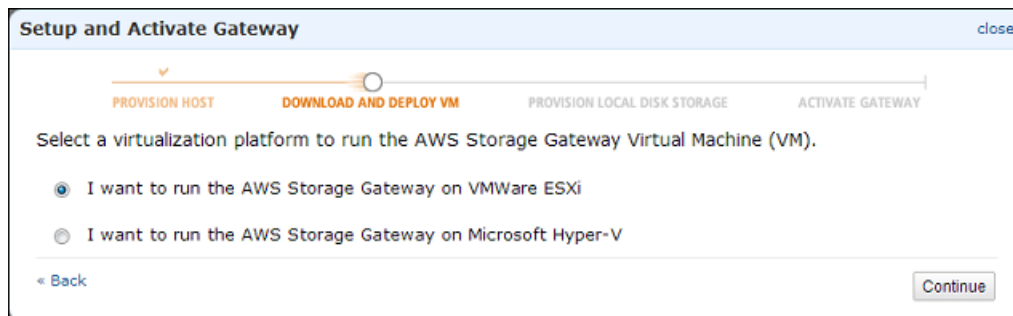
Important

Synchronizing the VM time with the host time is required for successful gateway activation.

Download the AWS Storage Gateway VM

To download the VM

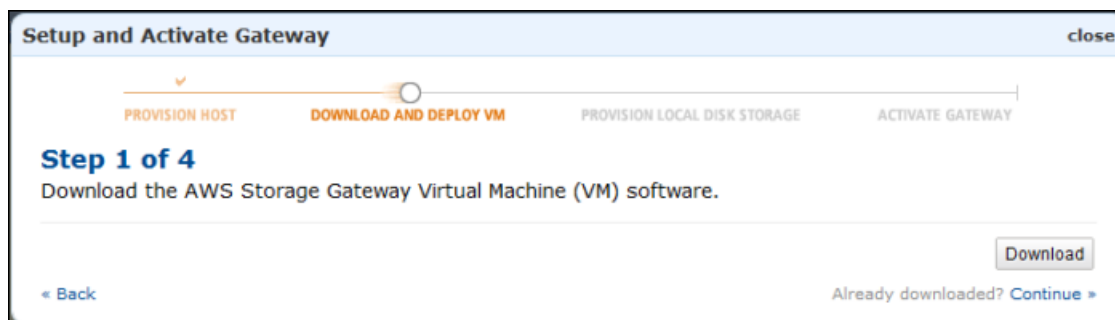
1. In the [AWS Storage Gateway](#) console, in the **Setup and Activate Gateway** wizard, navigate to the **DOWNLOAD AND DEPLOY VM** page.
2. Select **I want to run the AWS Storage Gateway on VMware ESXi** and click **Continue**.



3. Click **Download** to download a .zip file that contains the .ova file. Save the .zip file to a location on your computer.

Note

The .zip file is over 500 MB in size and may take some time to download, depending on your network connection.

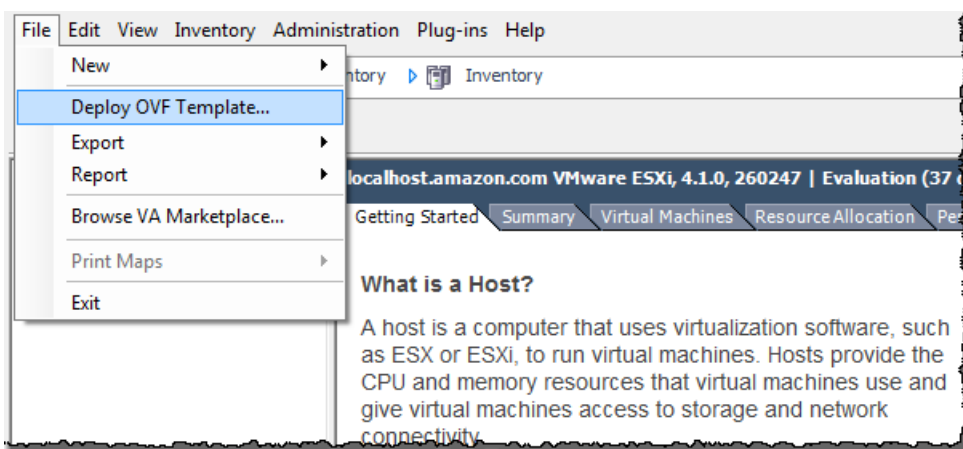


Deploy the AWS Storage Gateway VM to Your Host

1. Connect to your hypervisor host:
 - a. Start the VMware vSphere client on your Windows client.
 - b. In the login dialog box, enter the IP address of your host and your login credentials in the corresponding fields.
 - c. Click **Login**.

Your vSphere client is now connected to your host computer.

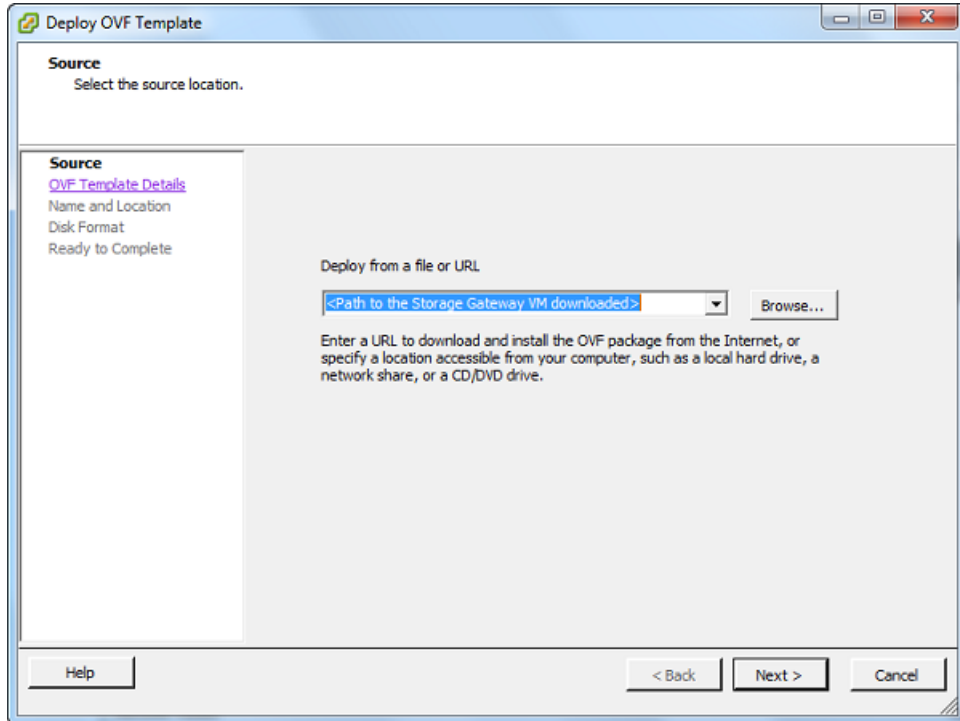
2. Deploy the AWS Storage Gateway VM on the host:
 - a. From the **File** menu of the vSphere client, click **Deploy OVF Template**.



This opens the **Deploy OVF Template** wizard. The wizard is a series of steps for you to provide the required information to deploy the VM.

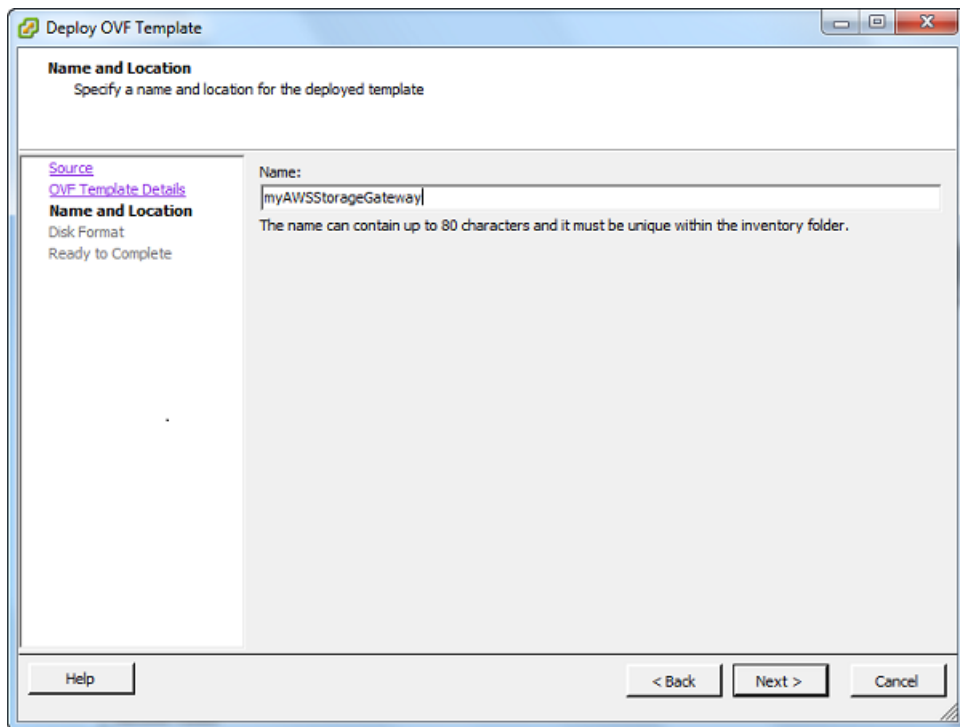
- b. In the **Source** pane, provide the file path to the AWS Storage Gateway .ova package and click **Next**.

AWS Storage Gateway User Guide
Step 2.1: Set Up and Activate a Gateway



- c. In the **OVF Template Details** pane, click **Next**.
- d. In the **Name and Location** pane, enter the VM name in the **Name** field, and then click **Next**.

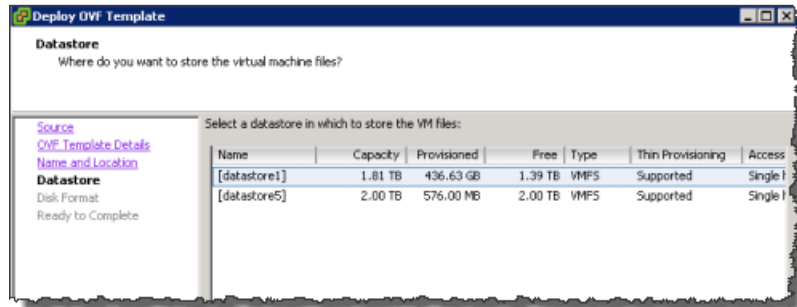
This VM name appears in the vSphere client. However, this name is not used anywhere by AWS Storage Gateway.



AWS Storage Gateway User Guide
Step 2.1: Set Up and Activate a Gateway

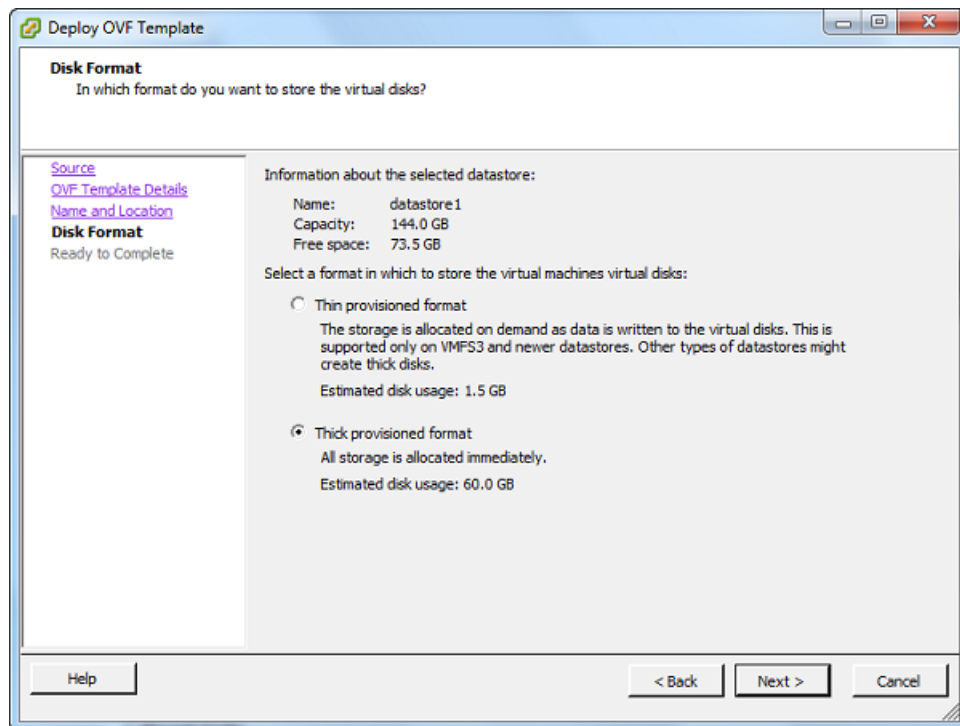
- e. The following **Datastore** pane is displayed only if your host has multiple data stores. In this pane, you select a data store where you want to deploy the VM and click **Next**. Skip to the next step if your host has only one datastore.

A datastore is a virtual representation of underlying physical storage resources. The following example shows a host that has two datastores: datastore1 and datastore2.



- f. In the **Disk Format** pane, select **Thick provisioned format** and click **Next**.


When you use thick provisioning, the disk storage is allocated immediately, resulting in better performance. In contrast, thin provisioning allocates storage on demand.



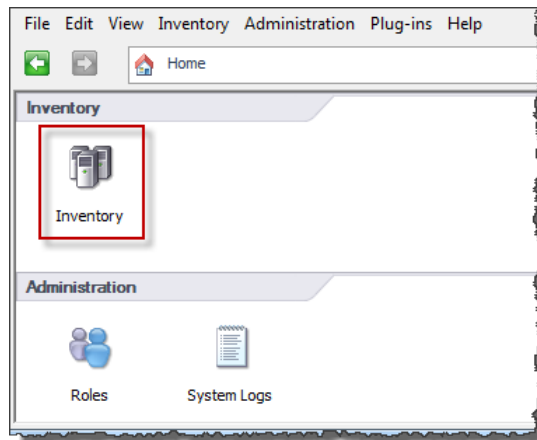
- g. In the **Ready to Complete** pane, click **Finish**.

The AWS Storage Gateway VM starts deploying to your host.

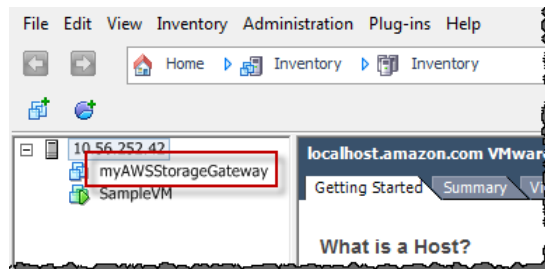
AWS Storage Gateway User Guide
Step 2.1: Set Up and Activate a Gateway



- h. View the details of the new VM.
- i. Depending on the state of your vSphere client, you may need to click the **Inventory** icon first to view the host object that contains the new VM.



- ii. Expand the host object to view the details of the new VM.



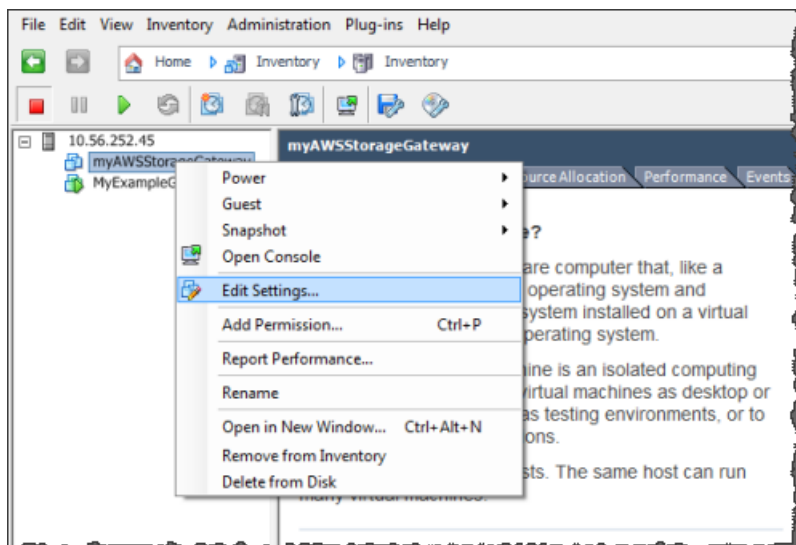
Synchronize VM Time with Host Time

You must ensure that your VM time is synchronized to the host time, and that the host time is correctly set. Synchronizing VM and host times is required for successful gateway activation. In this procedure, you first synchronize the time on the VM to the host time. You then check the host time and, if needed, set the host time and configure the host to synchronize its time automatically to a Network Time Protocol (NTP) server.

To synchronize VM time with host time

1. Configure your VM time.
 - a. In the vSphere client, right-click the name of your gateway VM and select **Edit Settings**.

The **Virtual Machine Properties** dialog box opens.

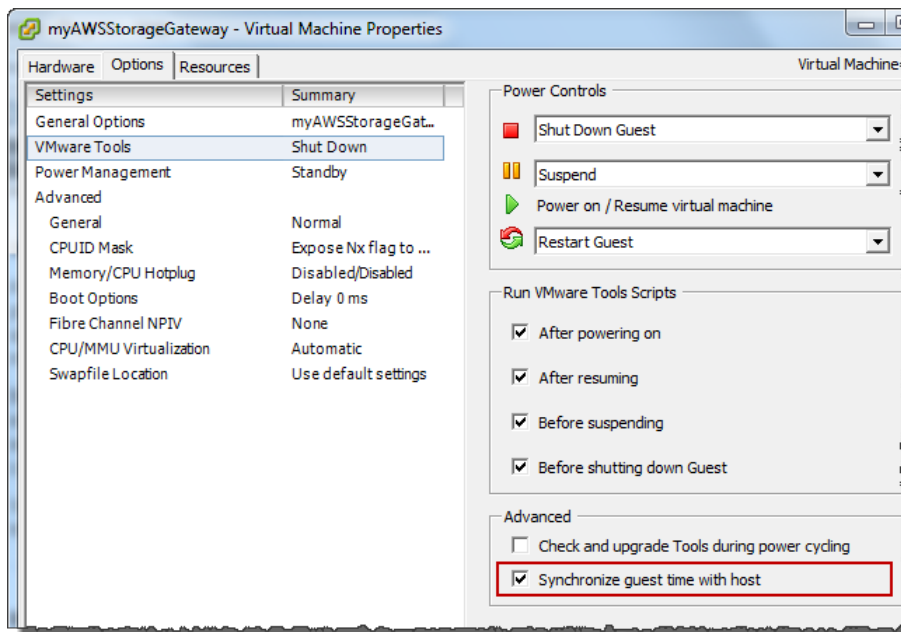


- b. In the **Options** tab, select **VMware Tools** from the options list.
 - c. Check the **Synchronize guest time with host** option and click **OK**.

The VM synchronizes its time with the host.

AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway



2. Configure the host time.

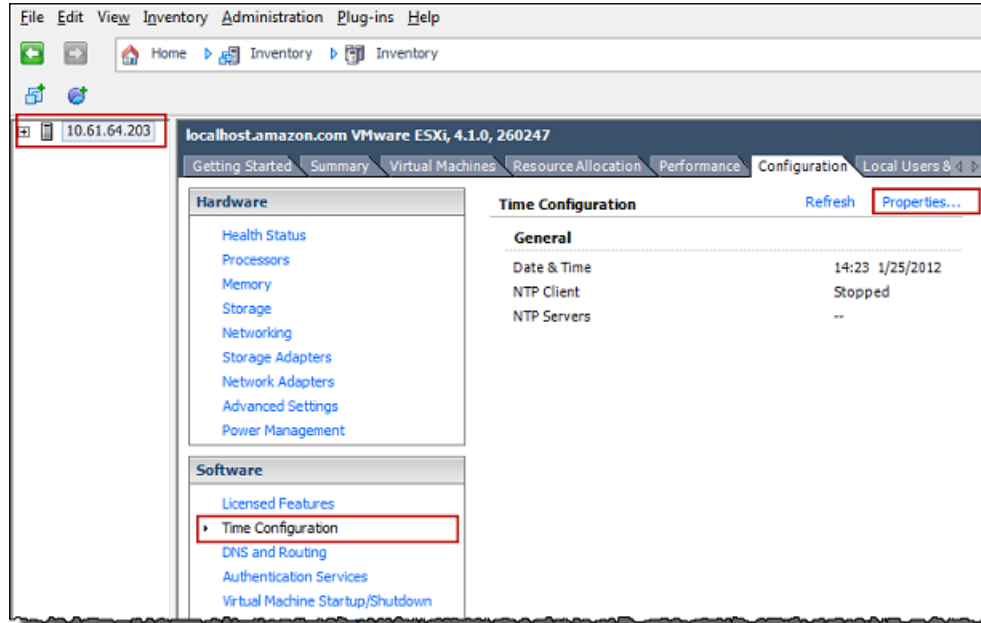
It is important to make sure that your host clock is set to the correct time. If you have not configured your host clock, use the following steps to set and synchronize it with an NTP server.

- In the VMware vSphere Client, select the vSphere host node in the left pane, and select the **Configuration** tab.
- Select **Time Configuration** in the **Software** panel.
- Click the **Properties** link.

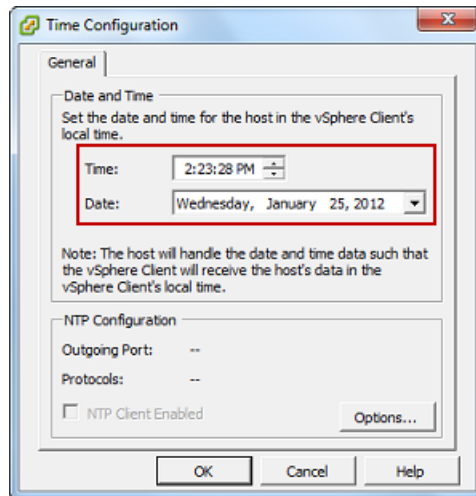
The **Time Configuration** dialog box appears.

AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway




- d. Set the date and time in the **Date and Time** pane.



- e. Configure the host to synchronize its time automatically to a Network Time Protocol (NTP) server:
 - i. Click **Options** in the **Time Configuration** dialog box.
 - ii. In the **NTP Daemon (ntpd) Options** dialog box, select **NTP Settings** in the left pane.


AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway



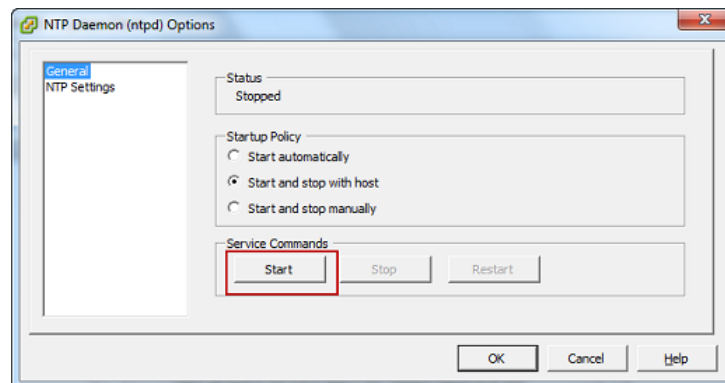
- iii. Click **Add** to add a new NTP server.
- iv. In the **Add NTP Server** dialog box, enter the IP address or the fully qualified domain name of an NTP server and click **OK**.

You can use `pool.ntp.org` as shown in the example.



- v. In the **NTP Daemon (ntpd) Options** dialog box, click **General** in the left pane.
- vi. In the **Service Commands** pane, click **Start** to start the service.

Note that if you change or add another NTP server reference later, you will need to restart the service to use the new server.



- f. Click **OK** to close the **NTP Daemon (ntpd) Options** dialog box.
- g. Click **OK** to close the **Time Configuration** dialog box.

Provision Local Disk Storage for Your AWS Storage Gateway VM

In the AWS Storage Gateway console, in the **Setup and Activate Gateway** wizard, navigate to the **PROVISION LOCAL DISK STORAGE** step. At this step in the console, you will see the following screen shot.



Select the type of iSCSI storage volumes to create on your gateway. You can choose either **Gateway-Cached volumes** or **Gateway-Stored volumes**. Gateway-cached volumes are ideal for corporate file share and backup use cases, where you want to store your volume data in Amazon S3, and just keep recently accessed data on-premises for low-latency access. Gateway-stored volumes are ideal for off-site backups and disaster recovery use cases, where you want to store all your volume data locally for low-latency access to your entire data set, while uploading backups to AWS. For additional information, see [How AWS Storage Gateway Works \(p. 3\)](#).

Depending on the gateway architecture (gateway-cached or gateway-stored) you plan to test, click one of the following links for the next step of instructions.

To...	Do This...
Provision local disks for gateway-cached volumes	Follow the steps in Provision Local Disk Storage (Gateway-Cached Architecture) (p. 19) .
Provision local disks for gateway-stored volumes	Follow the steps in Provision Local Disk Storage (Gateway-Stored Architecture) (p. 24) .

Provision Local Disk Storage (Gateway-Cached Architecture)

In the following steps, you allocate local disks to your deployed gateway VM. After completing these steps, you will have added two virtual disks.

For this Getting Started exercise, you allocate 20 GiB as cache storage and 10 GiB as upload buffer to the VM for exclusive use by the gateway.

Important

In this tutorial, the sizes of the virtual disks you allocate for your VM to use as cache storage and upload buffer are not suitable for real workloads. We strongly recommend that you allocate at least 150 GiB of upload buffer. The size of the cache storage should be based on the size of the upload buffer. In a later step in this tutorial ([Sizing Your Gateway's Storage for Real-World Workloads \(p. 88\)](#)), you will learn about sizing both cache storage and upload buffer appropriately for real workloads.

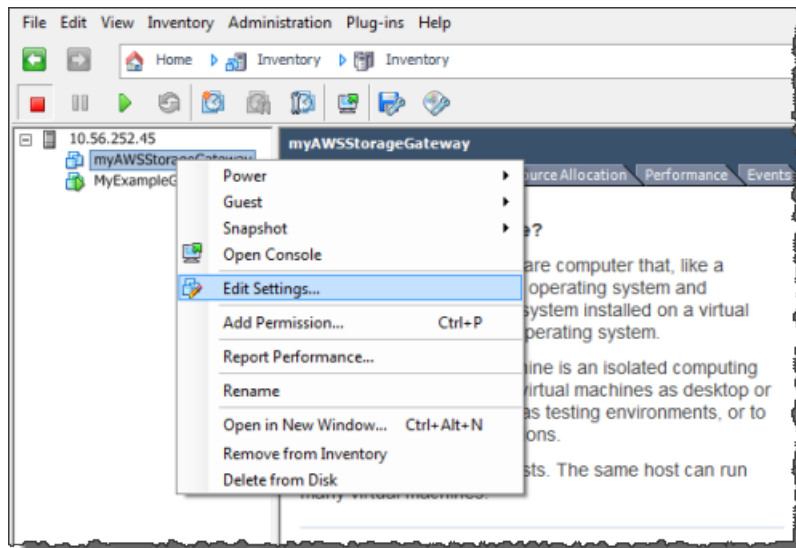
Allocate a Local Disk for Cache Storage

Your frequently accessed application data is maintained locally. You must allocate a disk on the VM as a cache to store this data. This section provides instructions to add a virtual disk from a Direct Attached Storage (DAS) disk. Use the following instructions to provision one virtual disk to store your application data. For instructions on attaching iSCSI volumes from an existing storage area network (SAN) so you can use them in this step, see [To add a new iSCSI target \(p. 418\)](#).

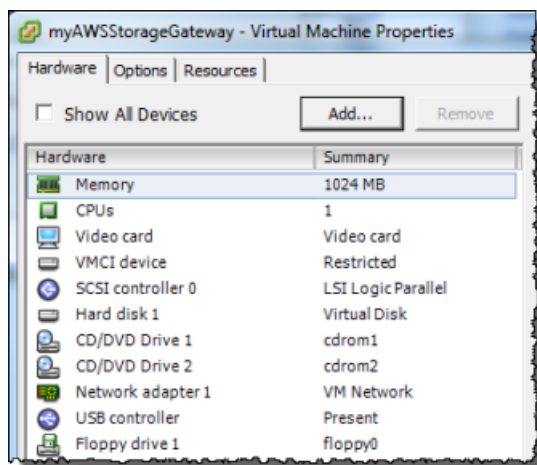
For this getting started exercise, you allocate a 20 GiB virtual disk to the VM.

To allocate a local disk as a cache


1. Start the VMware vSphere client and connect to your host.
2. In the client, right-click the name of your gateway VM and click **Edit Settings**.



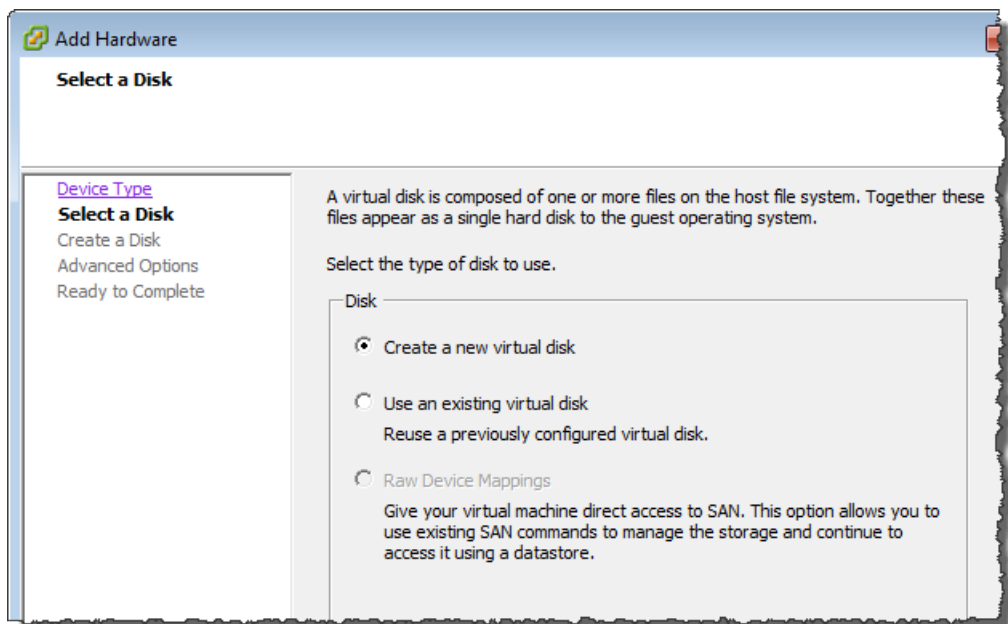
3. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, click **Add** to add a device.



4. Follow the **Add Hardware** wizard to add a disk:
 - a. In the **Device Type** pane, click **Hard Disk** to add a disk, and click **Next**.



- b. In the **Select a Disk** pane, select **Create a new virtual disk** and click **Next**.

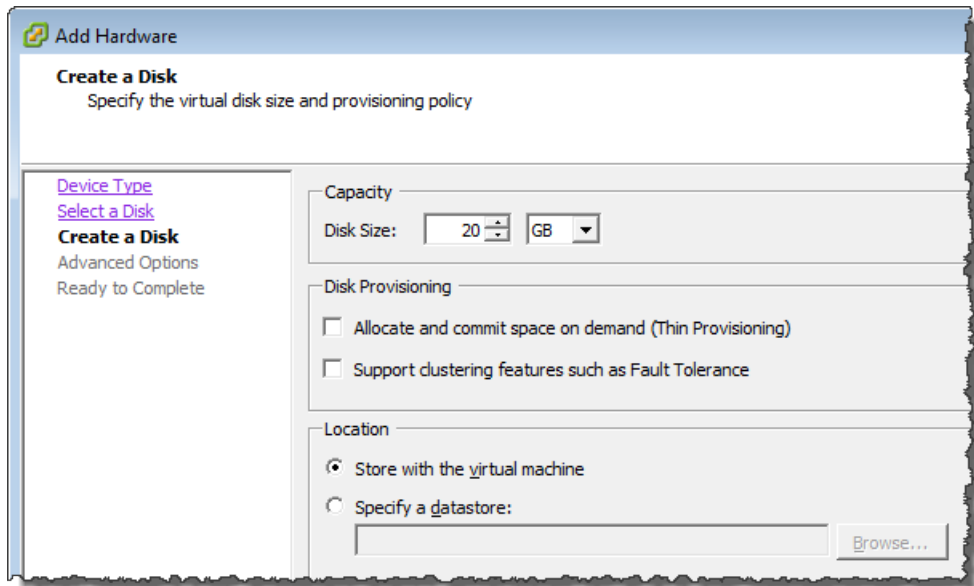


- c. In the **Create a Disk** pane, specify the size of the disk as 20 GiB, and click **Next**.

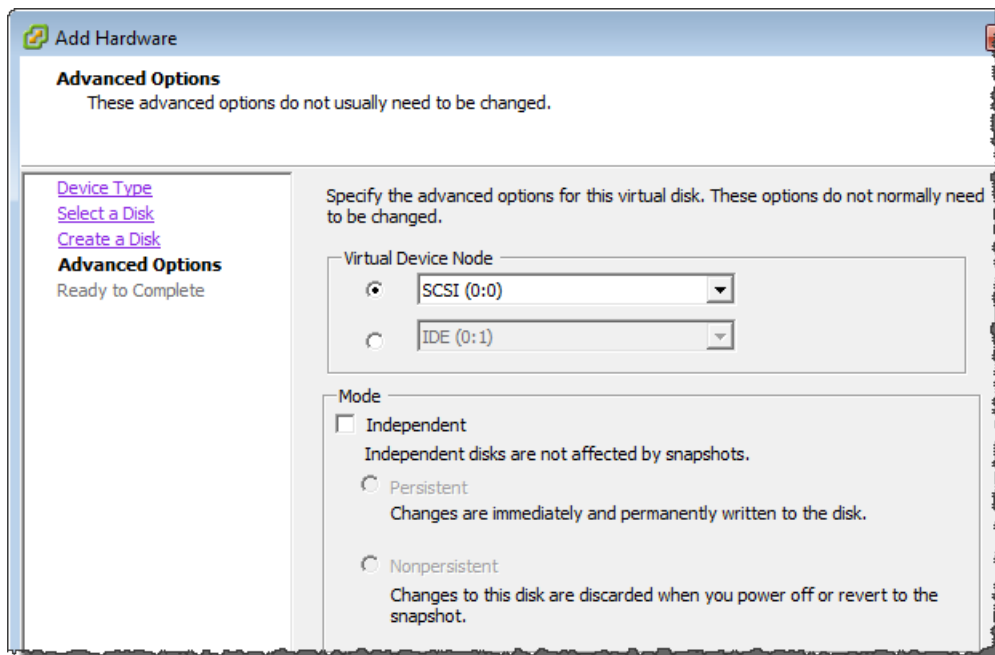
Note

In this example setup, you store the **Location** of the disk with the virtual machine. For real-world workloads, we strongly recommend that you do not provision local disks using the same underlying physical storage disk. Depending on your hosting environment, it may be better to select a different datastore for the disk you provision in this step. For more information, see [Provisioning Local Disks \(Gateway-Cached\)](#) (p. 93).

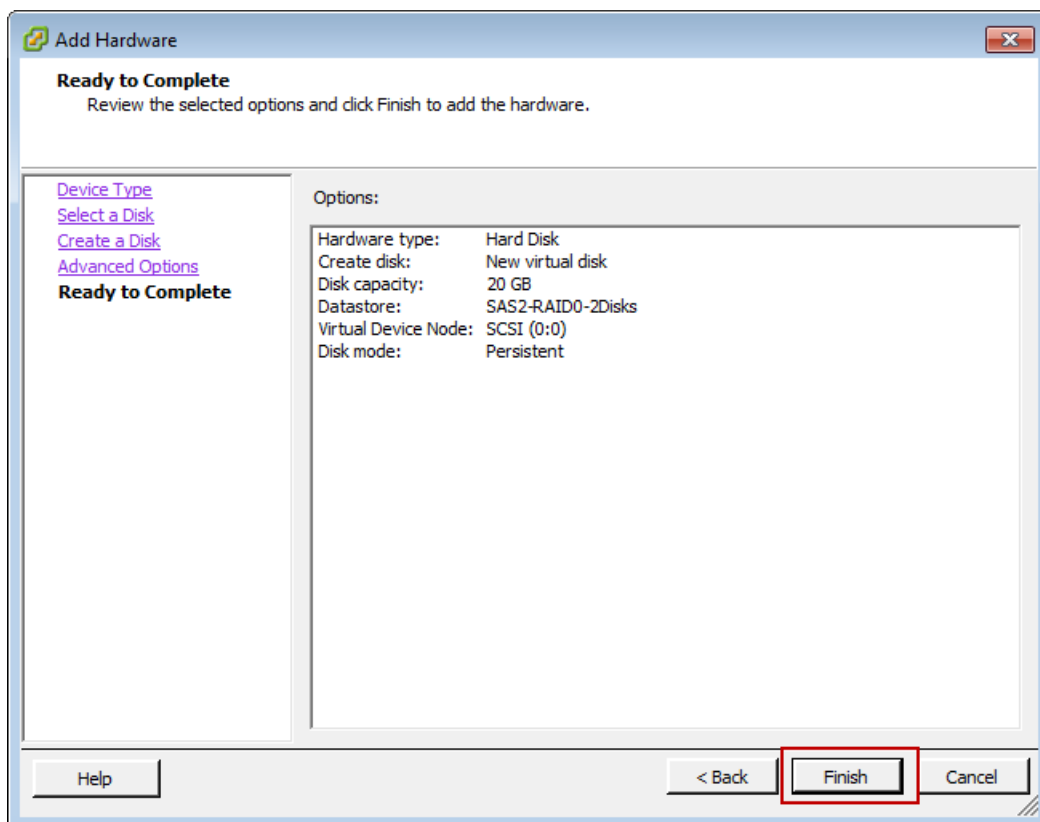
AWS Storage Gateway User Guide
Step 2.1: Set Up and Activate a Gateway



- d. In the **Advanced Options** pane, accept the default values, and click **Next**.



- e. In the **Ready to Complete** pane, accept the default values, and click **Finish**.



- f. In the **Virtual Machine Properties** dialog box, click **OK** to complete adding the disk.

Allocate a Local Disk for an Upload Buffer

The gateway needs buffer space to temporarily store data as it uploads snapshots to AWS. This is referred to as the upload buffer. You must add virtual disks to the VM exclusively for use by the VM. The size of the upload buffer the gateway needs depends on the cache of frequently-accessed data you specified. For related guidelines, see [Sizing the Upload Buffer \(Gateway-Cached\) \(p. 98\)](#).

For this Getting Started exercise, you allocate a 10 GiB virtual disk to the VM for exclusive use by the gateway. In the **Create a Disk** pane of the wizard, enter 10 GiB for the disk size.

To allocate a local disk as an upload buffer

- Repeat the steps in the preceding section ([To allocate a local disk as a cache \(p. 20\)](#)) to add another virtual disk to the gateway.

Verify the Gateway VM Has Two Disks

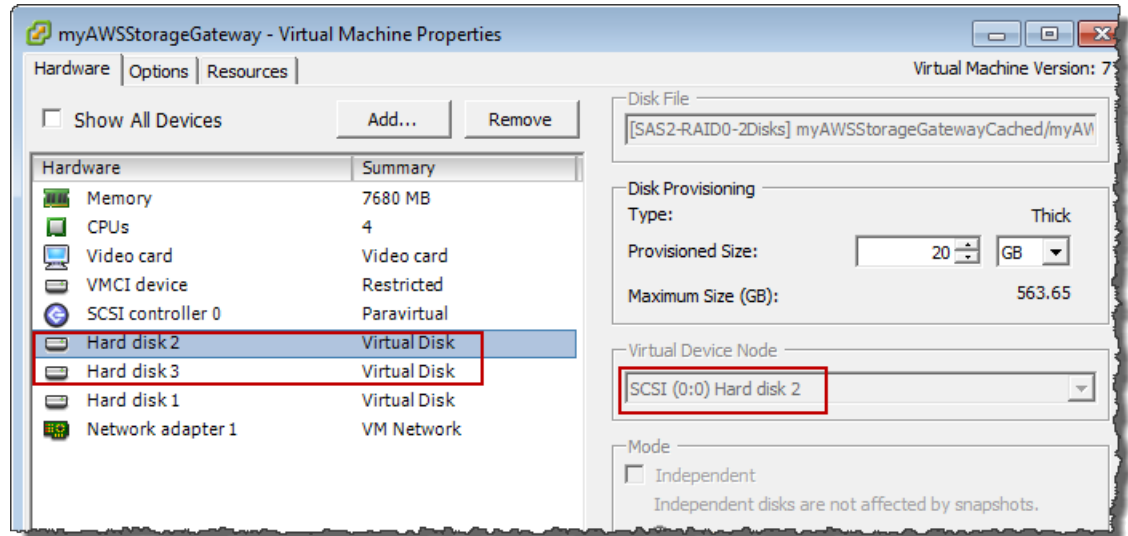
The remainder of the Getting Started exercise requires that you have allocated two disks to your gateway VM. You can use the following optional procedure to verify that you have allocated two disks to your gateway VM. If you need to allocate another disk, repeat the steps in the [To allocate a local disk as a cache \(p. 20\)](#) procedure.

To verify the VM has two disks

1. In the client, right-click the name of your gateway VM and click **Edit Settings**.

2. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, verify that **Hard disk 2** and **Hard disk 3** appear in hardware list.

These two disks will be used later in the AWS Storage Gateway console and appear as SCSI (0:0) and SCSI (0:1) in drop-down lists.



Provision Local Disk Storage (Gateway-Stored Architecture)

In the following steps, you allocate local disks to your deployed gateway VM. After completing these steps, you will have added two virtual disks.

Allocate a Local Disk for Volume Storage (for Your Application Data)

All your application data is maintained locally. You must allocate a disk on the VM to store your application data. This section provides instructions to add a virtual disk from a Direct Attached Storage (DAS) disk. Use the following instructions to provision one virtual disk to store your application data. For instructions on attaching iSCSI volumes from an existing storage area network (SAN) so you can use them in this step, see [To add a new iSCSI target \(p. 418\)](#).

For this getting started exercise, you allocate a 2 GiB virtual disk to the VM for storing application data and a 10 GiB upload buffer to the VM for exclusive use by the gateway.

Important

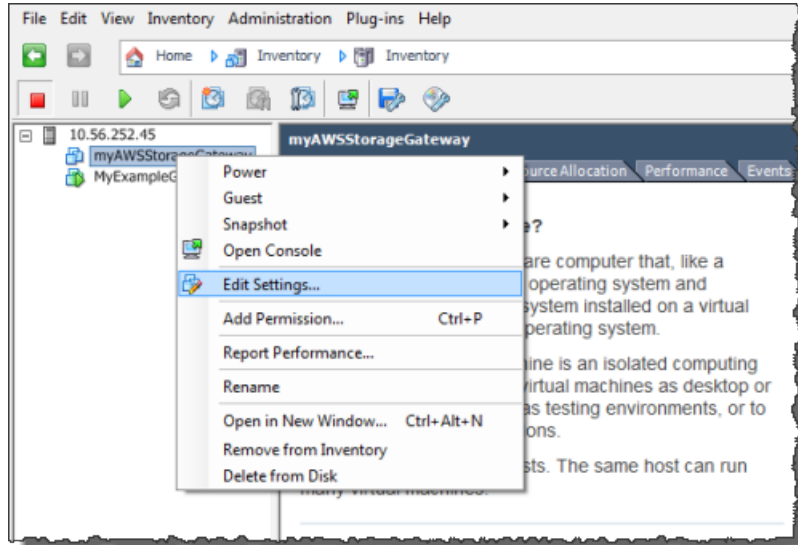
The 10 GiB virtual disk you allocate for your VM to use as the upload buffer in this tutorial is not suitable for real-world workloads. It is strongly recommended that you allocate at least 150 GiB of upload buffer. In a later step in this tutorial ([Sizing Your Gateway's Storage for Real-World Workloads \(p. 88\)](#)), you will learn about sizing the upload buffer appropriately for real workloads.

To allocate a local disk to store your application data

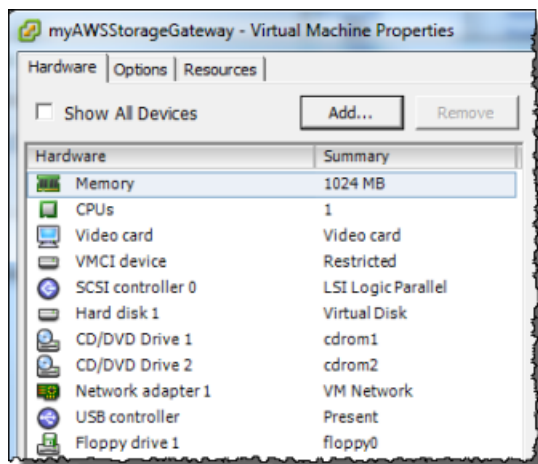
1. Start the VMware vSphere client and connect to your host.
2. In the client, right-click the name of your gateway VM and click **Edit Settings**.

AWS Storage Gateway User Guide

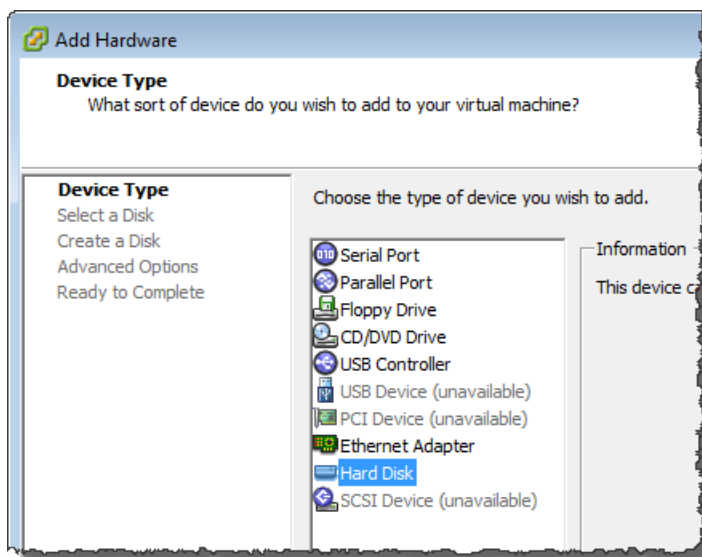
Step 2.1: Set Up and Activate a Gateway



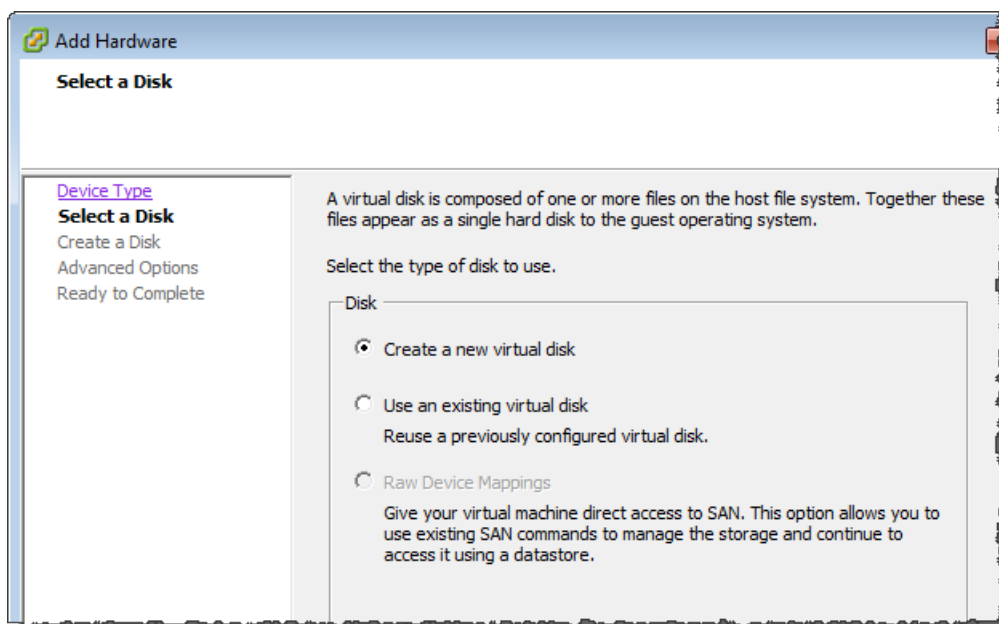
3. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, click **Add** to add a device.



4. Follow the **Add Hardware** wizard to add a disk:
 - a. In the **Device Type** pane, click **Hard Disk** to add a disk, and click **Next**.



- b. In the **Select a Disk** pane, select **Create a new virtual disk**, and click **Next**.

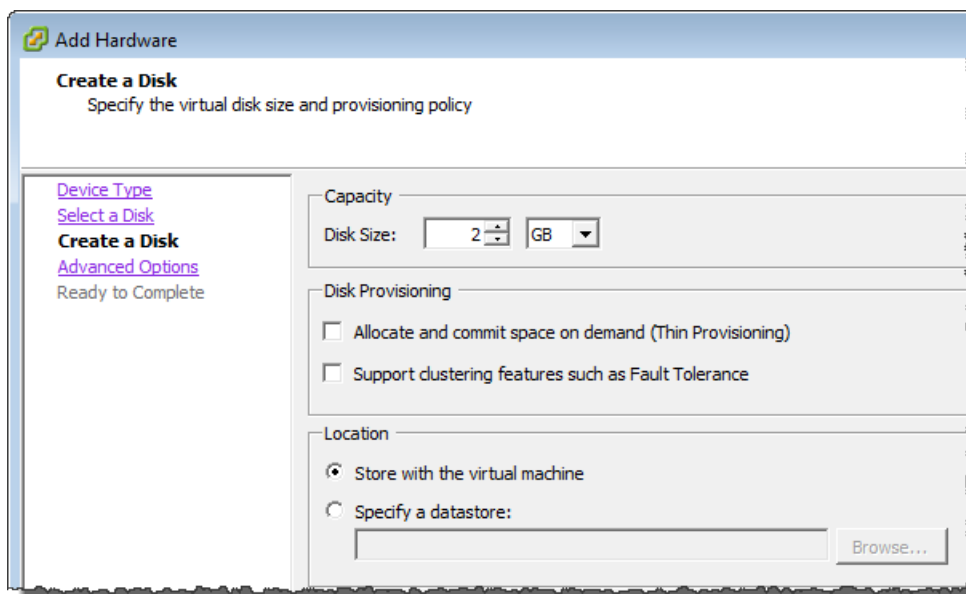


- c. In the **Create a Disk** pane, specify the size of the disk as 2 GiB, and click **Next**.

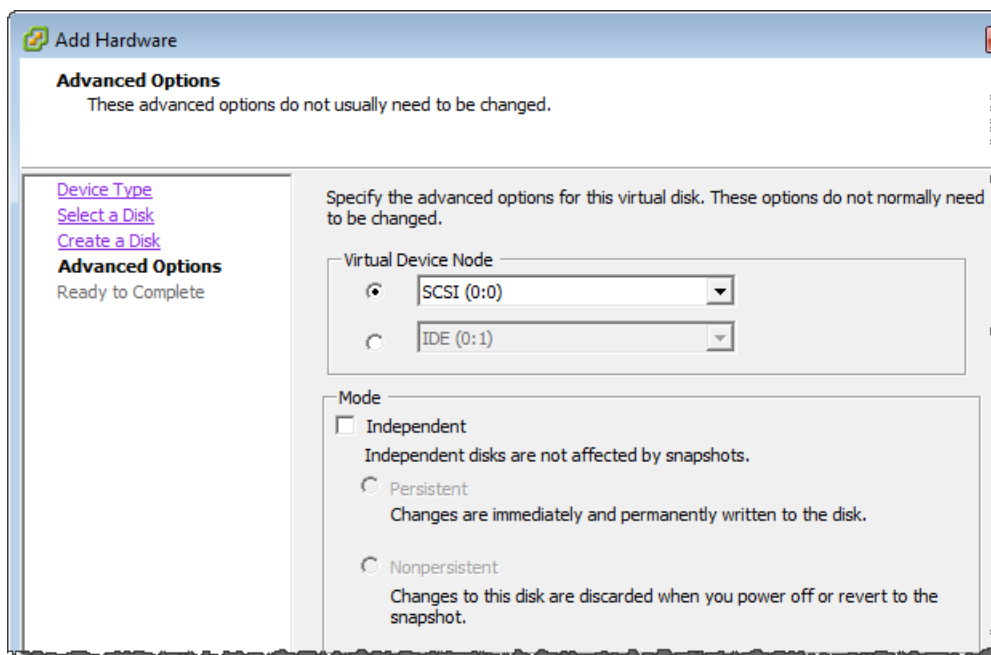
Note

In this example setup, you store the **Location** of the disk with the virtual machine. For real-world workloads, we strongly recommend that you do not provision local disks using the same underlying physical storage disk. Depending on your hosting environment, it might be better to select a different datastore for the disk you provision in this step. For more information, see [Provisioning Local Disks \(Gateway-Stored\)](#) (p. 102).

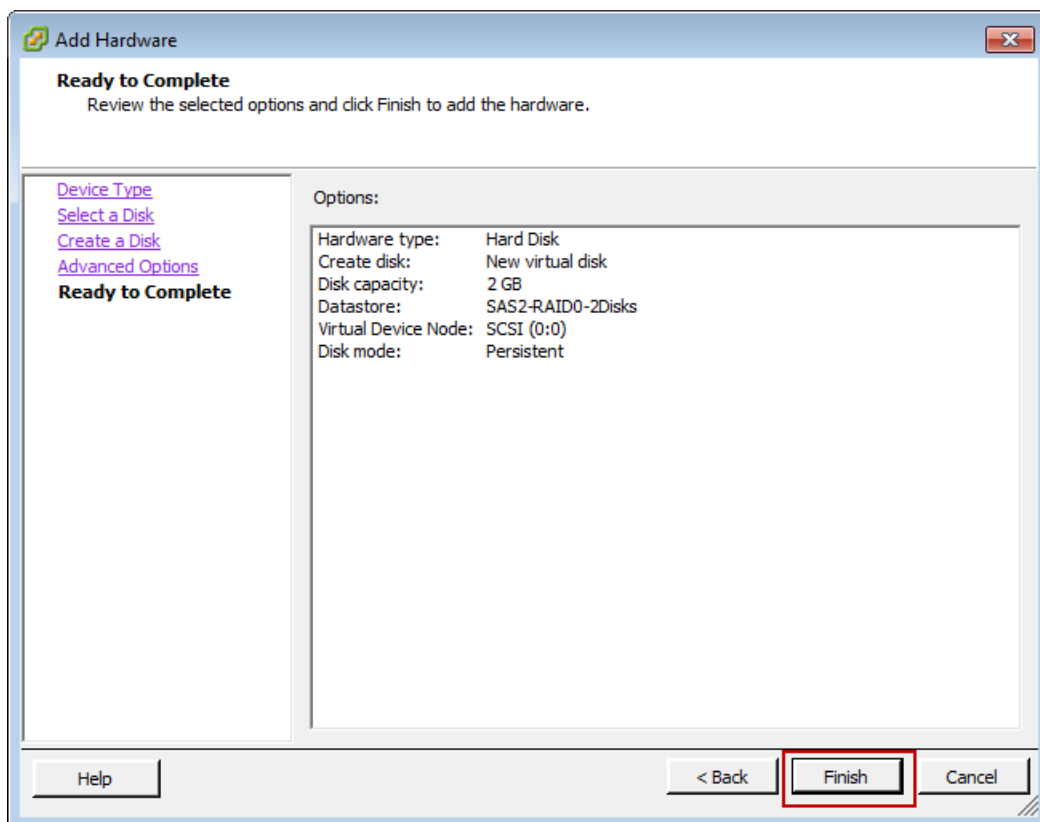
AWS Storage Gateway User Guide
Step 2.1: Set Up and Activate a Gateway



- d. In the **Advanced Options** pane, accept the default values, and click **Next**.



- e. In the **Ready to Complete** pane, accept the default values, and click **Finish**.



- f. In the **Virtual Machine Properties** dialog box, click **OK** to complete adding the disk.

Allocate a Local Disk for an Upload Buffer

The gateway needs buffer space to temporarily store data as it uploads snapshots to AWS. This is referred to as the upload buffer. You must add virtual disks to the VM exclusively for use by the VM. The size of the upload buffer that the gateway needs depends on the size of the disks that you allocate for storing your data. For related guidelines, see [Sizing the Upload Buffer \(Gateway-Stored\) \(p. 106\)](#).

For this tutorial, you allocate a 10 GiB virtual disk to the VM for exclusive use by the gateway. In the **Create a Disk** pane of the wizard, enter 10 GiB for the disk size.

To allocate a local disk for the upload buffer

- Repeat the steps in the [To allocate a local disk to store your application data \(p. 24\)](#) procedure to add another virtual disk to the gateway.

Verify the Gateway VM Has Two Disks

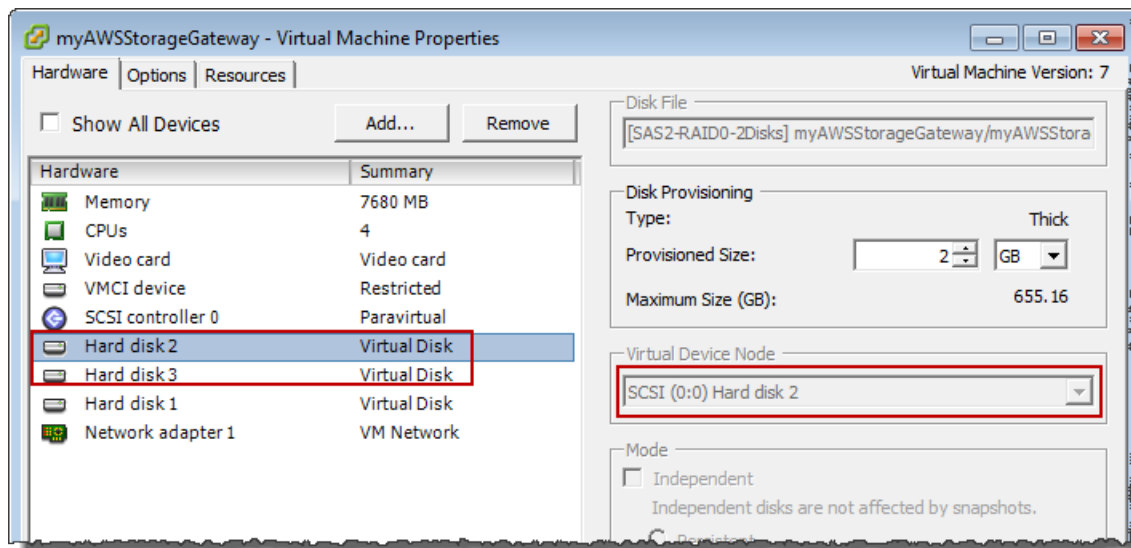
The remainder of this tutorial requires that you have allocated two disks to your gateway VM; use the following optional procedure to verify this. If you need to allocate another disk, repeat the steps in the [To allocate a local disk to store your application data \(p. 24\)](#) procedure.

To verify the VM has two disks

1. In the client, right-click the name of your gateway VM, and click **Edit Settings**.

2. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, verify that **Hard disk 2** and **Hard disk 3** appear in the hardware list.

These two disks will be used later in the AWS Storage Gateway console and appear as SCSI (0:0) and SCSI (0:1) in drop-down lists.



Configure the AWS Storage Gateway VM to Use Paravirtualized Disk Controllers

In this task, the iSCSI controller is set so that the VM uses paravirtualization, a mode where the gateway VM works with the host OS, to enable the console to properly identify the virtual disks that you add to your VM.

Note

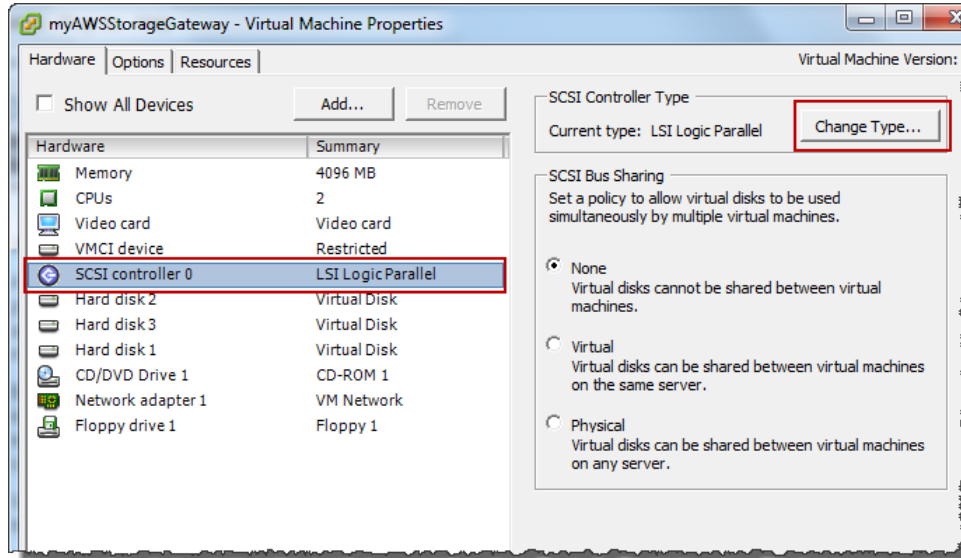
You must complete this step to avoid issues in identifying these disks in the gateway console later when you configure them.

To configure your VM to use paravirtualized controllers

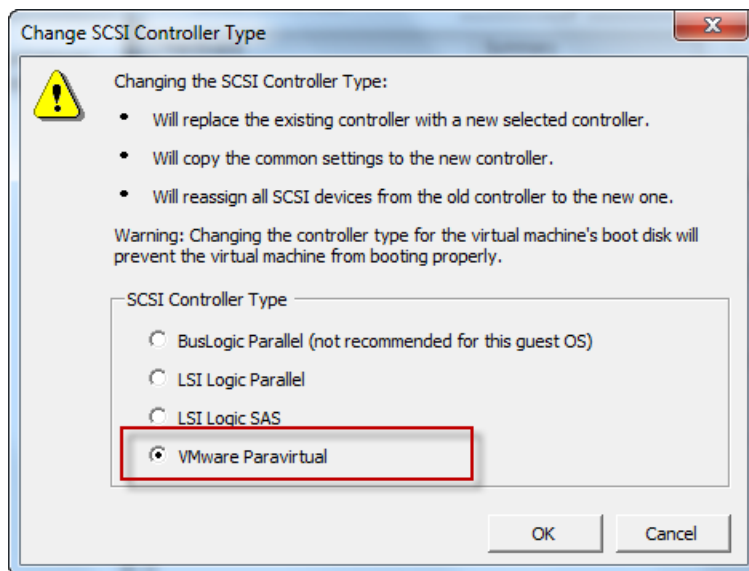
1. In the VMware vSphere client, right-click the name of your gateway VM.
2. Select **Edit Settings**.
3. In the **Virtual Machine Properties** dialog box, click the **Hardware** tab, select the **SCSI controller 0**, and then click **Change Type**.

AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway



4. In the **Change SCSI Controller Type** dialog box, select the **VMware Paravirtual** SCSI controller type, and click **OK**.

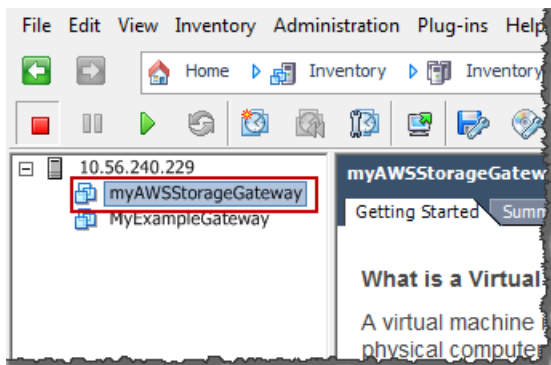


Activate Your Gateway

Now, you are ready to activate your gateway. The activation process associates your gateway with your AWS account. You must power on the gateway VM before you activate your gateway.

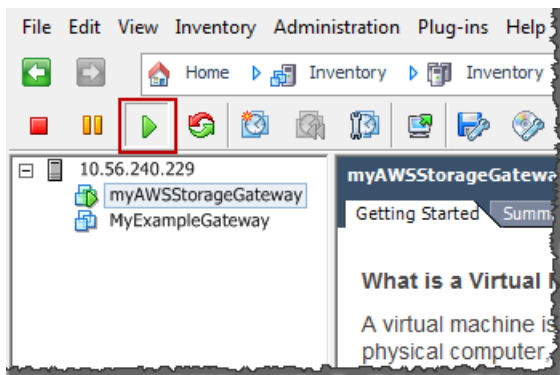
To activate your gateway

1. Power on the VM.
 - a. In the vSphere client, select the gateway VM.



- b. On the **Toolbar** menu, click the **Power On** icon.

Your gateway VM icon now includes a green arrow icon indicating that you have powered on the VM.




2. Activate your gateway.
 - a. Obtain the IP address of your gateway.
 - i. In the vSphere client, select the deployed gateway VM.
 - ii. Click the **Summary** tab for the IP address.

Note

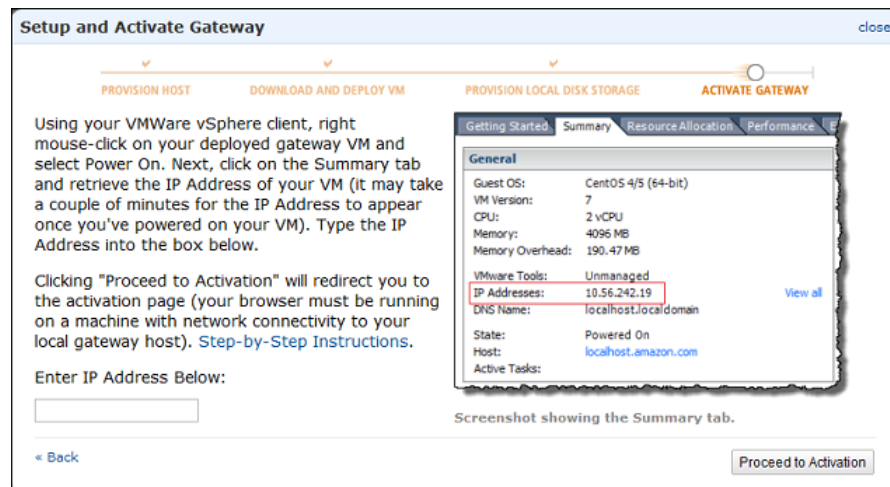
The IP address of your gateway appears as part of the summary. After powering on the VM, it might take a few moments for the IP address to appear.

AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway



- b. Associate your gateway to your AWS account.
 - i. In the AWS Storage Gateway console, in the **Setup and Activate Gateway** wizard, navigate to the following **ACTIVATE GATEWAY** page.
 - A. If the wizard is not already started, click the **Set up and Activate a New Gateway** button.
 - B. Click **Continue** in each wizard step until you reach the **ACTIVATE GATEWAY** page.
 - ii. Enter the IP address of your gateway, and click **Proceed to Activation**.



Note

During activation, your browser connects to the gateway. If activation fails, then check that the IP address you entered is correct. If the IP address is correct, then confirm that your network is configured to allow your browser to access the gateway VM.

- iii. On the activation page fill in the requested information to complete activation.

AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway

The **AWS Region** determines where AWS stores your snapshots. If you choose to restore a snapshot to an Amazon EBS volume, then the Amazon EBS volume must be in the same region as the snapshot. You cannot change the region after the gateway is activated.

The gateway name identifies your gateway in the console. You use this name to manage your gateway in the console, and you can change it post-activation. This name must be unique to your account.

Activating Your AWS Storage Gateway Virtual Machine (VM)

Below is the type and IP address of the gateway you are activating:

Gateway Type: Gateway-Cached Volumes

Activated gateways are billed at \$125 per month, prorated daily. Upon activation of your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies solely to the gateway price. Storage pricing and data transfer pricing continue to apply. The AWS Service Terms are available [here](#).

Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway.

AWS Region: US East (Virginia)

Gateway Time Zone: (GMT -8:00) Pacific Time (US & Canada)

Gateway Name: MyNewGateway

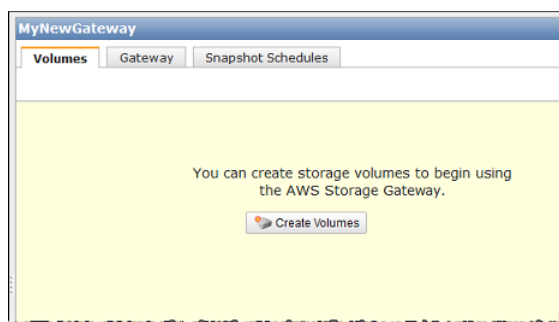
[Activate My Storage Gateway](#)

[Click here](#) if you need to exit the activation process.

- iv. Click **Activate My Storage Gateway**.

Upon successful activation, the **AWS Storage Gateway** console displays a link to the activated gateway under the **Gateways** section of the **Navigation** pane. Click the gateway you just added.

The **Create Volumes** button is displayed.



Set Up and Activate (Hyper-V Host)

In this section, you will provision an on-premises Microsoft Hyper-V host, download and deploy the gateway VM to the host, configure the gateway VM, and then activate it.

Provision a Hyper-V Host to Deploy the AWS Storage Gateway VM

In this procedure, you create a Hyper-V host in your data center on which you deploy the gateway virtual machine (VM).

To provision a host

1. Review the minimum host requirements. For more information, see the [Requirements \(p. 6\)](#).
2. Set up a host in your data center with the Microsoft Hyper-V host.

An appendix in this guide provides the minimum instructions to install the hypervisor OS. For more information, see [Appendix D: Configuring a Microsoft Hyper-V Host for AWS Storage Gateway \(p. 422\)](#).

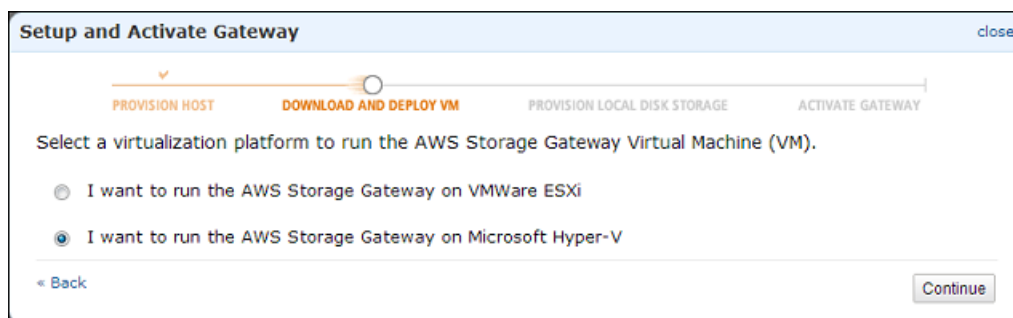
Download and Deploy the AWS Storage Gateway VM on Your Host

The AWS Storage Gateway virtual machine is available as a Hyper-V downloadable .zip file. This section explains how to download the file locally and import it to your host.

Download the AWS Storage Gateway VM

To download the VM

1. In the [AWS Storage Gateway](#) console, in the **Setup and Activate Gateway** wizard, navigate to the **DOWNLOAD AND DEPLOY VM** page.
2. Select **I want to run the AWS Storage Gateway on Microsoft Hyper-V** and click **Continue**.

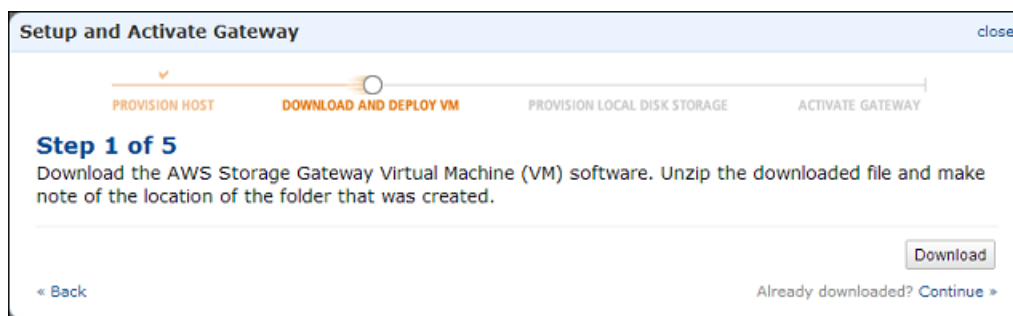


3. Click **Download** to download a .zip file that contains the VM.

Save the .zip file to a location on your computer. Unzip the downloaded file and make note of the location of the folder that was created.

Note

The .zip file is over 500 MB in size and may take some time to download, depending on your network connection.

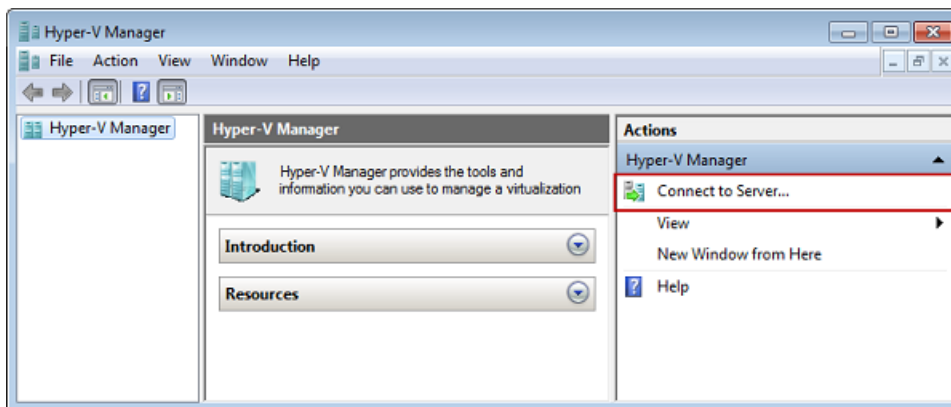


Deploy the AWS Storage Gateway VM to Your Host

To work with your hypervisor host, you must connect to it. After you connect to it, you will specify locations where the VM is stored, import the VM, and then configure a network for the VM.

To connect to the hypervisor host

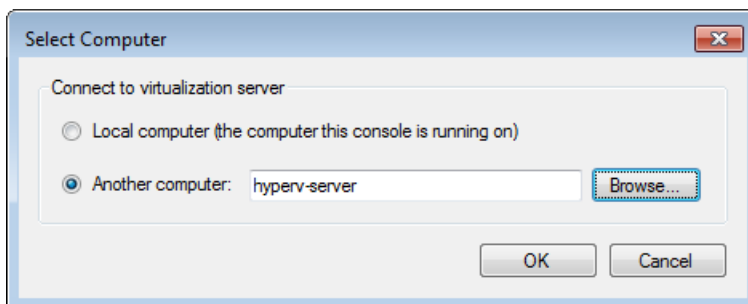
1. Start the Microsoft Hyper-V Manager on your Windows client.
2. In the **Actions** pane, click **Connect to Server...**



3. In the **Select Computer** dialog box, select **Another Computer** and enter the IP address or host name and click **OK**.

Note

In this getting started exercise we use `hyperv-server` as a host. Your host name will be different. If your host name can not be found when you use the **Select Computer** dialog box, you may need to make an entry in your hosts file so that Hyper-V Manager can resolve the server name.



Your Microsoft Hyper-V Manager is now connected to your host computer.

Now that you are connected to your host, the next step is to create folders on the host to store the downloaded source VM, the imported running VM, and associated virtual hard disks for the running VM.


To specify a location for the virtual hard disks and VM

1. Create locations on the hypervisor host for the gateway virtual hard disks and VM.
 - a. Navigate to the hypervisor drive.

For example, using the name of the host in this getting started exercise and assuming that the `c` drive is valid for your host, in the start menu you can type `\\hyperv-server\c$`.

AWS Storage Gateway User Guide
Step 2.1: Set Up and Activate a Gateway


- b. Create a folder called `getting-started` with two subfolders, `unzippedSourceVM` and `gateway`.



- 2. Configure the Hyper-V Manager to point to the `gateway` folder you created.

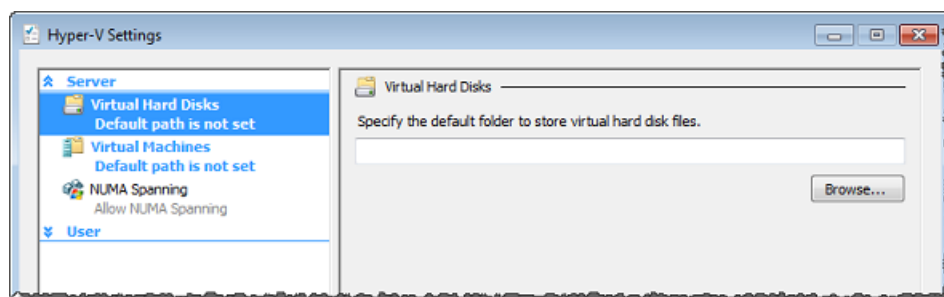
This is the folder that the running VM will use to store its configuration.

- a. In the **Actions** menu, click **Hyper-V Settings....**



- b. In the **Hyper-V Settings** dialog box, configure the location of the virtual hard disks and virtual machines.

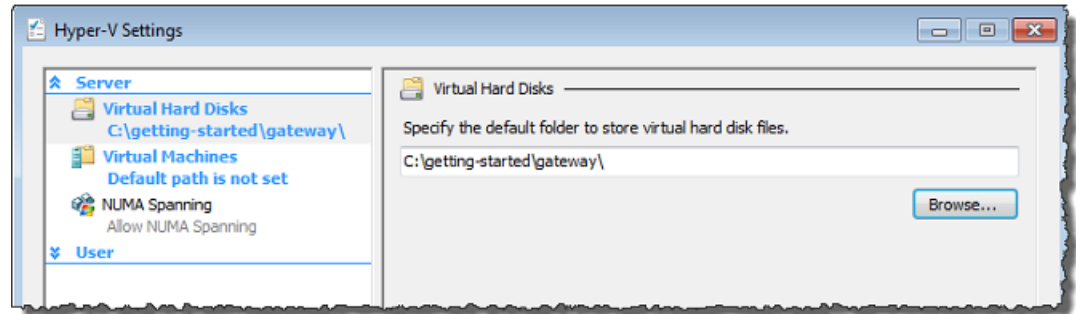
- i. In the left pane, under **Server**, select the **Virtual Hard Disks** setting.



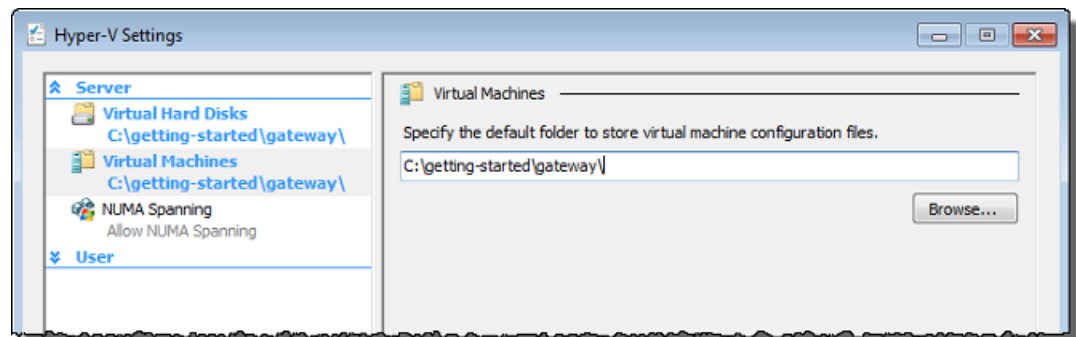
- ii. Browse to find the `gateway` folder you created earlier.

You are browsing on the hypervisor (host) server.

AWS Storage Gateway User Guide
Step 2.1: Set Up and Activate a Gateway




- iii. In the left pane, under **Server**, select the **Virtual Machines** setting.
- iv. Browse to set the location to the same `gateway` folder.



To import the VM

- 1. Copy the unzipped source VM files to the folder you created on the host computer. In this getting started exercise, the path is
`\\hyperv-server\c$\getting-started\unzippedSourceVM\AWS-Storage-Gateway.`
- 2. Import the AWS Storage Gateway VM to the host.
 - a. In the Hyper-V Manager, in the left console tree pane, select the host **hyperv-server**.
 - b. In the **Actions** menu, click **Import Virtual Machine....**



- c. In the **Import Virtual Machine** dialog box:

- i. In the **Location** box, find the location you created previously
`\\hyperv-server\c$\getting-started\unzippedSourceVM\AWS-Storage-Gateway.`

Caution

You must point to the correct folder for the import to succeed. The correct folder to select (`AWS-Storage-Gateway`) will contain three other folders (`Snapshots`, `Virtual Hard Disks`, `Virtual Machines`) and one file (`config.xml`). Depending on how you unzip the gateway source files, you may end up with an extra folder level. For help troubleshooting imports see [Troubleshooting Your Microsoft Hyper-V Setup](#) (p. 429).

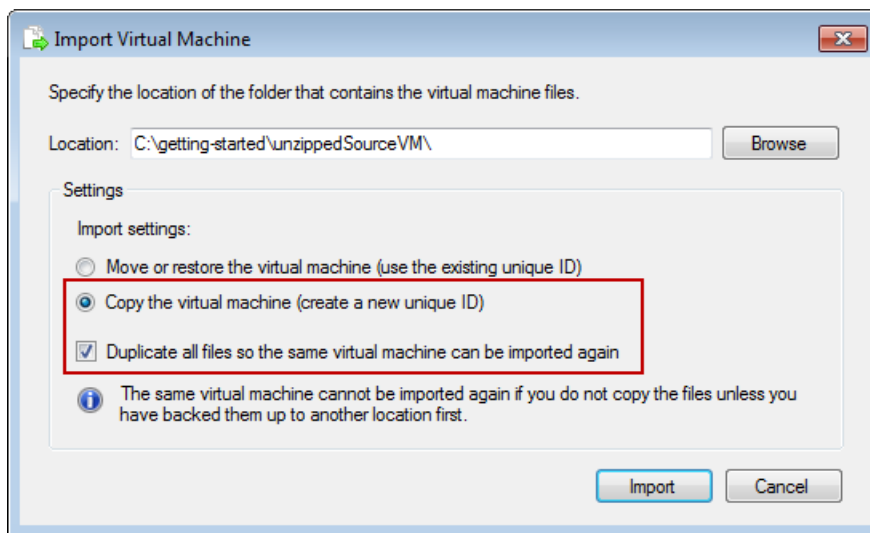
- ii. Select **Copy the virtual machine (create a new unique ID)**.
- iii. Check **Duplicate all files so the same virtual machine can be imported again**.
- iv. Click **Import**.

Caution

It is important to select the **Copy the virtual machine (create a new unique ID)** and **Duplicate all files so the same virtual machine can be imported again** options especially if you intend to reuse the unzipped gateway source files.

Important

You must have 75 GiB of disk space for installation of the VM image and system data.



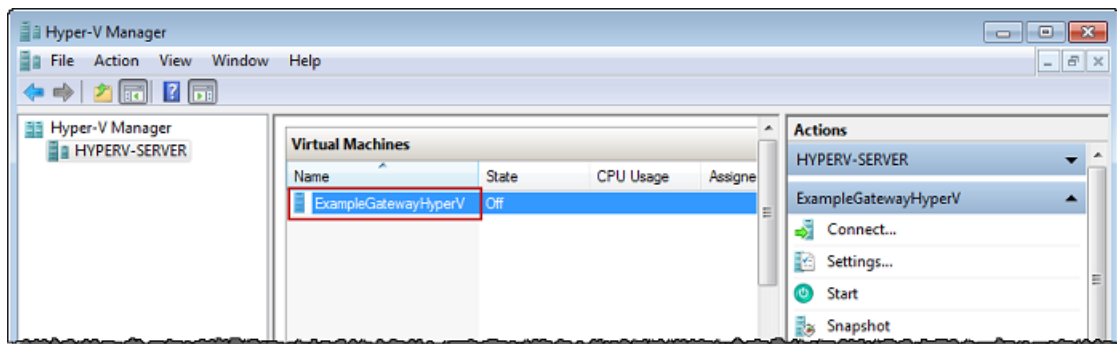
After the import is complete, a virtual machine named `AWS-Storage-Gateway` is created.

3. Rename the virtual machine to avoid confusion with other virtual machines that you might import on the host.
 - a. Select the virtual machine, right-click and select **Rename**.
 - b. Provide a new name for the virtual machine.

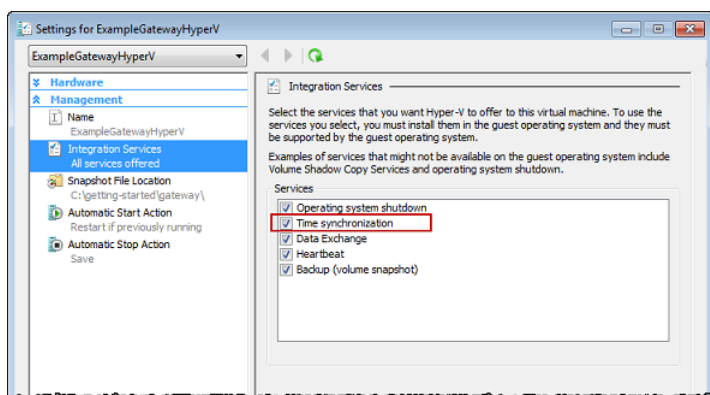
In this getting started exercise, we'll use `ExampleGatewayHyperV`.

AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway



4. Confirm that **Time synchronization** for the VM is selected in **Integration Services**.
 - a. In the **Virtual Machines** list pane, select the virtual machine **ExampleGatewayHyperV**.
 - b. In the **Actions** menu, click **Settings...**
 - c. In the **Settings** dialog box, under **Management**, select **Integration services** and confirm that **Time synchronization** is checked.

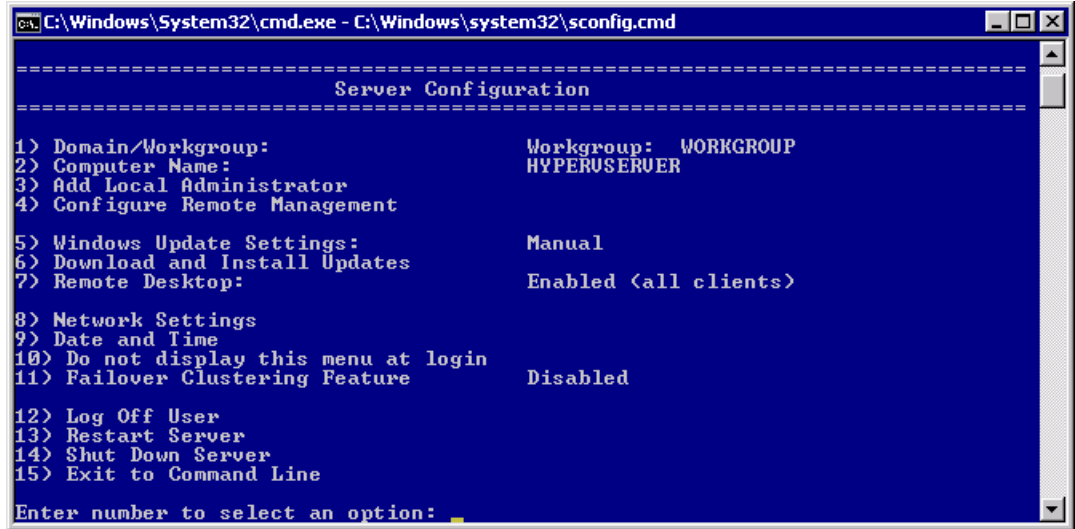


5. Configure the host time if you have not already done so.

It is important to make sure that your host clock is set to the correct time. The following steps show you how to set the time by using the Server Configuration Tool (Sconfig.cmd). For more information on Sconfig.cmd, go to [Configure a Server Core Server with Sconfig.cmd](#). (Depending on the version of Microsoft Hyper-V you are running, you may be able to set the time in other ways.)

- a. Access the Sconfig.cmd tool by either accessing the hypervisor host console or logging in remotely.

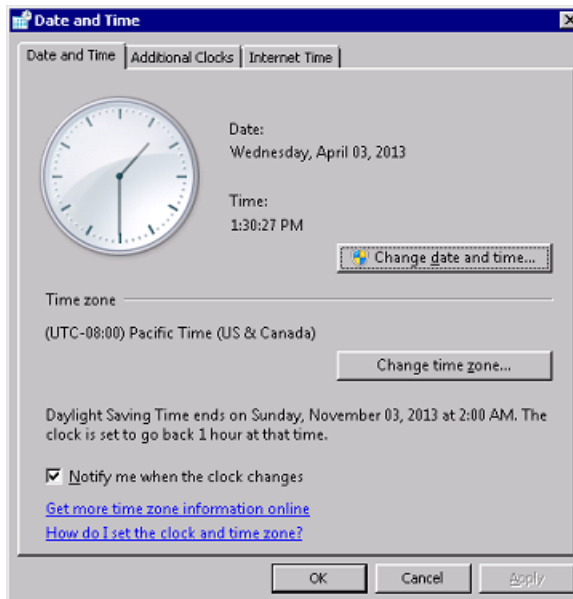
AWS Storage Gateway User Guide
Step 2.1: Set Up and Activate a Gateway



- b. Enter option **9 Date and Time**.

The **Date and Time** control panel is displayed.

- c. Configure the time and click **OK**.



To configure a virtual network and use it for the VM

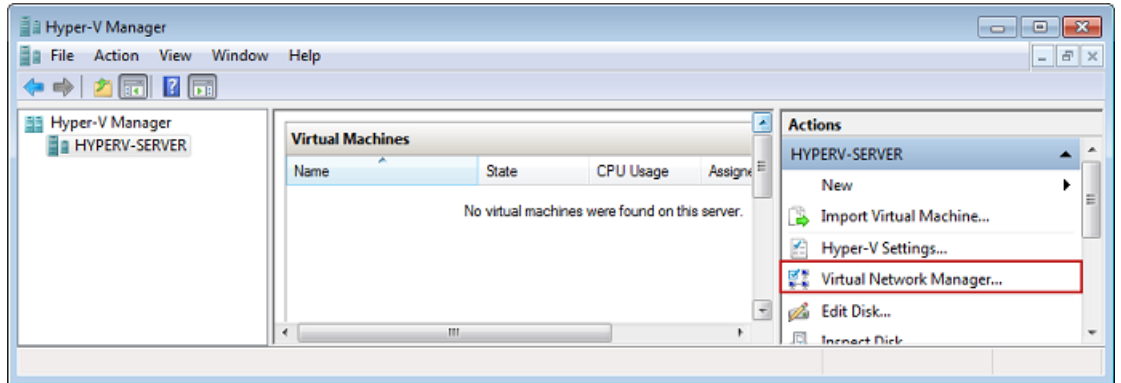
1. Configure virtual network settings for the Hyper-V host.

Note

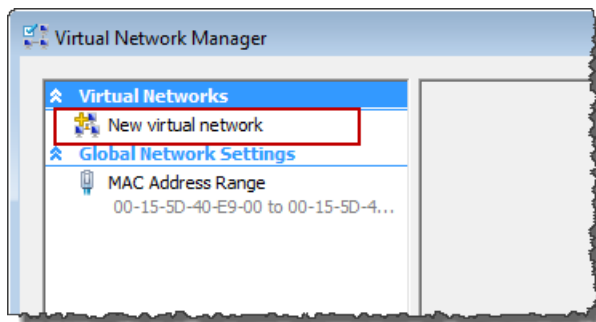
In this Getting Started exercise, we assume the host has not had virtual network settings configured. If you already have a virtual network configured, go to step 2.

- a. In the **Actions** menu, under the hypervisor host name (e.g., `hyperv-server`), click **Virtual Network Manager....**


AWS Storage Gateway User Guide
Step 2.1: Set Up and Activate a Gateway



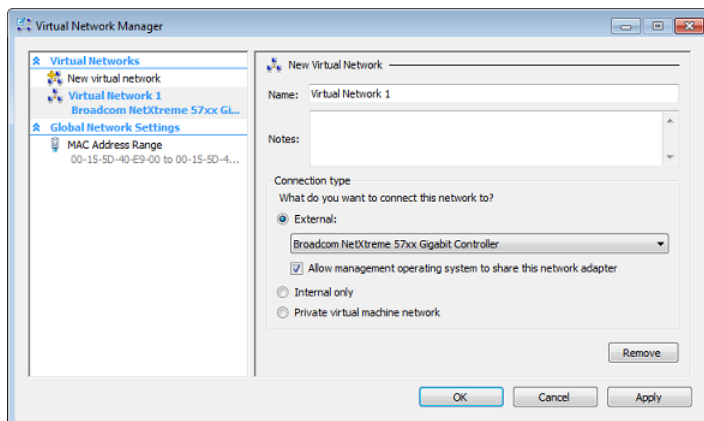
- b. In the **Virtual Network Manager** dialog box, select **New virtual network**.



- c. Select **External** as the virtual network type and click **Add**.



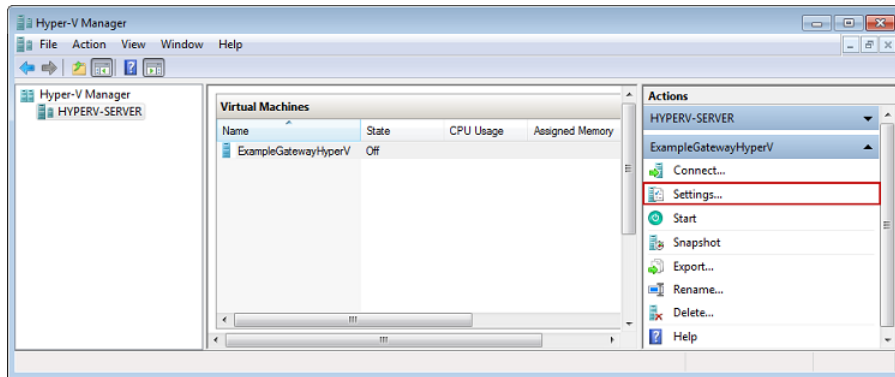
- d. Provide a name for the network, and click **OK**.



AWS Storage Gateway User Guide

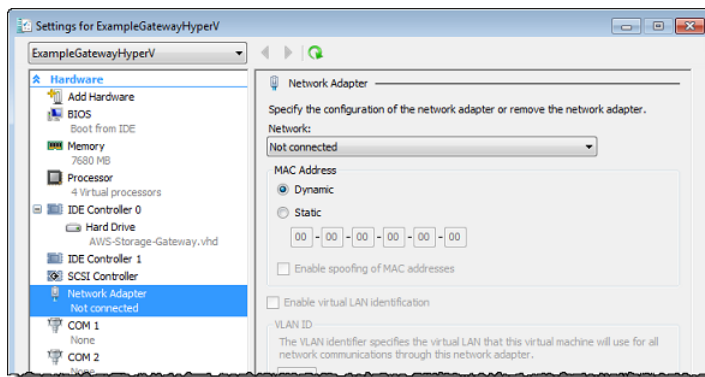
Step 2.1: Set Up and Activate a Gateway

2. Configure the virtual machine to use a virtual network.
 - a. In the **Virtual Machines** list pane, select the virtual machine **ExampleGatewayHyperV**.
 - b. In the **Actions** pane, select **Settings...**



- c. In the **Settings** window, select **Network Adapter**.

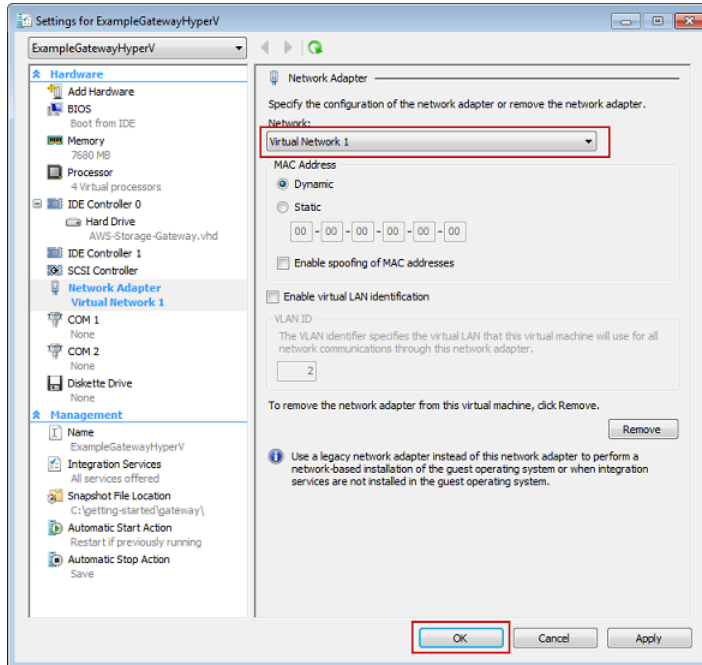
The **Network Adapter** should have a status of **Not connected**.



- d. In the right pane in the **Network** box, select a network.
In the following example, **Virtual Network 1** is selected.

AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway



- e. Click **OK**.

Provision Local Disk Storage for Your AWS Storage Gateway VM

In the AWS Storage Gateway console, in the **Setup and Activate Gateway** wizard, navigate to the **PROVISION LOCAL DISK STORAGE** step. At this step in the console, you will see the following screen shot.



Select the type of iSCSI storage volumes to create on your gateway. You can choose either **Gateway-Cached volumes** or **Gateway-Stored volumes**. Gateway-cached volumes are ideal for corporate file share and backup use cases, where you want to store your volume data in Amazon S3, and just keep recently accessed data on-premises for low-latency access. Gateway-stored volumes are ideal for off-site backups and disaster recovery use cases, where you want to store all your volume data

locally for low-latency access to your entire data set, while uploading backups to AWS. For additional information, see [How AWS Storage Gateway Works \(p. 3\)](#).

Depending on the gateway architecture (gateway-cached or gateway-stored) you plan to test, click one of the following links for the next step of instructions.

To...	Do This...
Provision local disks for gateway-cached volumes	Follow the steps in Provision Local Disk Storage (Gateway-Cached Architecture) (p. 44) .
Provision local disks for gateway-stored volumes	Follow the steps in Provision Local Disk Storage (Gateway-Stored Architecture) (p. 50) .

Provision Local Disk Storage (Gateway-Cached Architecture)

In the following steps, you allocate local disks to your deployed gateway VM. After completing these steps, you will have added two virtual disks.

For this Getting Started exercise, you allocate 20 GiB as cache storage and 10 GiB as upload buffer to the VM for exclusive use by the gateway.

Important

In this tutorial, the sizes of the virtual disks you allocate for your VM to use as cache storage and upload buffer are not suitable for real workloads. We strongly recommend that you allocate at least 150 GiB of upload buffer. The size of the cache storage should be based on the size of the upload buffer. In a later step in this tutorial ([Sizing Your Gateway's Storage for Real-World Workloads \(p. 88\)](#)), you will learn about sizing both cache storage and upload buffer appropriately for real workloads.

Allocate a Local Disk for Cache Storage

Your frequently accessed application data is maintained locally. You must allocate a disk on the VM as a cache to store this data. This section provides instructions to add a virtual hard disk on the host's physical disk. In a real-world application, you should consider using a separate physical disk as the backing storage. For instructions on using a separate physical disk to back a virtual hard disk, see [Adding a Virtual Disk Backed by a Hard Disk \(p. 428\)](#).

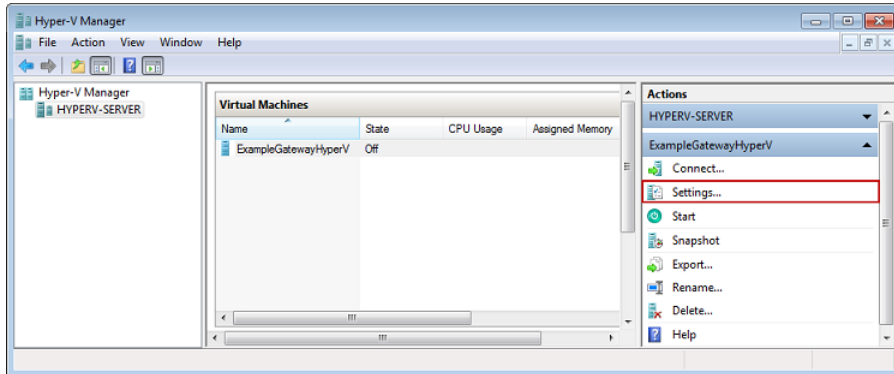
For this getting started exercise, you allocate a 20 GiB virtual disk to the VM.

To allocate a local disk as a cache

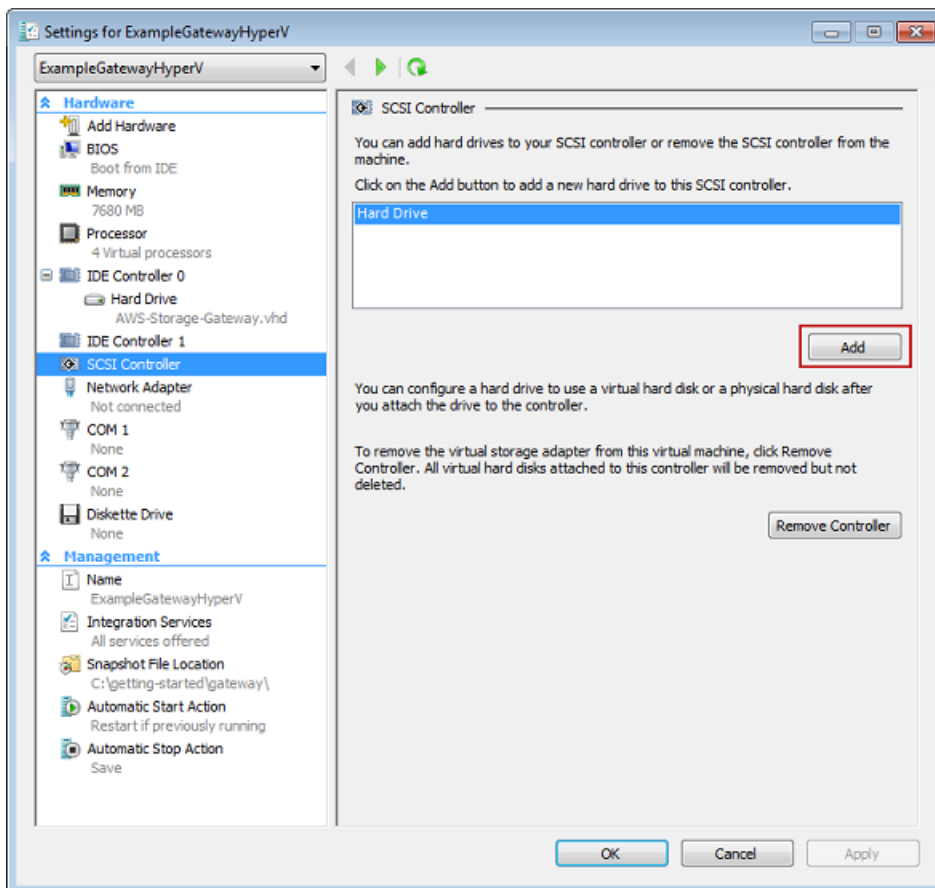
1. Start the Microsoft Hyper-V Manager and connect to the hypervisor.
2. In the **Virtual Machines** list pane, select the virtual machine **ExampleGatewayHyperV**.
3. In the **Actions** pane, select **Settings....**

AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway




4. In the **Settings** window, select **SCSI Controller**, and click **Add**.



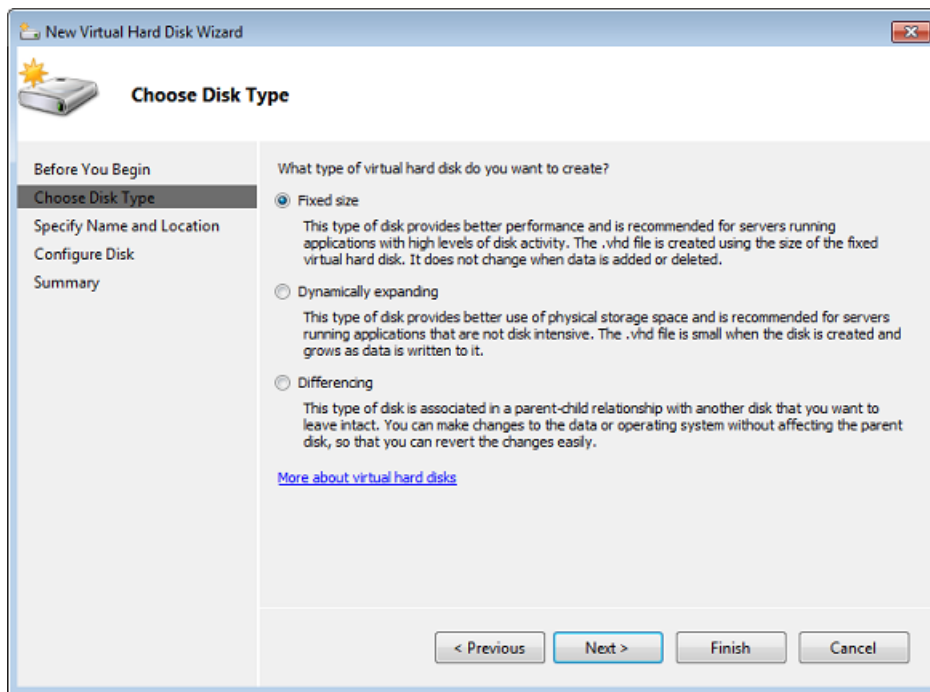
5. In the **Hard Drive** pane, under **Media**, click **New**.

AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway



6. In the **New Virtual Hard Disk Wizard** create a new virtual hard disk.
 - a. On the **Before You Begin** page, click **Next**.
 - b. On the **Choose Disk Type** page, choose **Fixed size**, and click **Next**.



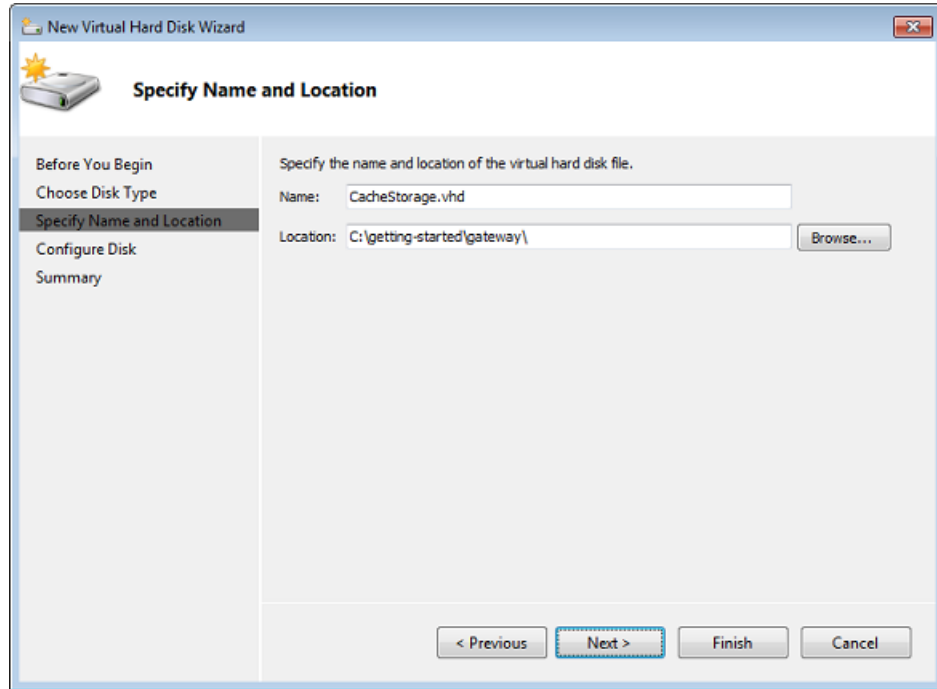
- c. On the **Specify Name and Location** page, specify a name and location for the virtual hard disk.
 - i. Specify `CacheStorage.vhd` as the name.
 - ii. Specify the location as `c:\getting-started\gateway`.

Note

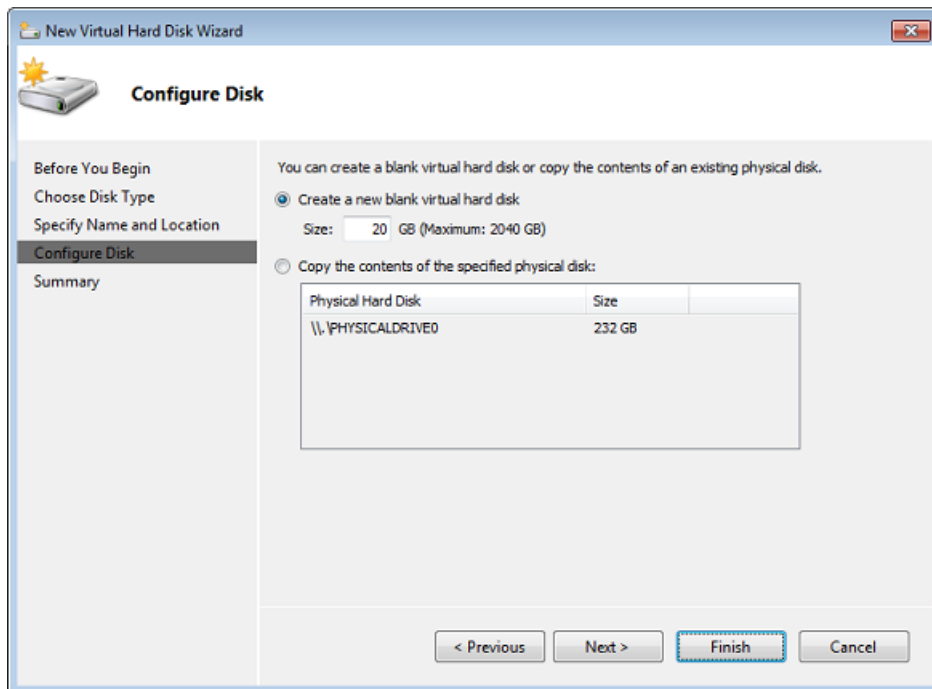
In this example setup, you store the virtual disk with the virtual machine. For real-world workloads, it is strongly recommended that you do not provision local disks using the same underlying physical storage disk. Depending on your hosting environment, performance and portability requirements, it might be better to select a different physical disk in this step. For more information, see [Provisioning Local Disks \(Gateway-Cached\)](#) (p. 93).

- iii. Click **Next**.

AWS Storage Gateway User Guide
Step 2.1: Set Up and Activate a Gateway



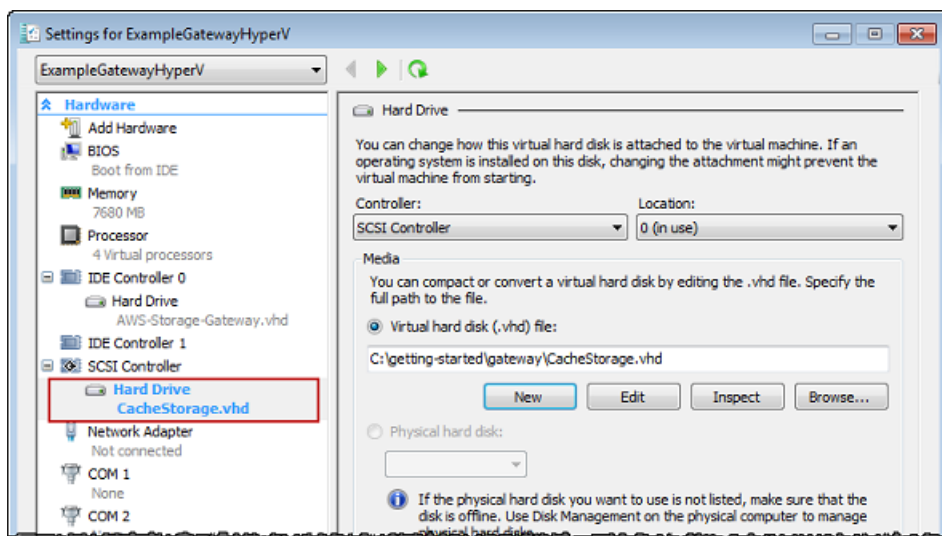
- d. In the **Configure Disk** page, specify the size of the disk as 20 GiB, and click **Finish**.



- e. After the virtual disk is created, verify the **Hard Drive** shows up under **SCSI controller**.
f. Click **SCSI Controller** to prepare to add another hard drive.

Warning

When you add another hard drive, you will need to first click **SCSI Controller** and then follow the steps in this procedure. Clicking **New** when viewing the details of an existing hard drive will replace the existing drive.



7. Click **Ok**.

Allocate a Local Disk for an Upload Buffer

The gateway needs buffer space to temporarily store data as it uploads snapshots to AWS. This is referred to as the upload buffer. You must add virtual disks to the VM exclusively for use by the VM. The size of the upload buffer the gateway needs depends on the cache of frequently-accessed data you specified. For related guidelines, see [Sizing the Upload Buffer \(Gateway-Cached\)](#) (p. 98).

For this Getting Started exercise, you allocate a 10 GiB virtual disk to the VM for exclusive use by the gateway. In the **Create a Disk** pane of the wizard, enter 10 GiB for the disk size.

To allocate a local disk as an upload buffer

- Repeat the steps in the preceding section ([To allocate a local disk as a cache](#) (p. 44)) to add another virtual disk to the gateway. Follow the steps exactly except use the name `UploadBuffer.vhd` for the disk name and 10 GiB for the disk size.

Verify the Gateway VM Has Two Disks

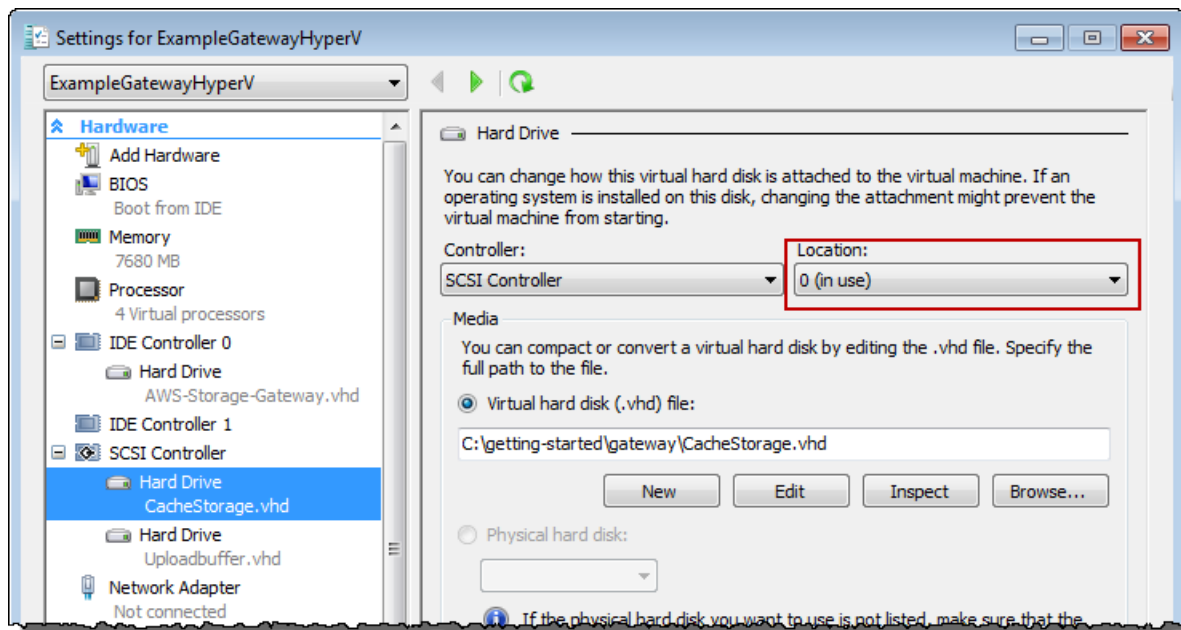
The remainder of the Getting Started exercise requires that you have allocated two disks to your gateway VM. You can use the following optional procedure to verify that you have allocated two disks to your gateway VM. If you need to allocate another disk, repeat the steps in the [To allocate a local disk as a cache](#) (p. 44) procedure.

To verify the VM has two disks

1. Start the Microsoft Hyper-V Manager and connect to the hypervisor.
2. In the **Virtual Machines** list pane, select the virtual machine **ExampleGatewayHyperV**.
3. In the **Actions** pane, select **Settings...**

4. In the **Settings** window, select **SCSI Controller**, and verify that there are two disks.

The two disks you created will be used later in the AWS Storage Gateway console and appear as SCSI (0:0) and SCSI (0:1) in drop-down lists. In the example below, the **CacheStorage.vhd** disk is selected and is SCSI (0:0).



Provision Local Disk Storage (Gateway-Stored Architecture)

In the following steps, you allocate local disks to your deployed gateway VM. After completing these steps, you will have added two virtual disks.

Allocate a Local Disk for Volume Storage (for Your Application Data)

All your application data is maintained locally. You must allocate a disk on the VM to store your application data. This section provides instructions to add a virtual hard disk on the host's physical disk. In a real-world application, you should consider using a separate physical disk as the backing storage. For instructions on using a separate physical disk to back a virtual hard disk, see [Adding a Virtual Disk Backed by a Hard Disk](#) (p. 428).

For this getting started exercise, you allocate a 2 GiB virtual disk to the VM for storing application data and a 10 GiB upload buffer to the VM for exclusive use by the gateway.

Important

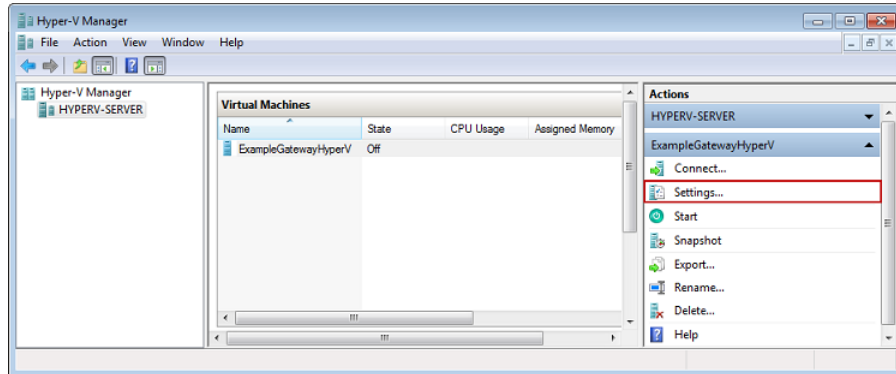
The 10 GiB virtual disk you allocate for your VM to use as the upload buffer in this tutorial is not suitable for real-world workloads. It is strongly recommended that you allocate at least 150 GiB of upload buffer. In a later step in this tutorial ([Sizing Your Gateway's Storage for Real-World Workloads](#) (p. 88)), you will learn about sizing the upload buffer appropriately for real workloads.

To allocate a local disk to store your application data

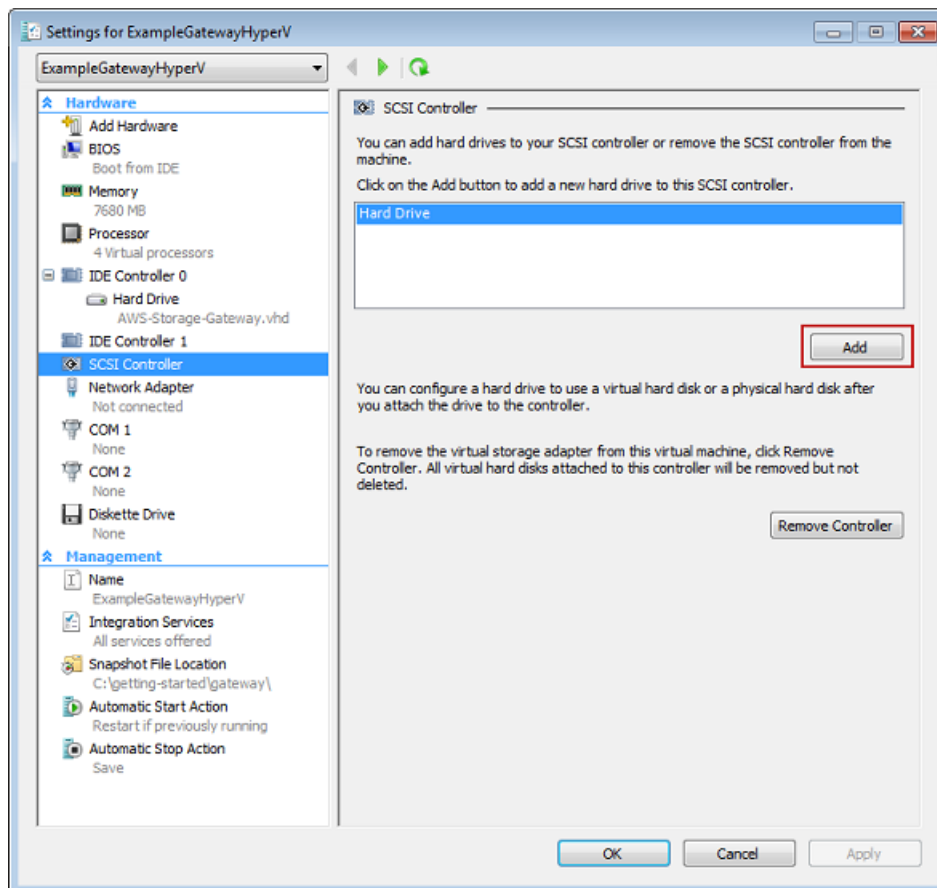
1. Start the Microsoft Hyper-V Manager and connect to the hypervisor.
2. In the **Virtual Machines** list pane, select the virtual machine **ExampleGatewayHyperV**.
3. In the **Actions** pane, select **Settings**....

AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway



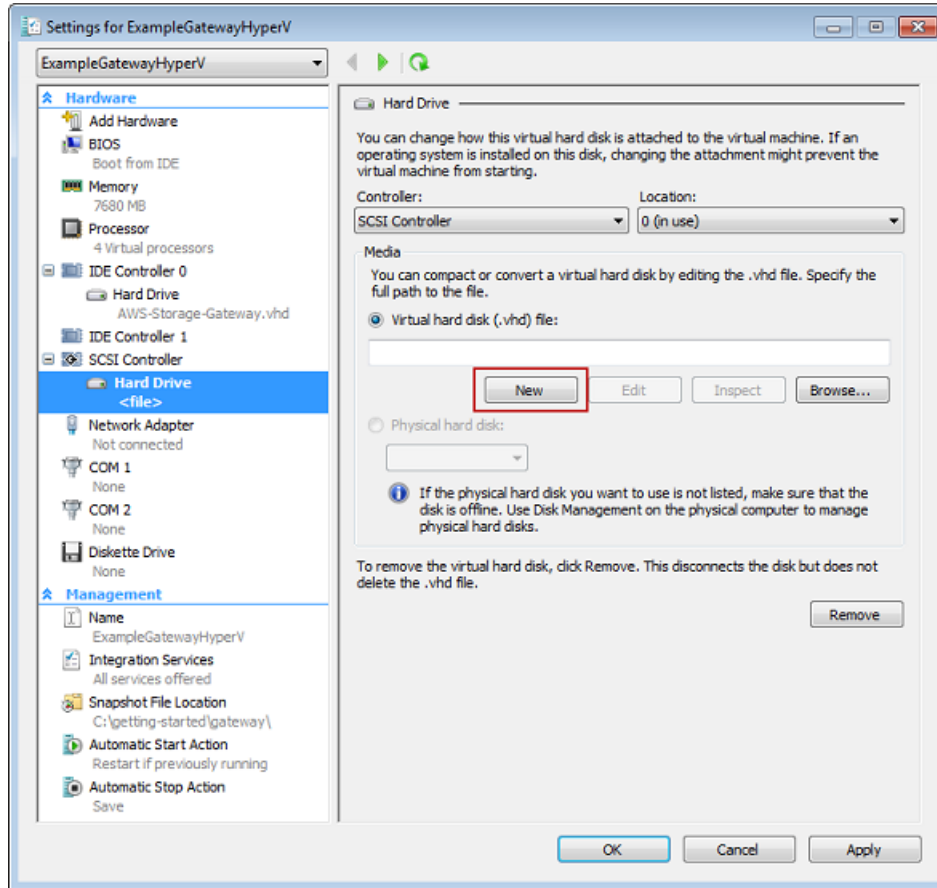
4. In the **Settings** window, select **SCSI Controller**, and click **Add**.




5. In the **Hard Drive** pane, under **Media**, click **New**.

AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway



6. In the **New Virtual Hard Disk Wizard** create a new virtual hard disk.
 - a. On the **Before You Begin** page, click **Next**.
 - b. On the **Choose Disk Type** page, choose **Fixed size**, and click **Next**.



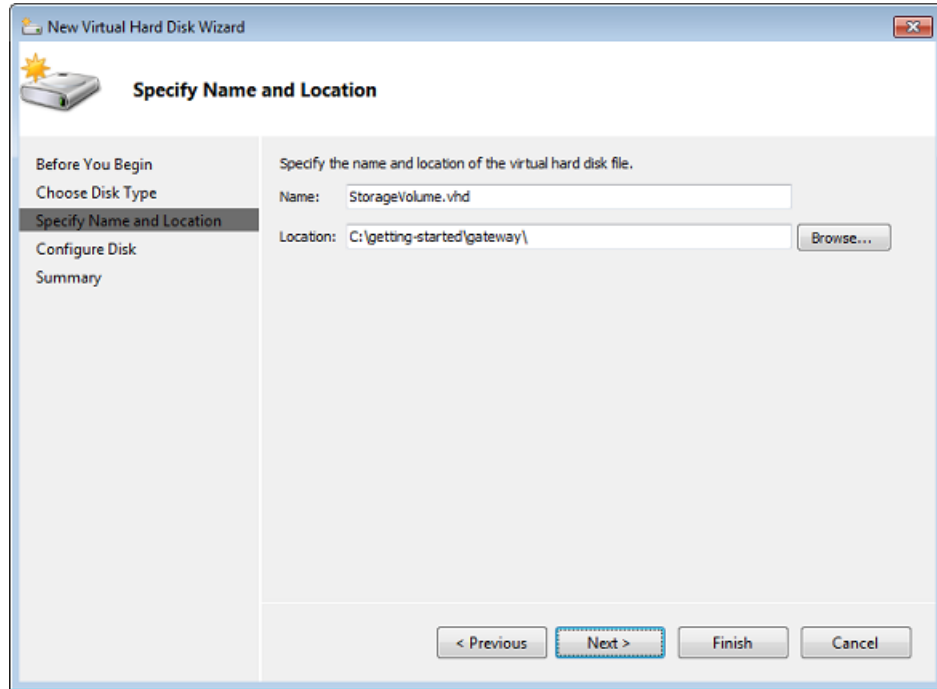
- c. On the **Specify Name and Location** page, specify a name and location for the virtual hard disk.
 - i. Specify `StorageVolume.vhd` as the name.
 - ii. Specify the location as `c:\getting-started\gateway`.

Note

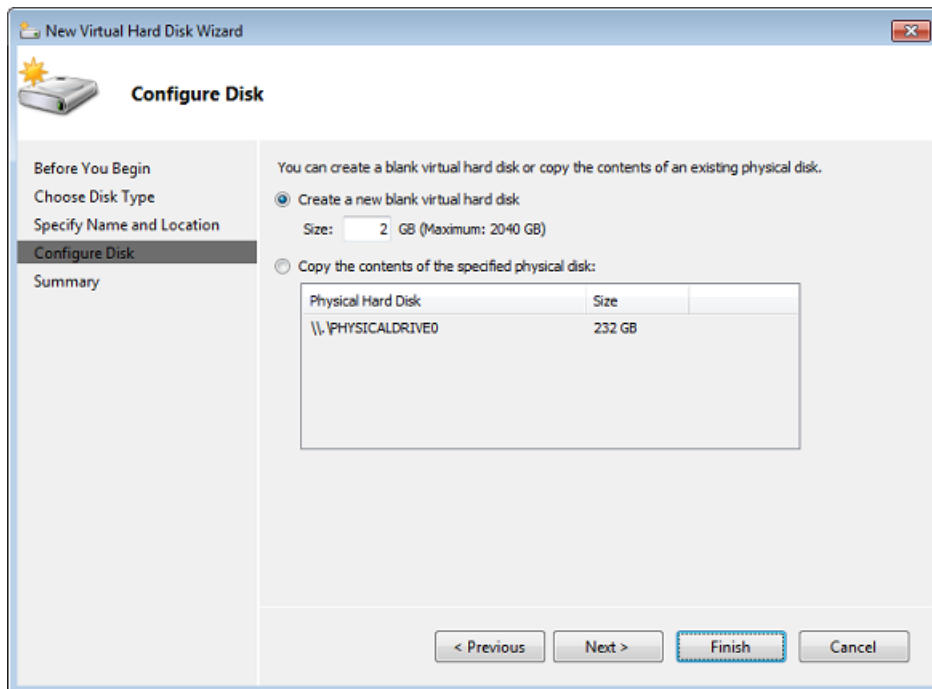
In this example setup, you store the virtual disk with the virtual machine. For real-world workloads, it is strongly recommended that you do not provision local disks using the same underlying physical storage disk. Depending on your hosting environment and performance and portability requirements, it might be better to select a different physical disk in this step. For more information, see [Provisioning Local Disks \(Gateway-Stored\)](#) (p. 102).

- iii. Click **Next**.

AWS Storage Gateway User Guide
Step 2.1: Set Up and Activate a Gateway



- d. In the **Configure Disk** page, specify the size of the disk as 2 GiB, and click **Finish**.



- e. After the virtual disk is created, verify the **Hard Drive** shows up under **SCSI controller**.
f. Click **SCSI Controller** to prepare to add another hard drive.

Allocate a Local Disk for an Upload Buffer

The gateway needs buffer space to temporarily store data as it uploads snapshots to AWS. This is referred to as the upload buffer. You must add virtual disks to the VM exclusively for use by the VM. The size of the upload buffer that the gateway needs depends on the size of the disks that you allocate for storing your data. For related guidelines, see [Sizing the Upload Buffer \(Gateway-Stored\) \(p. 106\)](#).

For this tutorial, you allocate a 10 GiB virtual disk to the VM for exclusive use by the gateway. In the **Create a Disk** pane of the wizard, enter 10 GiB for the disk size.

To allocate a local disk for the upload buffer

- Repeat the steps in the [To allocate a local disk to store your application data \(p. 50\)](#) procedure to add another virtual disk to the gateway. Follow the steps exactly except use the name `UploadBuffer.vhd` for the disk name and 10 GiB for the disk size.

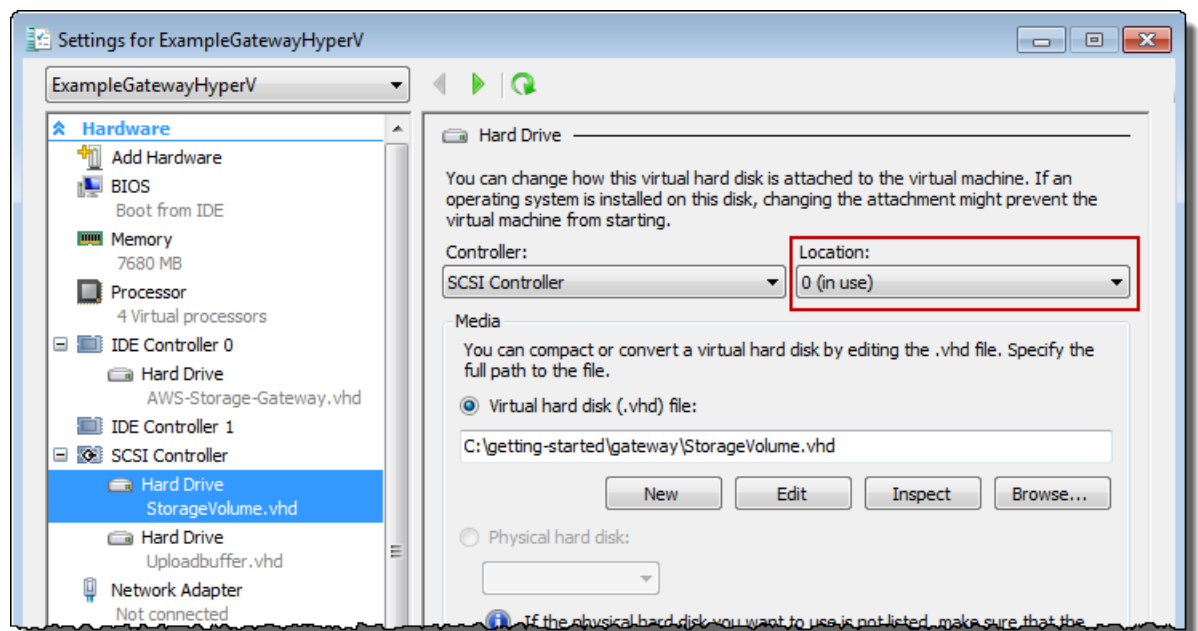
Verify the Gateway VM Has Two Disks

The remainder of this tutorial requires that you have allocated two disks to your gateway VM; use the following optional procedure to verify this. If you need to allocate another disk, repeat the steps in the [To allocate a local disk to store your application data \(p. 50\)](#) procedure.

To verify the VM has two disks

- Start the Microsoft Hyper-V Manager and connect to the hypervisor.
- In the **Virtual Machines** list pane, select the virtual machine **ExampleGatewayHyperV**.
- In the **Actions** pane, select **Settings...**
- In the **Settings** window, select **SCSI Controller**, and verify that there are two disks.

The two disks you created will be used later in the AWS Storage Gateway console and appear as SCSI (0:0) and SCSI (0:1) in drop-down lists. In the example below, the **StorageVolume.vhd** disk is selected and is SCSI (0:0).




Activate Your Gateway

Now, you are ready to activate your gateway. The activation process associates your gateway with your AWS account. You must power on the gateway VM before you activate your gateway.

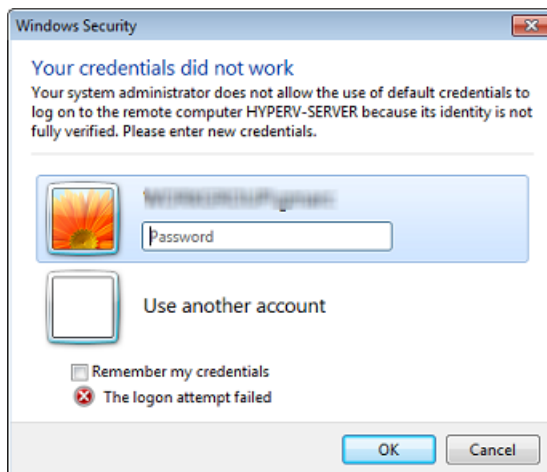
To activate your gateway

1. Power on the VM.
 - a. Start the Microsoft Hyper-V Manager and connect to the hypervisor.
 - b. In the **Virtual Machines** list pane, select the virtual machine **ExampleGatewayHyperV**.
 - c. In the **Actions** pane, select **Start**.

The **Virtual Machine Connection** window appears.

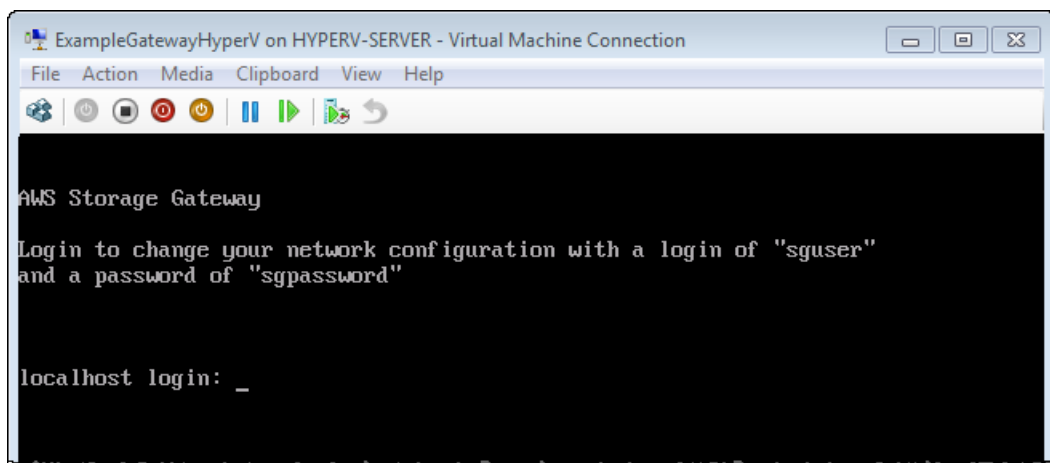


- d. If an authentication window appears, enter the user name and password provided to you by the hypervisor administrator.

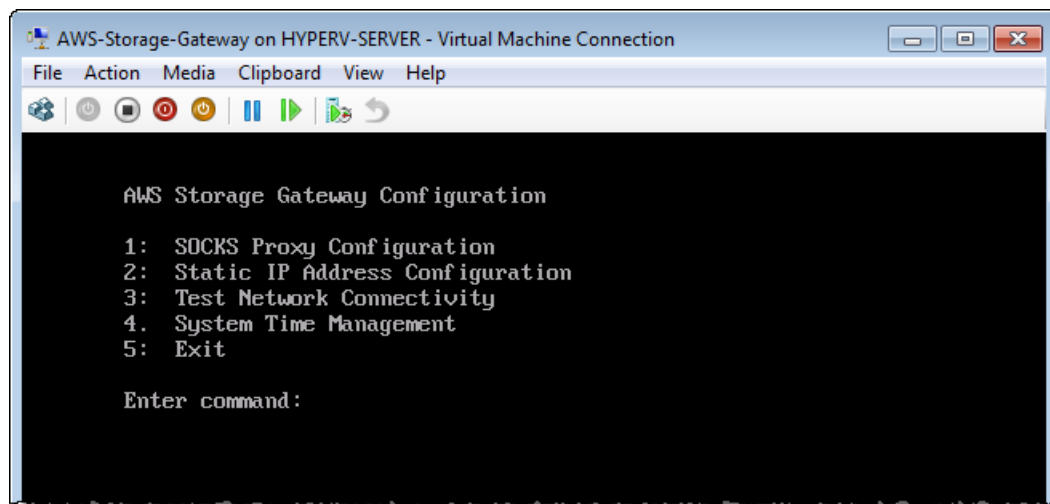


- e. After a few moments, the virtual machine is ready for you to log in.

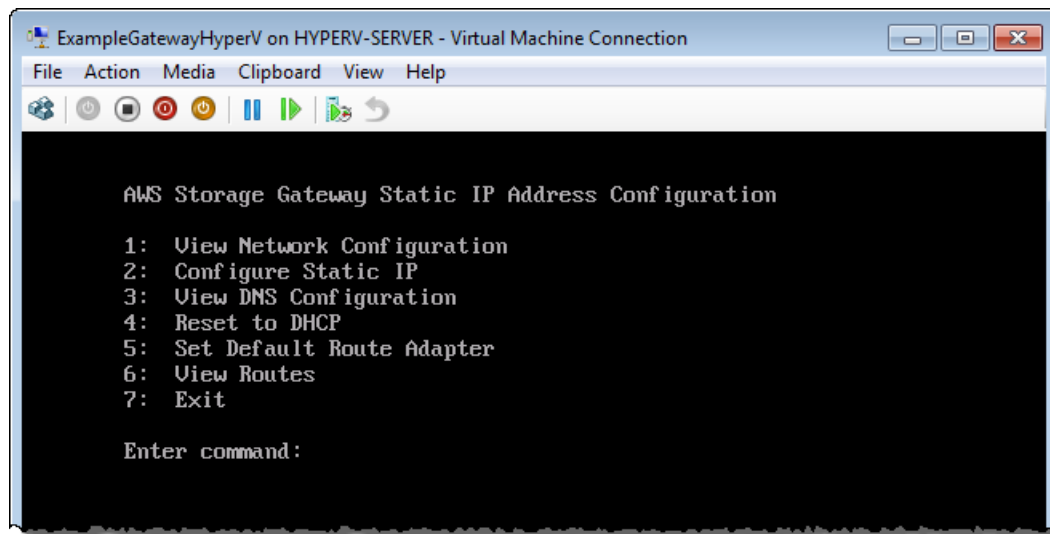
The example below shows the login prompt you see when the VM is ready.



2. Activate your gateway.
 - a. Obtain the IP address of your gateway.
 - i. In the Microsoft Hyper-V Manager, select the deployed gateway VM.
 - ii. In the **Virtual Machines** list pane, select the virtual machine **ExampleGatewayHyperV**.
 - iii. In the **Actions** pane, select **Connect...**The **Virtual Machine Connection** window appears.
 - iv. At the login prompt, enter the user name `sguser`, and the password `sgpassword`.
 - v. In the **AWS Storage Gateway Configuration** menu, select option 2, **Static IP Address Configuration**.

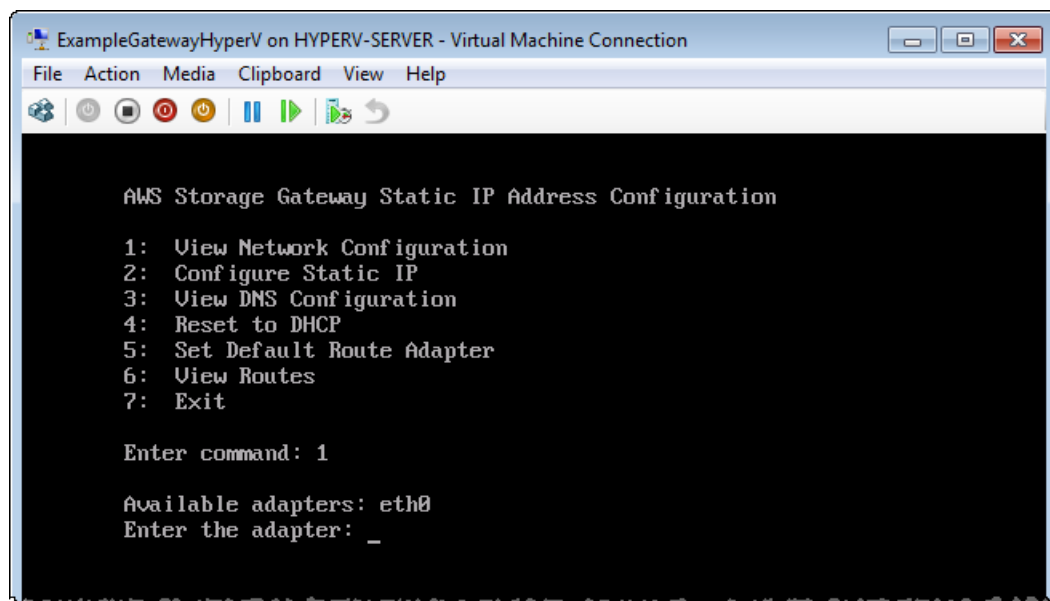


- vi. In the **AWS Storage Gateway Static IP Address Configuration** menu, select option 1, **View Network Configuration**.



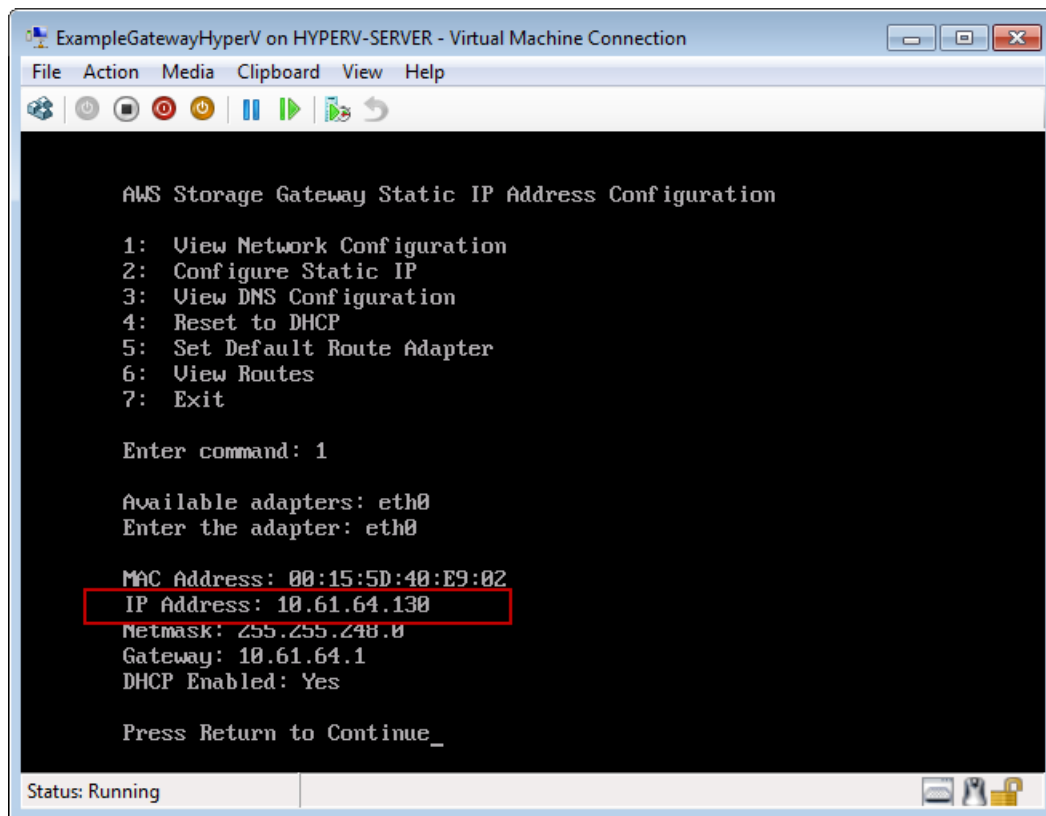
vii. Type the identifier of the adapter.

In most setups, `eth0` will be the adapter identifier.



viii. Get the IP address from the adapter information.

In the example below, the IP address is 10.61.64.130. Your gateway's IP address will be different.



- ix. Press **Return**, and follow the prompts to exit the configuration menu.
- b. Associate your gateway to your AWS account.
 - i. In the AWS Storage Gateway console, in the **Setup and Activate Gateway** wizard, navigate to the following **ACTIVATE GATEWAY** page.
 - A. If the wizard is not already started, click the **Set up and Activate a New Gateway** button.
 - B. Click **Continue** in each wizard step until you reach the **ACTIVATE GATEWAY** page.
 - ii. Enter the IP address of your gateway, and click **Proceed to Activation**.

AWS Storage Gateway User Guide

Step 2.1: Set Up and Activate a Gateway

Setup and Activate Gateway close

PROVISION HOST DOWNLOAD AND DEPLOY VM PROVISION LOCAL DISK STORAGE **ACTIVATE GATEWAY**

Using your Hyper-V Manager client, double-click on your imported gateway VM and select "Start". Follow the instructions that are subsequently shown to log in.

Select "2: Static IP Address Configuration". In the next screen, select "1: View Network Configuration". Enter the available network adapter. This will show you the IP address, which you type into the box below.

Clicking "Proceed to Activation" will redirect you to the activation page (your browser must be running on a machine with network connectivity to your local gateway host).

[Step-by-Step Instructions](#)

Enter IP Address Below:

[← Back](#) Proceed to Activation

```
MAC Address: 00:15:5D:63:00:10
IP Address: 192.168.99.221
Netmask: 255.255.255.0
Gateway: 192.168.99.1
DHCP Enabled: Yes
Press Return to Continue_
```

Screenshot of the IP address look-up. This action is taken outside of the AWS Management Console.

Note

During activation, your browser connects to the gateway. If activation fails, then check that the IP address you entered is correct. If the IP address is correct, then confirm that your network is configured to allow your browser to access the gateway VM.

- iii. On the activation page fill in the requested information to complete activation.

The **AWS Region** determines where AWS stores your snapshots. If you choose to restore a snapshot to an Amazon EBS volume, then the Amazon EBS volume must be in the same region as the snapshot. You cannot change the region after the gateway is activated.

The gateway name identifies your gateway in the console. You use this name to manage your gateway in the console, and you can change it post-activation. This name must be unique to your account.

AWS Storage Gateway

Activating Your AWS Storage Gateway Virtual Machine (VM)

Below is the type and IP address of the gateway you are activating:

Gateway Type: Gateway-Cached Volumes

Activated gateways are billed at \$125 per month, prorated daily. Upon activation of your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies solely to the gateway price. Storage pricing and data transfer pricing continue to apply. The AWS Service Terms are available [here](#).

Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway.

AWS Region: US East (Virginia)

Gateway Time Zone: (GMT -8:00) Pacific Time (US & Canada)

Gateway Name: MyNewGateway

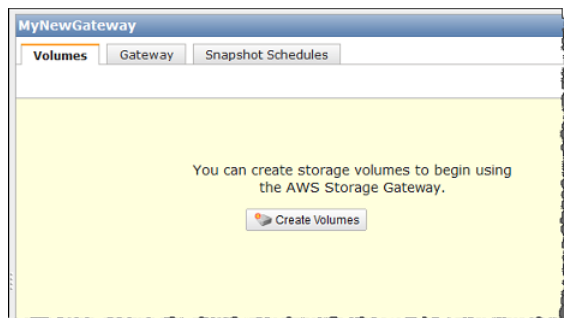
Activate My Storage Gateway

[Click here](#) if you need to exit the activation process.

- iv. Click **Activate My Storage Gateway**.

Upon successful activation, the **AWS Storage Gateway** console displays a link to the activated gateway under the **Gateways** section of the **Navigation** pane. Click the gateway you just added.

The **Create Volumes** button is displayed.



Step 2.2: Create Volumes Using the AWS Storage Gateway Console

So far you have deployed the gateway VM, allocated disks, and activated the gateway. Now you are ready to create iSCSI storage volumes. For this setup, you use the console to create these volumes.

- If you are testing setup for the gateway-cached architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)), you allocate the two disks that you previously added to the VM, one disk for cache storage and one for the upload buffer. You then create an iSCSI storage volume in Amazon S3. Data from your on-premises applications is written to this volume, which is stored in Amazon S3. The gateway maintains the volume's recently accessed data locally in the cache storage.
- If you are testing the setup for the gateway-stored architecture, you create an iSCSI storage volume mapped to one of the two disks you previously added to the VM. You allocate the remaining disk to your gateway's upload buffer. Data from your on-premises applications is written to this volume, which is stored locally. The gateway periodically takes snapshots (incremental backups) and uploads them to Amazon S3.

Click one of the following links and follow instructions to create the volumes.

To...	Do This...
Create volumes (Gateway-Cached architecture)	Follow the steps in Create Volumes (Gateway-Cached) (p. 61) .
Create volumes (Gateway-Stored architecture)	Follow the steps in Create Volumes (Gateway-Stored) (p. 67) .

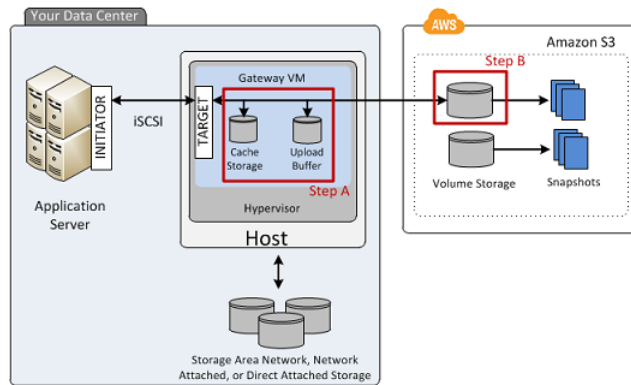
Create Volumes (Gateway-Cached)

Topics

- [Step A: Create Cache Storage and an Upload Buffer on Your Local Disks and Configure Optional Alarms \(p. 62\)](#)
- [Step B: Create a Volume in Amazon S3 \(p. 65\)](#)

In the gateway-cached setup (see [How AWS Storage Gateway Works \(p. 3\)](#)), you allocate the two disks, that you previously added to the VM, for cache storage and upload buffer. You then create an iSCSI storage volume in Amazon S3. Data that your on-premises applications write to this volume is stored in Amazon S3. The gateway maintains the volume's recently accessed data locally in the Cache Storage.

The following architectural overview diagram shows what part of the gateway-cached setup you are creating.



Step A: Create Cache Storage and an Upload Buffer on Your Local Disks and Configure Optional Alarms

In this step, you create the cache storage and upload buffer on your local disks and configure alarms.

To create local volumes (cache storage and upload buffer)

1. In the **Navigation** pane of the console, select your gateway, click the **Volumes** tab, and then click **Create Volumes**.



2. In the **Configure Your Activated Gateway** wizard, configure local working storage, which consists of the upload buffer and cache storage.



- Of the two disks you provisioned, select one as the upload buffer and the other as cache storage, as shown in the following example.



- Click **Next**.

Now you can create optional alarms to monitor the storage utilization of the two volumes you created.

To configure optional alarms

- In the upload buffer alarm dialog box, configure the alarm for upload buffer utilization.

Configure Your Activated Gateway

CONFIGURE LOCAL STORAGE CREATE VOLUME

Use Basic Alarms

Notify me when my gateway's Upload Buffer utilization exceeds 65% or 85% of the space available.

Notification will be sent when:

- Upload Buffer used > ~7 GiBs
- or
- Upload Buffer used > ~9 GiBs

Send notification to (email):
user@example.com

Additional charges may apply if you exceed the AWS Free Tier levels for CloudWatch or Simple Notification Service.

Advanced alarm settings are available in the CloudWatch Management Console.

Skip Continue

- a. Using the two drop-down boxes, select utilization percentages that are used to create two upload buffer alarms.

You can select the thresholds, for example, so that the first threshold (the lower percentage value) represents a upload buffer percentage utilization that, if exceeded, you want to be warned about. The second threshold can be selected to represent a upload buffer utilization that, if exceeded, is cause for action, such as adding more upload buffer space.

After you complete this step, you can go to the Amazon CloudWatch console at any time and change the alarm thresholds.

- b. Enter an email address.
- c. Click **Continue**.

Two alarms are created. For example, using the gateway name *MyNewGatewayCached*, the alarms created are *MyNewGatewayCached-UploadBufferUtilization-Alarm1* and *MyNewGatewayCached-UploadBufferUtilization-Alarm2*.

- d. Check for a subscription confirmation email that is sent to the email address you indicated, and follow the instructions in that email to confirm your subscription to the Amazon Simple Notification Service (Amazon SNS) topic. After you have confirmed your subscription, you will receive an email when either threshold you specified is exceeded.

2. Click **Continue** to create the alarm.
3. In the cached storage alarm dialog box, configure alarm for cache storage utilization.

Configure Your Activated Gateway

CONFIGURE LOCAL STORAGE CREATE VOLUME

Use Basic Alarms

Notify me when the dirty data in my gateway's cache storage exceeds 65% or 85% of the space available.

Notification will be sent when:

- Dirty data in Cache Storage used > ~13 GiBs
- or
- Dirty data in Cache Storage used > ~17 GiBs

Send notification to (email):
user@example.com

Additional charges may apply if you exceed the AWS Free Tier levels for CloudWatch or Simple Notification Service.

Advanced alarm settings are available in the CloudWatch Management Console.

Skip Continue

Follow instructions in the preceding step to configure this alarm.

For example, using the gateway name *MyNewGatewayCached*, the alarms created are *MyNewGatewayCached-CacheUtilization-Alarm1* and *MyNewGatewayCached-CacheUtilization-Alarm2*.

4. Click **Continue** to create the alarm.

Continue in the next section to create the volume storage in Amazon S3.

Step B: Create a Volume in Amazon S3

In the preceding section, you created volumes (cache storage and upload buffer) on your on-premises hardware. Now, you create volume storage in Amazon S3. This is where your application data will reside.

To create a cached volume in Amazon S3

1. In the **Configure Your Activated Gateway** dialog box, create an iSCSI storage volume in Amazon S3.

Configure Your Activated Gateway close

CONFIGURE LOCAL STORAGE CREATE VOLUME

Create an iSCSI storage volume up to 32 TBs in size. This volume will be stored in Amazon S3, with only a cache of recently accessed data kept locally. Your client applications will connect to this volume over an iSCSI interface. [Learn More](#).

Capacity: 50 TBs (Max: 32 TBs)
TBs
GBs

iSCSI Target Name: iqn.1997-05.com.amazon:myvolume

Based on Snapshot ID:

Host IP: 192.168.99.227

Port: 3260

Cancel Create Volume

- a. For **Capacity**, specify 50 GiB.

The maximum size you can specify is 32 TiB.

- b. Enter a name in the **iSCSI Target Name** field.

The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name in the **Targets** tab of the **iSCSI Microsoft Initiator** UI after discovery. For example, a name `target1` would appear as `iqn.1007-05.com.amazon:target1`. Ensure that the target name is globally unique within your SAN network.

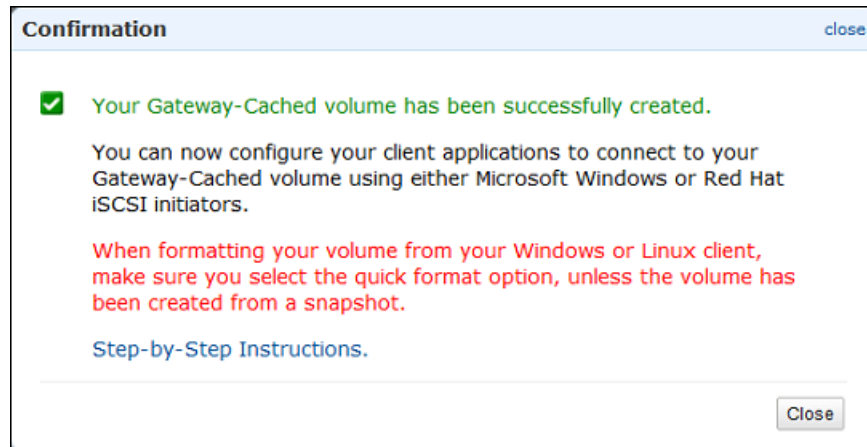
For this tutorial, use **myvolume**.

- c. Leave the **Based on Snapshot ID** field empty.

If you want to restore an existing Amazon EBS snapshot or a gateway snapshot on to the storage volume that you are creating, you must specify the snapshot ID. The gateway downloads your existing snapshot data to the storage volume.

- d. Verify that the **Host IP** field is the IP address of your gateway, and click **Create Cached Volume**.

2. In the **Configure iSCSI Initiators** dialog box, click **Close**.



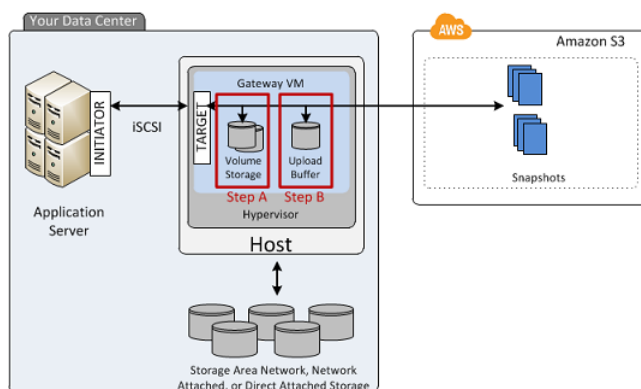
Create Volumes (Gateway-Stored)

Topics

- [Step A: Create a Storage Volume \(p. 67\)](#)
- [Step B: Create an Upload Buffer \(p. 69\)](#)

In the gateway-stored setup (see [How AWS Storage Gateway Works \(p. 3\)](#)), you create an iSCSI storage volume mapped to one of the two disks you previously added to the VM. You allocate the remaining disk to your gateway's upload buffer. Data that your on-premise applications write to this volume is stored locally. The gateway periodically takes snapshots (incremental backups) and uploads them to Amazon S3.

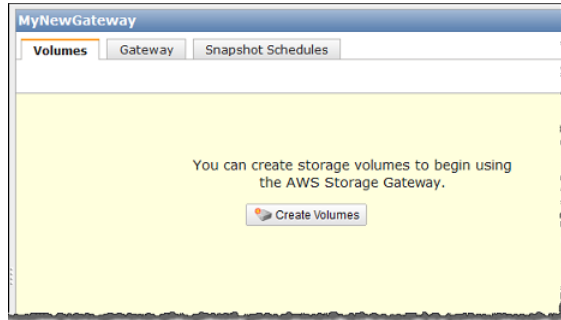
The following architectural overview diagram shows what part of the gateway-stored setup you are creating.



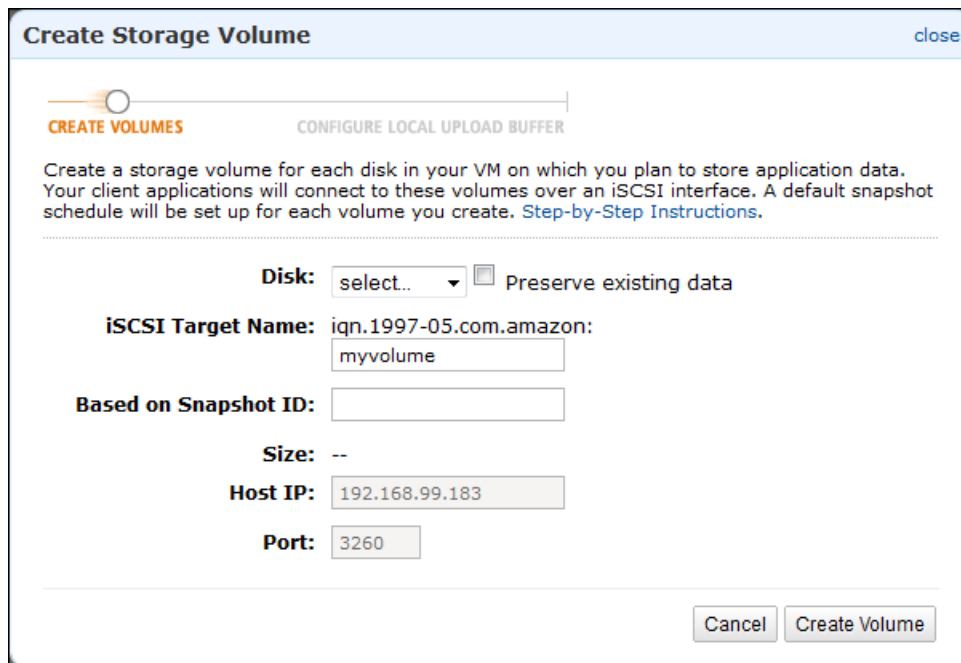
Step A: Create a Storage Volume

To create a storage volume

1. In the **Navigation** pane of the console, select your gateway, click the **Volumes** tab, and click **Create Volumes**.



2. In the **Create Storage Volume** wizard, provide storage volume information.



- a. In the drop-down list of the **Disk** field, select the 2 GiB virtual disk on your VM.

This drop-down list shows the virtual disks that you added to the gateway VM. Select the disk on which you plan to store data.

- b. Keep the **Preserve existing data** check box unchecked.

Caution

Make sure that you don't have any existing data on the virtual disks. Any existing data on the disk is lost.

- c. Enter a name in the **iSCSI Target Name** field.

The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name in the **Targets** tab of the **iSCSI Microsoft Initiator** UI after discovery. For example, a name `target1` would appear as `iqn.1007-05.com.amazon:target1`. Ensure that the target name is globally unique within your SAN.

- d. Leave the **Based on Snapshot ID** field empty.

If you want to restore an existing Amazon EBS snapshot or a gateway snapshot on to the storage volume that you are creating, you must specify the snapshot ID. The gateway downloads your existing snapshot data to the storage volume.

- e. Verify that the **Host IP** setting is the IP address of your gateway, and click **Create Volume**.

Create Storage Volume close

CREATE VOLUMES CONFIGURE LOCAL UPLOAD BUFFER

Create a storage volume for each disk in your VM on which you plan to store application data. Your client applications will connect to these volumes over an iSCSI interface. A default snapshot schedule will be set up for each volume you create. [Step-by-Step Instructions](#).

Disk: SCSI (0:0) Preserve existing data

iSCSI Target Name: iqn.1997-05.com.amazon:
myvolume

Based on Snapshot ID:

Size: 2 GiB

Host IP: 192.168.99.183

Port: 3260

Cancel Create Volume

Step B: Create an Upload Buffer

AWS Storage Gateway requires storage space to buffer your incoming application data before uploading it to AWS. In this step, you configure one virtual disk as an upload buffer.

Note

When you configure a disk as an upload buffer, you lose any existing data on the disk, so be careful to preserve your data.

To allocate the upload buffer for your AWS Storage Gateway VM

1. In the **Create Storage Volume** dialog box, click **Configure Local Upload Buffer**.

Create Storage Volume close

CREATE VOLUMES CONFIGURE LOCAL UPLOAD BUFFER

Create a storage volume for each disk in your VM on which you plan to store application data. Your client applications will connect to these volumes over an iSCSI interface. A default snapshot schedule will be set up for each volume you create. [Step-by-Step Instructions](#).

Disk: select.. Preserve existing data

iSCSI Target Name: iqn.1997-05.com.amazon:
myvolume

Based on Snapshot ID:

Size: 2 GiB

Host IP: 192.168.99.183

Port: 3260

1 Volume Created

Cancel Create Volume

To start using the volumes you have created, proceed to the next step of allocating local Upload Buffer for your gateway.

Configure Local Upload Buffer

2. Select the check box next to the remaining disk to allocate the disk as the upload buffer, and then click **Next**.

This dialog box lists all available disks on your VM. Earlier, you added two virtual disks to the VM and configured one of the disks as a storage volume. Therefore, the dialog box should show one available disk. Select the disk to allocate as the upload buffer. You can extend the upload buffer later without disrupting the iSCSI I/O.

Configure Local Upload Buffer close

CREATE VOLUMES CONFIGURE LOCAL UPLOAD BUFFER

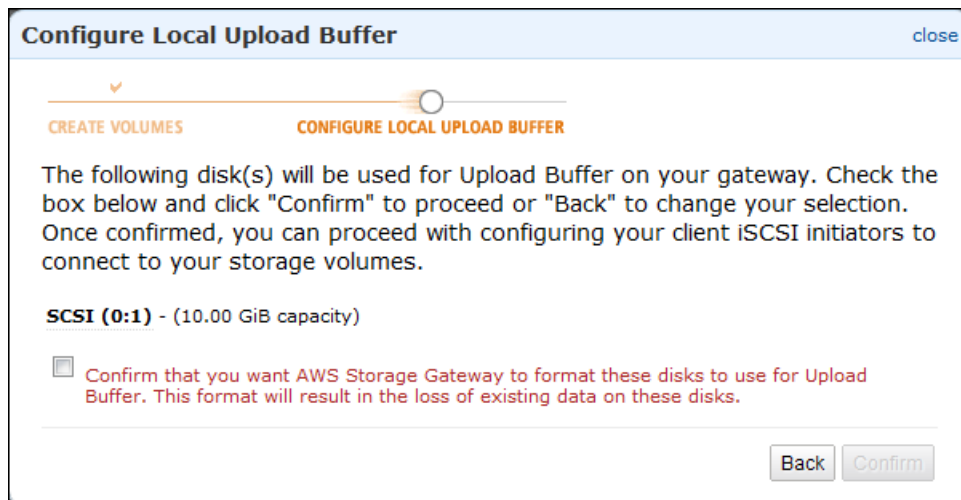
Please select which disks the VM can use for Upload Buffer. Please see our documentation for recommendations on how much space to provide given your workload and network connection. [Step-by-Step Instructions](#)

Local Disks

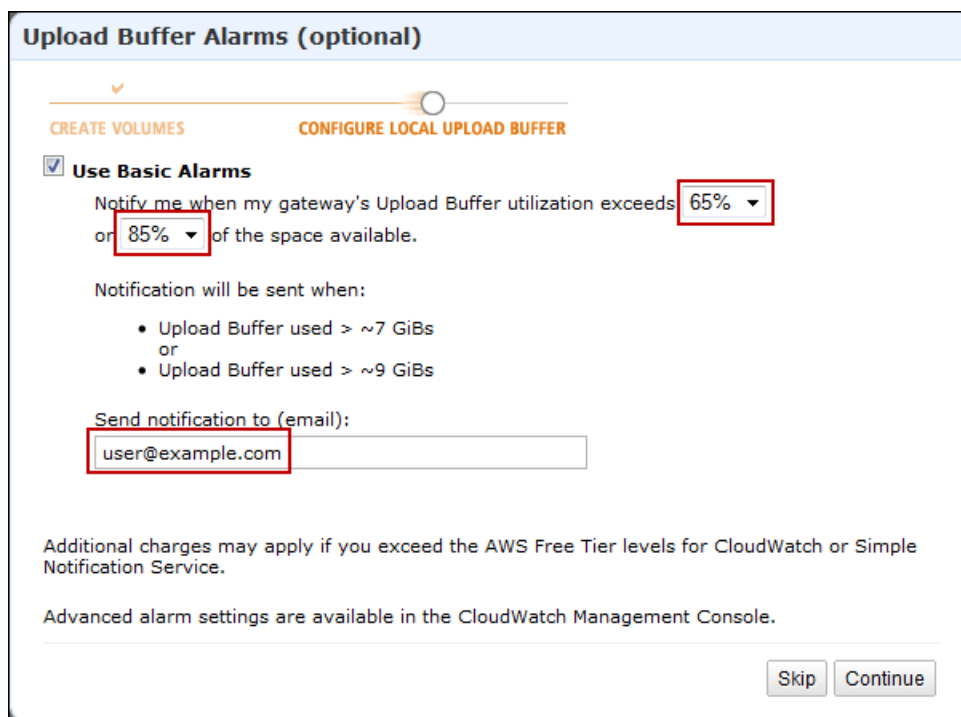
SCSI (0:0)	2.00 GiB	<input type="checkbox"/>	In Use for Storage Volume
SCSI (0:1)	10.00 GiB	<input checked="" type="checkbox"/>	Use for Upload Buffer Space

Cancel Next

3. In the confirmation dialog box, select the check box and click **Confirm**.



4. In the **Upload Buffer Alarms** dialog box, configure alarms for your upload buffer.



- a. Using the two drop-down boxes, select utilization percentages that are used to create two upload buffer alarms.

You can select the thresholds, for example, so that the first threshold (the lower percentage value) represents a upload buffer percentage utilization that, if exceeded, you want to be warned about. The second threshold can be selected to represent a upload buffer utilization that, if exceeded, is cause for action, such as adding more upload buffer capacity.

After you complete this step, you can go to the Amazon CloudWatch console at any time and change the alarm thresholds.

- b. Enter an email address.
- c. Click **Continue**.

Two alarms are created. Using the gateway name in this tutorial as an example, the alarms would be named *MyNewGateway-UploadBufferUtilization-Alarm1* and *MyNewGateway-UploadBufferUtilization-Alarm2*.

- d. Check for a subscription confirmation email that is sent to the email address you indicated and follow the instructions in that email to confirm your subscription to the Amazon Simple Notification Service (Amazon SNS) topic. After you have confirmed your subscription, you will receive an email when either threshold you specified is exceeded.

For more detailed information about creating upload buffer alarms, see [Monitoring the Upload Buffer](#) (p. 267).

5. In the **Configure iSCSI Initiators** dialog box, click **Close**.



Step 2.3: Access Your AWS Storage Gateway Volumes

You are now ready to connect your Windows client to your iSCSI storage volume. In this Getting Started exercise, you make this connection using the Microsoft iSCSI Initiator on your client. For instructions on accessing the iSCSI storage volume from Linux, see [Connecting from a Red Hat Client to Your Storage Volume](#) (p. 165).

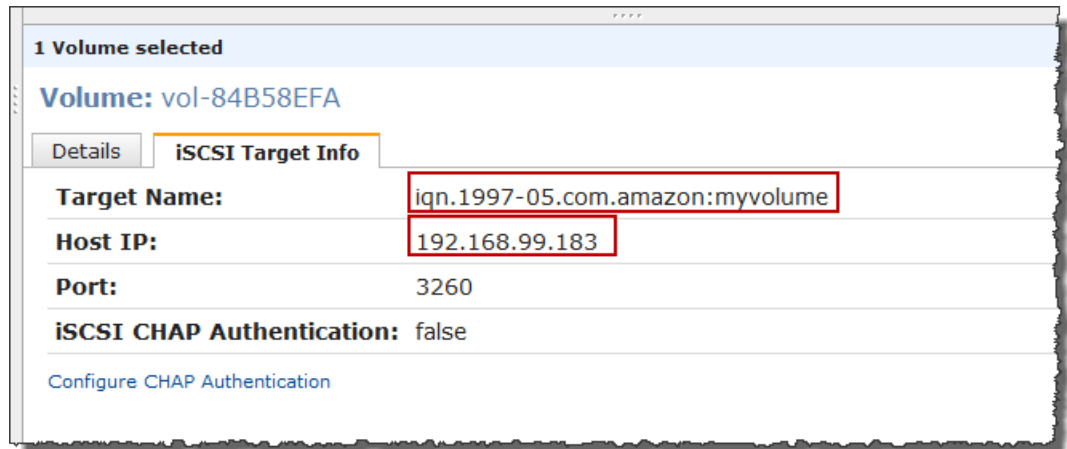
Note

You must have administrator rights to run the iSCSI Initiator.

To connect your Windows client to the storage volume

1. You need the host IP and the target name information for the storage volume you are connecting to. You can find this information in the AWS Storage Gateway console.
 - a. In the **Navigation** pane of the console, select your gateway.
 - b. In the **Volumes** tab, click the volume to connect to.

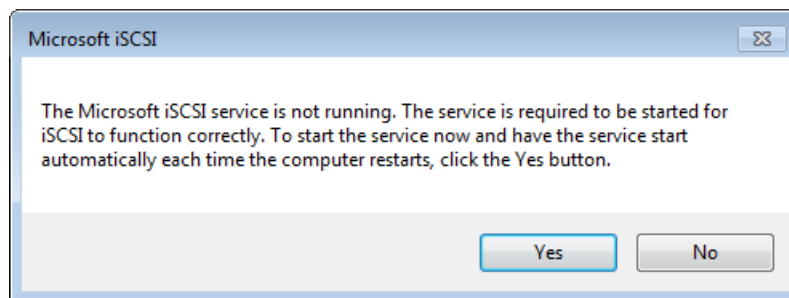
The **iSCSI Target Info** tab shows the information you need to connect your client to this volume.



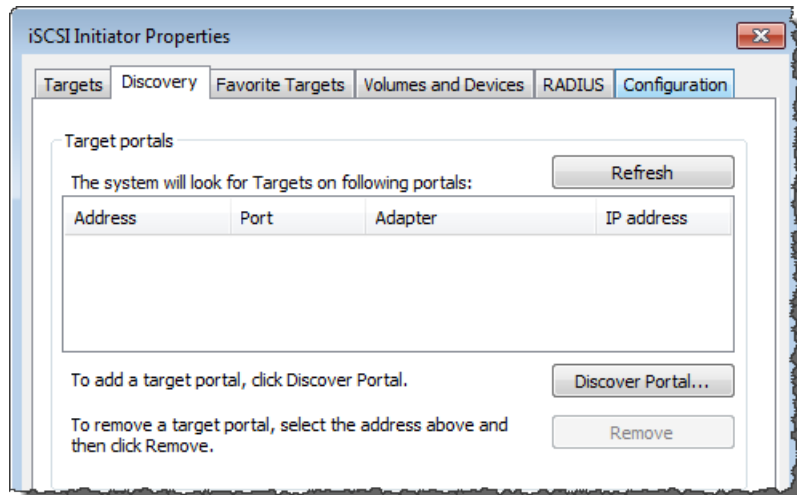
You will use the **Target Name** and **Host IP**, as highlighted in the preceding image. The **Host IP** is required for the following steps and the **Target Name** is used to verify that you are selecting the correct iSCSI target.

2. Start the iSCSI Initiator.
 - a. In the **Start** menu of your Windows client computer, type `iscsicpl.exe` and run the program.

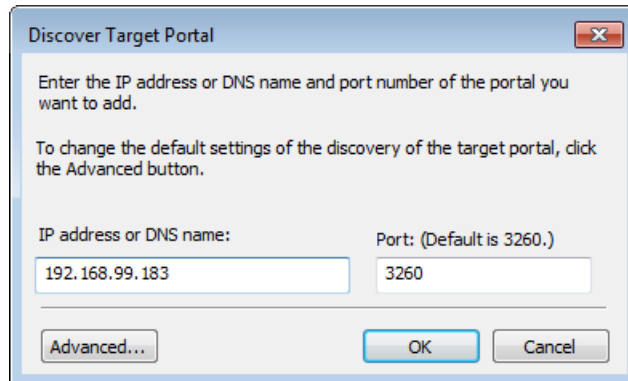
The **iSCSI Initiator Properties** dialog box appears if the iSCSI Initiator Service is running.
 - b. If the Microsoft iSCSI Initiator Service is not running, you are prompted to start the service and have the service start automatically each time the computer restarts. Click **Yes** in the **Microsoft iSCSI** dialog box to start the service.



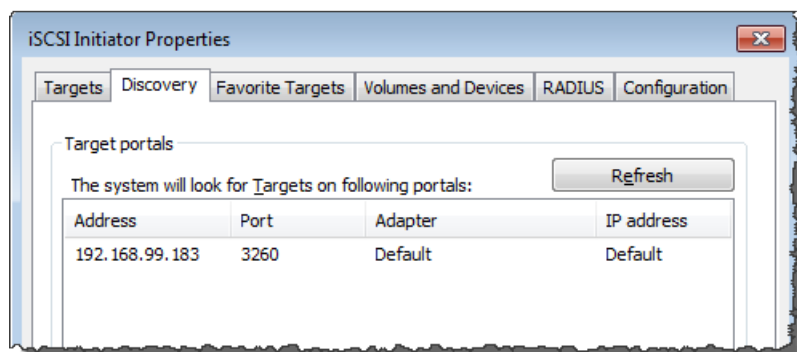
3. Discover the gateway:
 - a. In the **iSCSI Initiator Properties** dialog box, click the **Discovery** tab, and click the **Discovery Portal** button.



- b. In the **Discover Target Portal** dialog box, in the **IP address or DNS name** field, enter the IP address of your iSCSI target, and click **OK**.



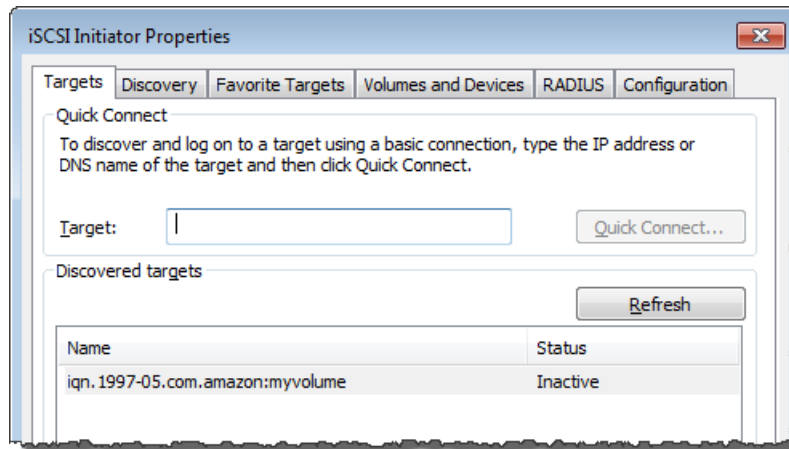
- c. The IP address is now displayed in the list of **Target portals** in the **Discovery** tab.



4. Connect to the storage volume target on the gateway:


- a. Click the **Targets** tab.

The target you just discovered is shown with an inactive status. Note that the target name shown should be the same as what you noted for your storage volume in step 1.

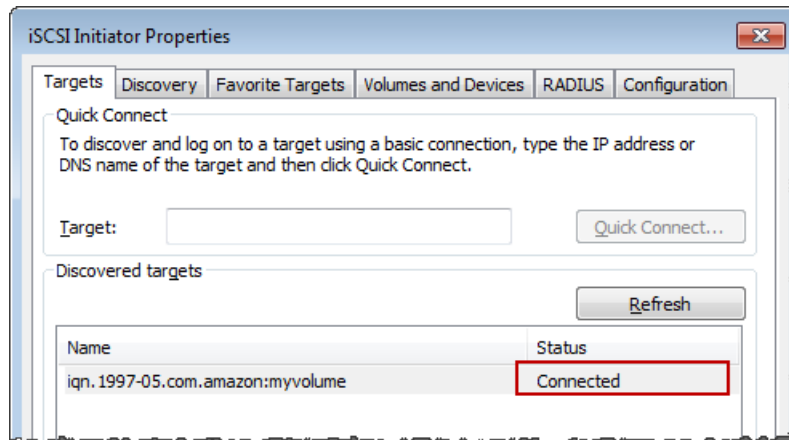


- b. Select the target and click **Connect**.

In the **Connect to Target** dialog box, if the target name is not prepopulated already, enter the name of the target name as shown in step 1, select the check box next to **Add this connection to the list of Favorite Targets**, and click **OK**.



- c. In the **Targets** tab, ensure that the target **Status** has the value **Connected** indicating the target is connected. Click **OK**.



You can now initialize and format this storage volume for Windows so you can begin saving data on it. You do this through the Windows Disk Management tool.

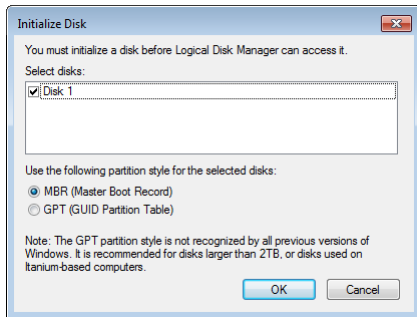
Note

While it is not required for this Getting Started exercise, we highly recommend that you customize your iSCSI settings for a real application as discussed in the topic [Customizing Your Windows iSCSI Settings \(p. 163\)](#).

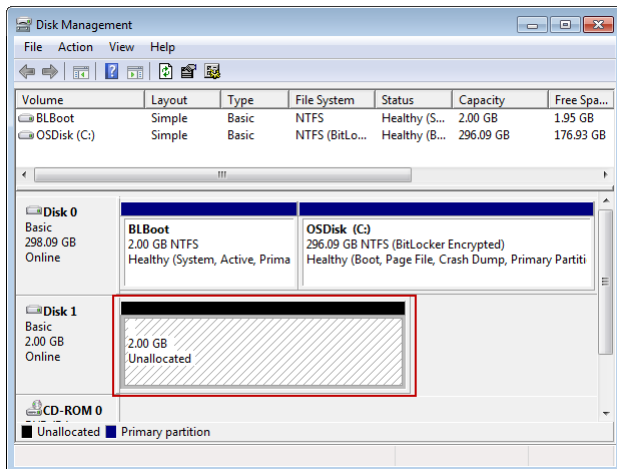
To initialize and format the storage volume you just mapped

1. In the **Start** menu, type `diskmgmt.msc` to open the **Disk Management** console.
2. In the **Initialize Disk** dialog box, select **MBR (Master Boot Record)** as the partition style and click **OK**. When selecting the partition style, you should take into account the type of volume you are connecting to—cached or stored—as shown in the following table.


Partition style	Use in the following conditions
MBR (Master Boot Record)	For all stored-volumes (which are limited to 1 TiB in size), or cached-volumes less than 2 TiB.
GPT (GUID Partition Table)	All stored-volume and cached-volumes. You must use GPT for cached-volumes greater than 2 TiB.



3. Create a simple volume:
 - a. If the disk is offline, you must bring it online before you can initialize it. After the disk is initialized, it is ready to be formatted as a simple volume. All the available volumes are displayed in the disk management console. In the following example, **Disk 1** is the storage volume. Notice that when you select the new volume, it displays hatch lines indicating that it is selected.

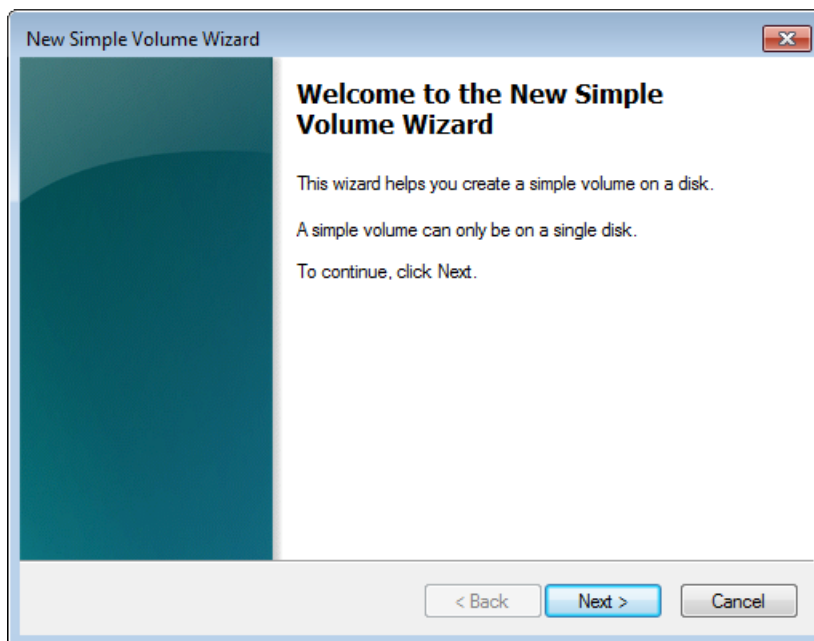


- b. Right-click the disk and select **New Simple Volume**.

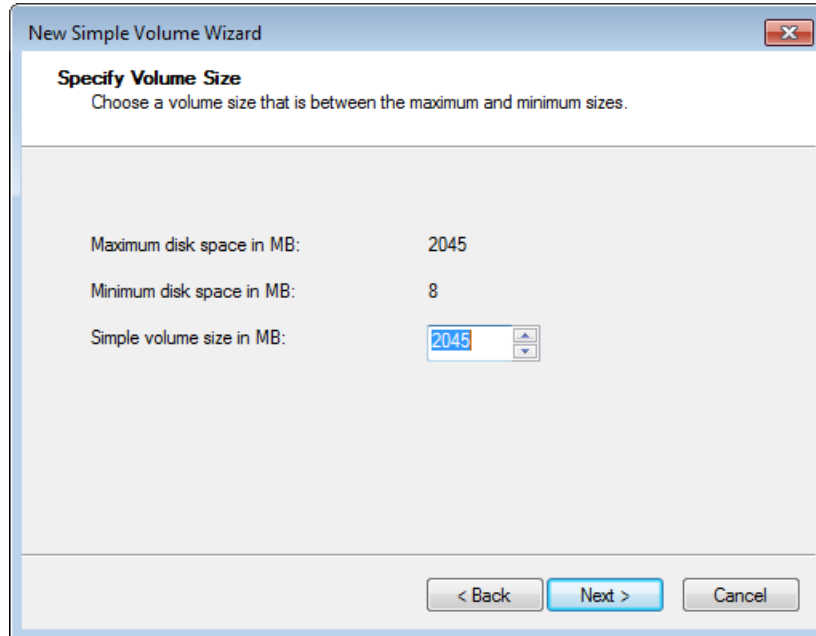


Important


Be careful not to format the wrong disk. Check to make sure that the disk you are formatting matches the size of the local disk you allocated to the gateway VM and that it has a status of **Unallocated**.



- c. In the **New Simple Volume Wizard**, click **Next**.
d. In the **Specify Volume Size** dialog box, leave the default values, and click **Next**.



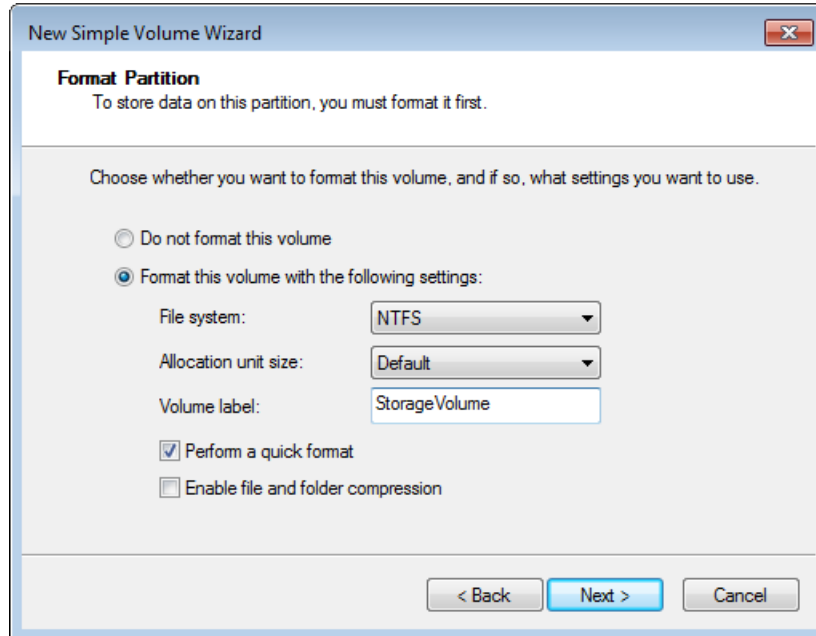
- e. In the **Assign Drive Letter or Path** dialog box, leave the default values, and click **Next**.



- f. In the **Format Partition** dialog box, specify a **Volume label** field, and ensure that **Perform a quick format** is selected. Click **Next**.

Caution

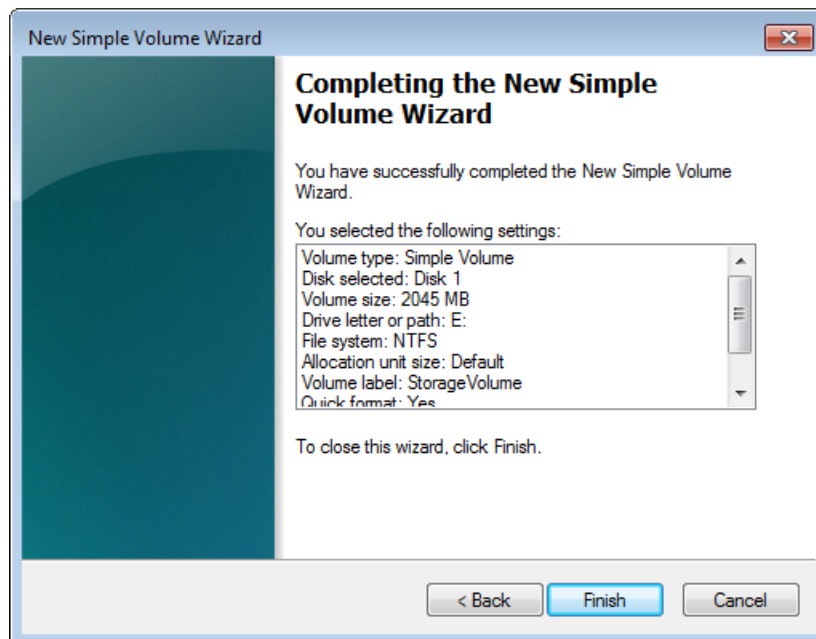
Selecting **Perform a quick format** is highly recommended for cached-volumes as it results in less initialization I/O, smaller initial snapshot size, fastest time to a usable volume, and avoids cached-volume usage that is due only to the full format process and not any application-data related activity.



- g. Click **Finish** to close the wizard.

Note

The time it takes to format the volume depends on the size of the volume and may take several minutes to complete.



Step 2.4: Test the Setup

By this point, you have an activated gateway with one iSCSI storage volume. Now you are ready to test your setup by writing data to the volume, taking a snapshot, and restoring the snapshot to another volume.

Click one of the following links and follow instructions to test your setup.

To...	Do This...
Test the Setup (Gateway-Cached architecture)	Follow the steps in Test the Setup (Gateway-Cached) (p. 80) .
Test the Setup (Gateway-Stored architecture)	Follow the steps in Test the Setup (Gateway-Stored) (p. 83) .

Test the Setup (Gateway-Cached)

In this section, you verify the setup by taking a snapshot backup of your gateway-cached volume. You then restore it on another volume.

This requires you to first create a snapshot of your volume. You then create another volume from this snapshot. Your gateway copies the data from the specified snapshot in AWS to the new volume.

To create a snapshot of a storage volume

1. On your Windows computer, copy some data to your mapped storage volume.

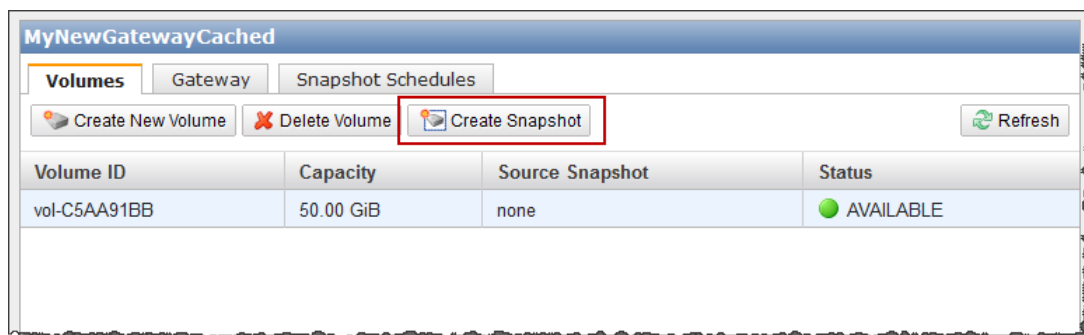
The amount of data copied doesn't matter for this demonstration. A small file is enough to demonstrate the restore.

2. In the **Navigation** pane of the AWS Storage Gateway console, select the gateway.
3. In the **Volumes** tab, select the storage volume created for the gateway.


There should be only one storage volume for this gateway. Selecting the volume displays its properties.

4. Click the **Create Snapshot** button to create a snapshot of the volume.

Depending on the amount of data on the disk and upload bandwidth, it may take a few seconds to complete the snapshot. Note the volume ID from which you create a snapshot. The ID will be used to find the snapshot.



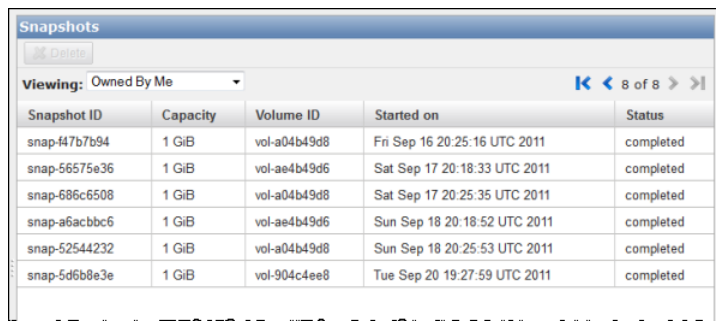
5. In the **Snapshot Taken** confirmation window, click **Close**.



6. In the **Navigation** pane, click **Snapshots**, and find the snapshot that you just created.

You can use the **Started on** column value and the volume ID you noted earlier to confirm the snapshot's source. Note the **Started on** time is UTC time.

The **Status** of your snapshot may be **pending**. In this case, you must wait for the snapshot **Status** to turn to **completed** before restoring the snapshot.



Snapshot ID	Capacity	Volume ID	Started on	Status
snap-f47b7b94	1 GiB	vol-a04b49d8	Fri Sep 16 20:25:16 UTC 2011	completed
snap-56575e36	1 GiB	vol-ae4b49d6	Sat Sep 17 20:18:33 UTC 2011	completed
snap-686c6508	1 GiB	vol-a04b49d8	Sat Sep 17 20:25:35 UTC 2011	completed
snap-a6acbbc6	1 GiB	vol-ae4b49d6	Sun Sep 18 20:18:52 UTC 2011	completed
snap-52544232	1 GiB	vol-a04b49d8	Sun Sep 18 20:25:53 UTC 2011	completed
snap-5d6b8e3e	1 GiB	vol-904c4ee8	Tue Sep 20 19:27:59 UTC 2011	completed

7. Copy the **Snapshot ID** so you can enter it in a subsequent step when you create a storage volume based on the snapshot.

To restore the snapshot

1. In the AWS Storage Gateway console, click the name of gateway in the navigation pane.
2. Click the **Volumes** tab, and then click **Create New Volume**.
3. In the **Create Storage Volume** dialog box, enter the following information.



Create Cached Volume close

Create an iSCSI storage volume up to 32TBs in size. This volume will be stored in Amazon S3, with only a cache of recently accessed data kept locally. Your client applications will connect to these volumes over an iSCSI interface. [object Object].

Capacity: TBs (Max: 32 TBs)

iSCSI Target Name:

Based on Snapshot ID:

Host IP:

Port:

- a. In the **Capacity** text box, enter the same capacity as the original volume from which you took the snapshot.
- b. In the **iSCSI Target Name** box, enter a name for your iSCSI target, for example, **myvolumerestored**.

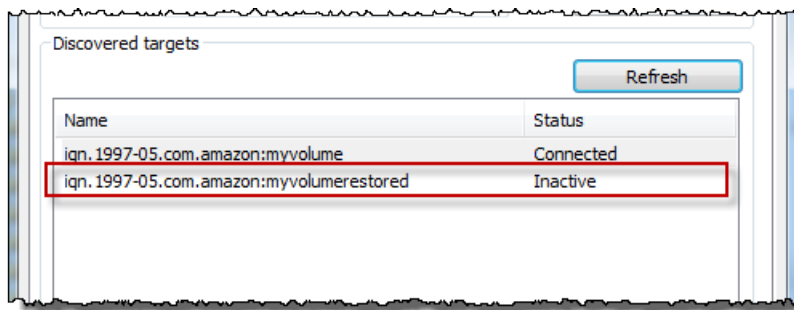
The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name in the **Targets** tab of the **iSCSI Microsoft Initiator** GUI after discovery. For example, a name `target1` would appear as `iqn.1007-05.com.amazon:target1`. Ensure that the target name is globally unique within your SAN network.

- c. In the **Based on Snapshot ID** field, enter the snapshot ID.
- d. Click **Create Cached Volume**.


This creates a storage volume based on your snapshot. The volume details appear in the AWS Storage Gateway console.

- 4. Connect to the new volume target.
 - a. In the **Start** menu of your Windows client computer, type `iscsicpl.exe` and run the program.
 - b. In the **iSCSI Initiator Properties** dialog box, click the **Targets** tab. If the new target does not appear in the **Discovered Targets** pane, click **Refresh**.

You should see both the original target and the new target. The new target will have a status of **Inactive**.

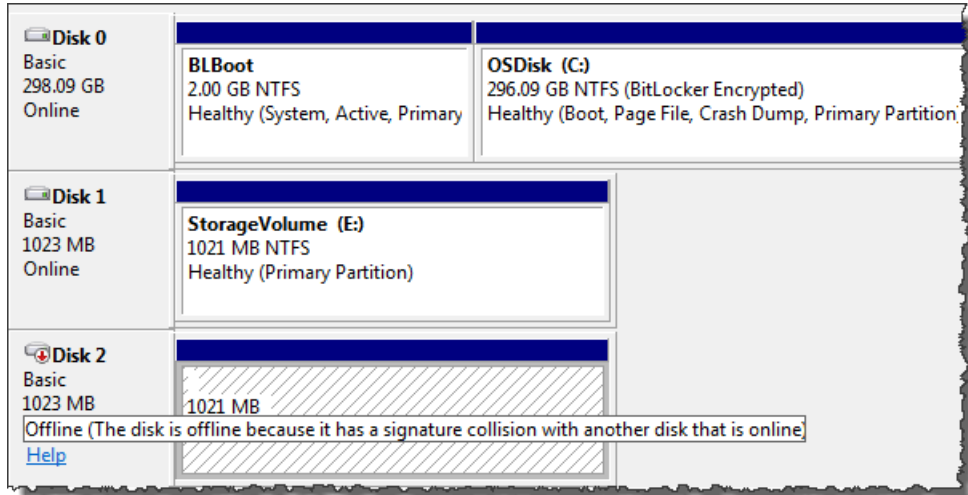


- c. Select the new target, and click **Connect**.
- d. In the **Connect to Target** dialog box, click **OK**.

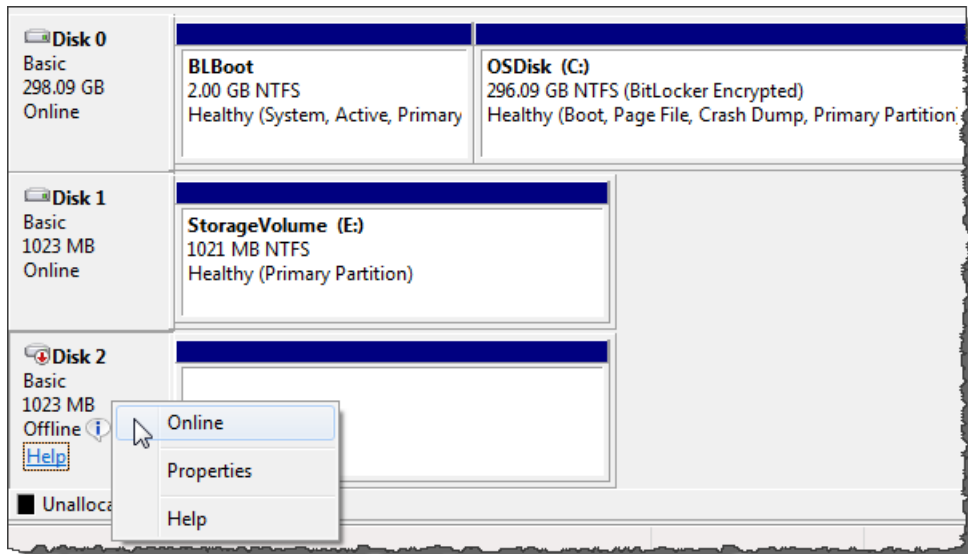


- 5. Bring the restored volume online.
 - a. If the **Disk Management** console is not already open, then in the **Start** menu, type `diskmgmt.msc`.

The restored storage volume is shown in the console with a warning.



- b. Right-click the restored volume and select **Online**. This brings the volume online and assigns it a different drive letter.



6. Open the restored volume and verify that the data you saved earlier is there.

Test the Setup (Gateway-Stored)

In this section, you verify the setup by taking a snapshot backup of your storage volume. You then restore it on another storage volume.

This requires you to first create a snapshot of your storage volume. You then add another local disk to your VM for a new storage volume, and create the new storage volume from this snapshot. Your gateway downloads the data from the specified snapshot in AWS to your storage volume's local disk.

To create a snapshot of a storage volume

1. On your Windows computer, copy some data to your mapped storage volume.

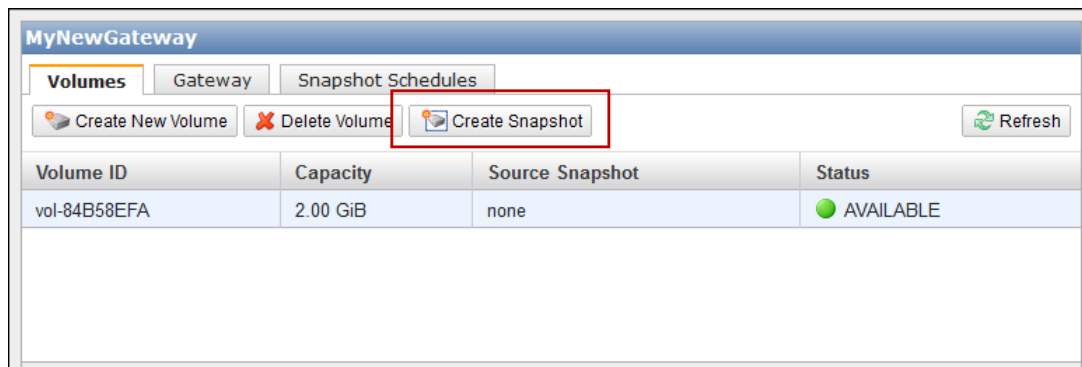
The amount of data copied doesn't matter for this demonstration. A small file is enough to demonstrate the restore.

2. In the AWS Storage Gateway console, select the gateway in the navigation pane.
3. In the **Volumes** tab, select the storage volume created for the gateway.


There should be only one storage volume for this gateway. Selecting the volume displays its properties.

4. Click the **Create Snapshot** button to create a snapshot of the volume.

Depending on the amount of data on the disk and upload bandwidth, it may take a few seconds to complete the snapshot. Note the volume ID from which you create a snapshot. The ID will be used to find the snapshot.



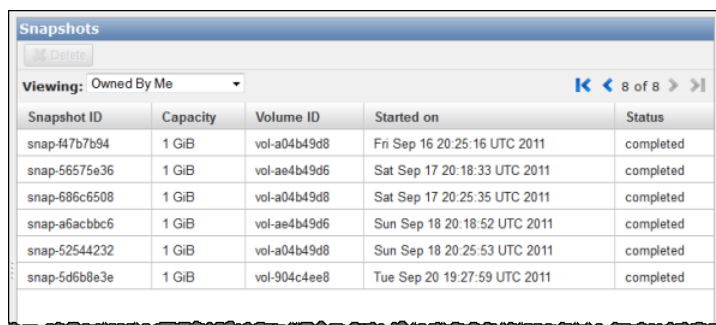
5. In the **Snapshot Taken** confirmation window, click **Close**.



6. In the navigation pane, click **Snapshots**, and find the snapshot that you just created.

You can use the **Started on** column value and the volume ID you noted earlier to confirm the snapshot's source. Note the **Started on** time is UTC time.

The **Status** of your snapshot may be **pending**. In this case, you must wait for the snapshot **Status** to turn to **completed** before restoring the snapshot.



7. Copy the **Snapshot ID** so you can enter it in a subsequent step when you create a storage volume based on the snapshot.

To restore the snapshot

1. Add another virtual disk to your VM that will become a new storage volume on which you restore the snapshot. Since this storage volume was 2 GiB in size, create a new virtual disk of the same size. For instructions, see [To allocate a local disk to store your application data \(p. 24\)](#).
2. In the AWS Storage Gateway console, click the name of the gateway in the navigation pane.
3. Click the **Volumes** tab, and then click **Create New Volume**.
4. In the **Create Storage Volume** dialog box, enter the following information.



- a. In the **Disk** drop-down list, select the virtual disk that you added in the preceding step.
- b. In the **iSCSI Target Name** box, enter a name for your iSCSI target, for example, `myvolumerestored`.

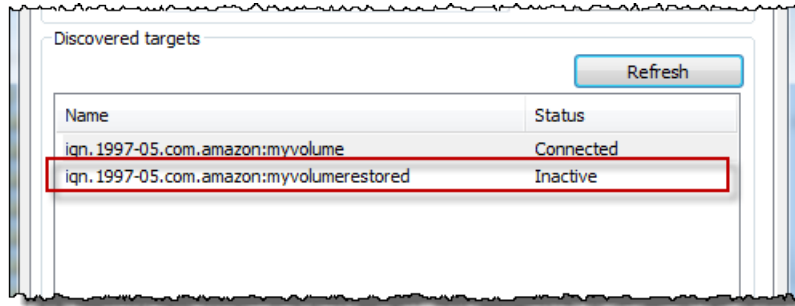
The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name in the **Targets** tab of the **iSCSI Microsoft Initiator** GUI after discovery. For example, a name `target1` would appear as `iqn.1007-05.com.amazon:target1`. Ensure that the target name is globally unique within your SAN network.

- c. In the **Based on Snapshot ID** box, enter the snapshot ID.
- d. Click **Create Volume**.


This creates a storage volume based on your snapshot. The storage volume details appear in the AWS Storage Gateway console.

5. Connect to the new volume target.
 - a. In the **Start** menu of your Windows client computer, type `iscsicpl.exe` and run the program.
 - b. In the **iSCSI Initiator Properties** dialog box, click the **Targets** tab. If the new target does not appear in the **Discovered Targets** pane, click **Refresh**.

You should see both the original target and the new target. The new target will have a status of **Inactive**.

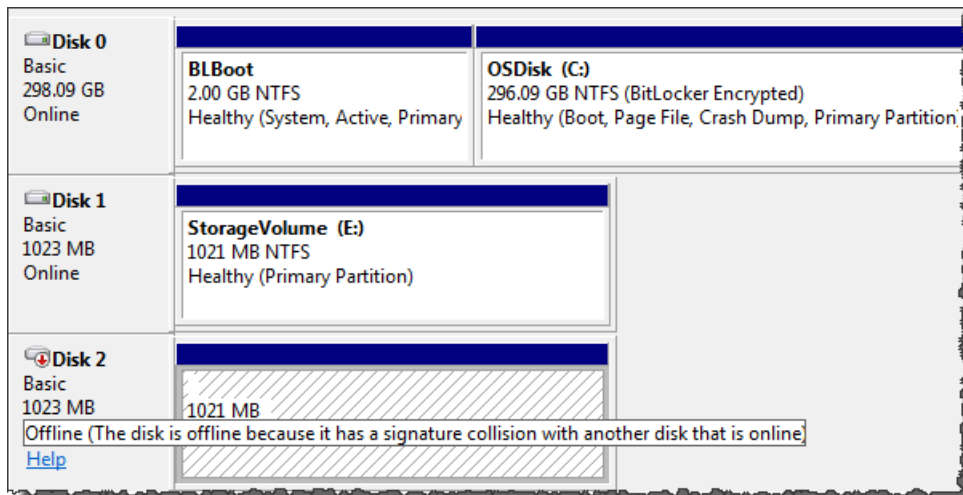


- c. Select the new target, and click **Connect**.
- d. In the **Connect to Target** dialog box, click **OK**.

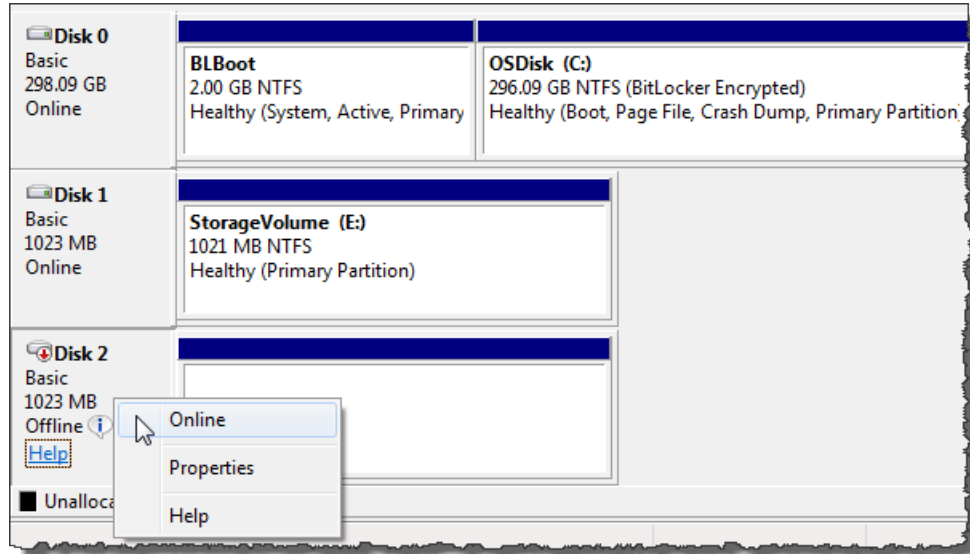


- 6. Bring the restored volume online.
 - a. If the **Disk Management** console is not already open, then in the **Start** menu, type `diskmgmt.msc`.

The restored storage volume is shown in the console with a warning.



- b. Right-click the restored volume, and select **Online**. This brings the volume online and assigns it a different drive letter.



7. Open the restored volume and verify that the data you saved earlier is there.

Where Do I Go from Here?

The AWS Storage Gateway service provides an easy way for you to back your application storage with the storage infrastructure of the AWS cloud. In [Getting Started with AWS Storage Gateway \(p. 7\)](#), you created and provisioned a gateway, and then connected your Windows host to the gateway's storage volume. You added data to the gateway's iSCSI volume, took a snapshot of the volume and restored it to a new volume, and connected to the new volume and verified that the data shows up on it.

After you finish the Getting Started exercise:

- If you plan on continuing to use your gateway, you should read about sizing the upload buffer more appropriately for real-world workloads. For more information, see [Sizing Your Gateway's Storage for Real-World Workloads \(p. 88\)](#).
- If you do not plan on continuing to use your gateway, consider deleting the gateway to avoid incurring any charges. For more information, see [Deleting a Gateway Using the AWS Storage Gateway Console \(p. 233\)](#).

Other sections of this guide include information about how to:

- Learn more about storage volumes and how to create them (see [Managing Storage Volumes in AWS Storage Gateway \(p. 176\)](#)).
- Troubleshoot gateway problems (see [Troubleshooting in AWS Storage Gateway \(p. 252\)](#)).
- Optimize your gateway (see [Optimizing AWS Storage Gateway Performance \(p. 260\)](#)).
- Understand Storage Gateway metrics and how you can monitor how your gateway performs (see [Monitoring Your AWS Storage Gateway \(p. 261\)](#)).
- Connect to the gateway's iSCSI targets to store data (see [Configuring Your Application Access to Storage Volumes \(p. 161\)](#)).

Sizing Your Gateway's Storage for Real-World Workloads

By this point, you have a simple, working gateway. However, because the assumptions used to create this gateway are not appropriate for real-world workloads, you need to do two things: Size your upload buffer appropriately, and set up monitoring for your upload buffer, if you haven't done so already. This step shows how to do both of these tasks. If you activated a gateway for cached-volumes, you also need to size cache storage for real-world workloads.

To size your upload buffer and cache storage for a gateway-cached setup

- Use the formula shown in [Sizing the Upload Buffer \(Gateway-Cached\) \(p. 98\)](#) for sizing the upload buffer and the formula in [Sizing Cache Storage \(Gateway-Cached\) \(p. 94\)](#) for cache storage. We strongly recommend that you allocate at least 150 GiB for the upload buffer. Therefore, if the upload buffer formula yields a value less than 150 GiB, use 150 GiB as your allocated upload buffer.

The upload buffer formula takes into account the difference between throughput from your application to your gateway and throughput from your gateway to AWS, multiplied by how long you expect to write data. For example, assume that your applications write text data to your gateway at a rate of 40 MB per second for 12 hours a day and your network throughput is 12 MB per second. Assuming a compression factor of 2:1 for the text data, the formula specifies that you need to allocate approximately 675 GiB of upload buffer space.

To size your upload buffer for a gateway-stored setup

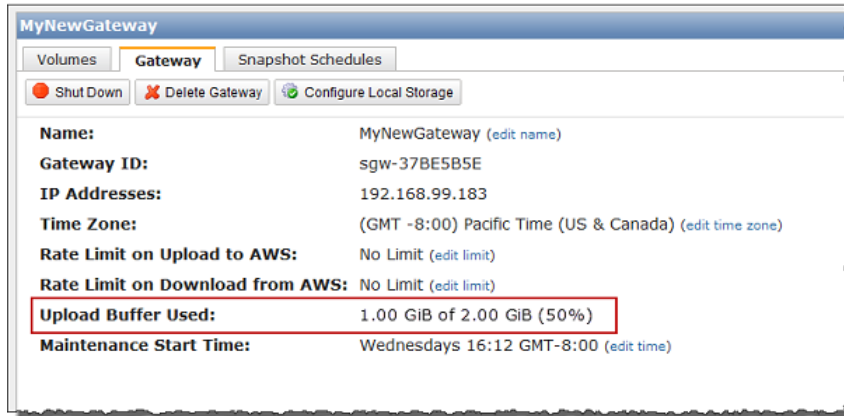
- Use the formula discussed in [Sizing the Upload Buffer \(Gateway-Stored\) \(p. 106\)](#). We strongly recommend that you allocate at least 150 GiB for your upload buffer. Therefore, if the upload buffer formula yields a value less than 150 GiB, use 150 GiB as your allocated upload buffer.

The upload buffer formula takes into account the difference between throughput from your application to your gateway and throughput from your gateway to AWS, multiplied by how long you expect to write data. For example, assume that your applications write text data to your gateway at a rate of 40 MB per second for 12 hours a day and your network throughput is 12 MB per second. Assuming a compression factor of 2:1 for the text data, the formula specifies that you need to allocate approximately 675 GiB of upload buffer space.

To monitor your upload buffer

1. View your gateway's current upload buffer.
 - In the **Gateway** tab in the AWS Storage Gateway console, find the **Upload Buffer Used** field.

The following example shows the upload buffer at three percent.



2. Set an alarm on upload buffer.

We highly recommend that you create a upload buffer alarm in the Amazon CloudWatch console. For more information, see [To set an upper threshold alarm for a gateway's upload buffer \(p. 268\)](#).

Setting Up AWS Storage Gateway

Topics

- [Deploying and Activating Up AWS Storage Gateway On-Premises \(p. 90\)](#)
- [Deploying and Activating AWS Storage Gateway on Amazon EC2 \(p. 137\)](#)
- [Configuring Upload Buffer and Cache Storage \(p. 150\)](#)
- [Creating Storage Volumes \(p. 157\)](#)
- [Configuring Your Application Access to Storage Volumes \(p. 161\)](#)

In this section, we show you how to set up AWS Storage Gateway in two different hosting environments.

- **On-premises**—Host your gateway locally, using an on-premises virtualization environment. To get started, see [Deploying and Activating Up AWS Storage Gateway On-Premises \(p. 90\)](#).
- **Cloud**—Use an Amazon EC2 instance to host a gateway in the cloud for disaster recovery and on-demand compute scenarios. To get started, see [Deploying and Activating AWS Storage Gateway on Amazon EC2 \(p. 137\)](#).

After you deploy and activate a gateway into a hosting environment, the remaining steps for setting up a gateway are the same regardless of the hosting environment and include configuring upload buffer and cache storage, creating storage volume, and configuring access to the storage volumes.

Deploying and Activating Up AWS Storage Gateway On-Premises

In this section, we show you how to download the AWS Storage Gateway and deploy the gateway VM on a virtualization host. We show you how to use both the VMware ESXi and Microsoft Hyper-V virtualization environments to host your gateway. The gateway deployment process for the two host types is conceptually similar. After setting up and activating your gateway, the remaining steps (creating and accessing volumes) are the same for both host types. We recommend that you review the getting started section (see [Getting Started with AWS Storage Gateway \(p. 7\)](#)), before continuing in this section.

To begin setup

1. Go to [AWS Storage Gateway](#) and click **Sign Up**.

AWS Storage Gateway User Guide


Deploying and Activating a Gateway on a VMware ESXi Host

You must sign up for the service before you can download and deploy AWS Storage Gateway.

2. If you don't have a gateway activated under your AWS account, your AWS Storage Gateway console experience starts with the following page. Click the **Setup and Activate a New Gateway** button to start the **Setup and Activate Gateway** wizard.



If you already have one or more gateways activated, the console shows a list of your gateways. In the **Navigation** pane, click **Deploy a new Gateway** to start the **Setup and Activate New Gateway** wizard.




Deploying and Activating AWS Storage Gateway On-Premises on a VMware ESXi Host

This section explains how you can use VMware ESXi to create an on-premises virtual machine to host AWS Storage Gateway. The tasks described here assume you have already provisioned a VMware ESXi host. If you have not already done so, see [Provision a VMware Host to Deploy the AWS Storage Gateway VM \(p. 9\)](#) which was part of the getting started exercise.

Downloading and Deploying AWS Storage Gateway VM

After provision a VMware ESXi host, the next steps in the **Setup and Activate Gateway** wizard are to select the VMware ESXi platform, download the VM software for this platform, and then deploy the VM.



For instructions, see [Download and Deploy the AWS Storage Gateway VM on Your Host \(p. 10\)](#).

Using AWS Storage Gateway with VMware High Availability

VMware High Availability (HA) is a component of vSphere that can provide protection from failures in your infrastructure layer supporting a gateway VM. VMware HA does this by using multiple hosts configured as a cluster so that if a host running a gateway VM fails, the gateway VM can be restarted automatically on another host within the cluster. For more information about VMware HA, go to [VMware HA: Concepts and Best Practices](#).

AWS Storage Gateway should be used with VMware HA with the following recommendations.

- Deploy the gateway Open Virtualization Application (OVA) on only one host in a cluster.
- When deploying the OVA, select a datastore that is not local to one host. Instead, use a datastore that is accessible to all hosts in the cluster. If you select a datastore that is local to a host and the host fails, then the data source may not be accessible to other hosts in the cluster and the failover may not succeed.
- Follow the recommended iSCSI settings to prevent your initiator from disconnecting from storage volume targets during failover. In a failover event, it could take between a few seconds to several minutes for a gateway VM to start in a new host in the failover cluster. The recommended iSCSI timeouts for Windows clients (see [Customizing Your Windows iSCSI Settings \(p. 163\)](#)) and Linux clients (see [Customizing Your Linux iSCSI Settings \(p. 166\)](#)) are greater than the typical time it takes for failover to occur.
- With clustering, if you deploy the OVA to the cluster, you will be asked to select a host. Alternately, you can deploy directly to a host in a cluster.

Provisioning Local Disk Storage for an AWS Storage Gateway VM

Topics

- [About the Disk the Gateway VM Uses to Store System Data \(p. 93\)](#)
- [Provisioning Local Disks \(Gateway-Cached\) \(p. 93\)](#)
- [Provisioning Local Disks \(Gateway-Stored\) \(p. 102\)](#)
- [Configure AWS Storage Gateway VM to Use Paravirtualization \(p. 110\)](#)

Before you provision local disk storage for the gateway VM you deployed, you should decide the configuration for your AWS Storage Gateway. You have the following options:

- **Use Gateway-Cached volumes**— In this configuration, the gateway stores your volume data in Amazon S3.

The gateway maintains a cache storage for recently accessed data to provide low-latency access. The gateway persistently holds the data that has not been uploaded to Amazon S3 in the cache storage; therefore, you must allocate disks on-premises for the cache storage. You must also allocate disks for the upload buffer to temporarily buffer your data prior to uploading to AWS. The cache storage should be larger than the upload buffer (see [How AWS Storage Gateway Works \(p. 3\)](#)).

- **Use Gateway-Stored volumes**—In this configuration, the gateway stores your volume data on your on-premises storage hardware.

You must allocate disks on-premises to hold all your data. The gateway then securely uploads data snapshots to Amazon S3 for cost-effective backup and rapid disaster recovery. You must also allocate disks for the gateway's upload buffer (see [How AWS Storage Gateway Works \(p. 3\)](#)).

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a VMware ESXi Host

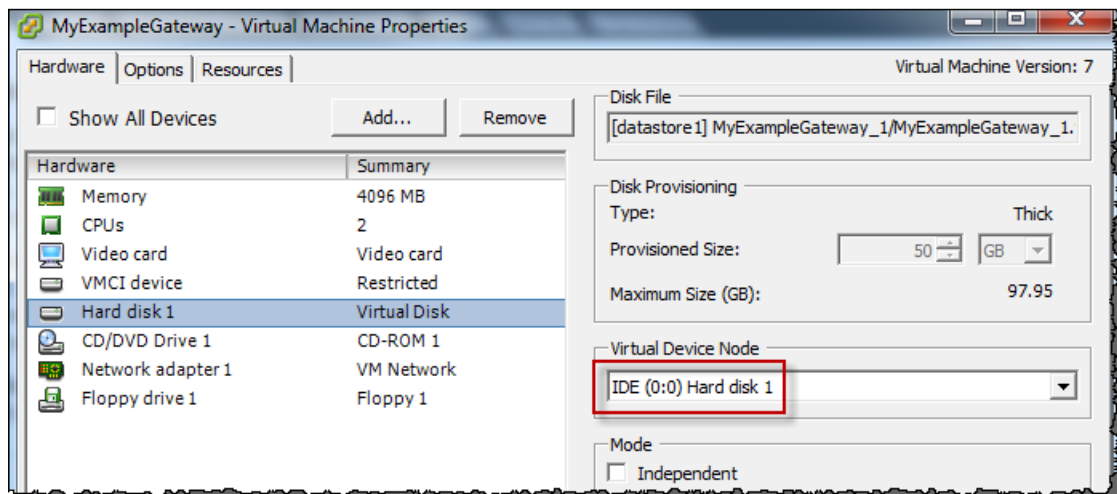
If you follow the **Setup and Activate Gateway** console wizard, the console shows the following prompt for you to choose the volume type.



After selecting the volume type, you must provision local disks to the gateway VM required to support the volume type you selected before activating the gateway.

About the Disk the Gateway VM Uses to Store System Data

After you deploy the gateway VM, it includes preconfigured processors, memory, and an IDE disk with the VM infrastructure on it. This IDE disk appears as IDE (0:0) Hard disk1 in the **Virtual Machine Properties** window, in the vSphere client, as shown in the following example screen shot. However, you cannot access or use this disk directly. The gateway uses it to store system data.



Provisioning Local Disks (Gateway-Cached)

Topics

- [Adding Local Disks for Cache Storage \(Gateway-Cached\) \(p. 94\)](#)
- [Adding Local Disks for the Upload Buffer \(Gateway-Cached\) \(p. 98\)](#)

In the gateway-cached architecture, the gateway stores your volume data in Amazon S3. However, you must provision disks to the gateway VM for cache storage and the upload buffer. For more information about how the gateway works, see [How AWS Storage Gateway Works \(p. 3\)](#)).

Note

When you provision disks, we strongly recommend that you do not provision local disks for upload buffer and cache storage that use the same underlying physical storage resource (disk). Underlying physical storage resources are represented as a datastore in VMware. When you deploy the gateway VM, you choose a datastore on which to store the VM files. When you provision a local disk (e.g., to use as cache storage or upload buffer), you have the option to store the virtual disk in the same datastore as the VM or a different datastore. If you have more than one datastore, we strongly recommend that you choose one datastore for the cache storage and another for the upload buffer. A datastore that is backed by only one underlying physical disk, or that is backed by a less-performant RAID configuration such as RAID 1, may lead to poor performance in some situations when used to back both the cache storage and upload buffer.

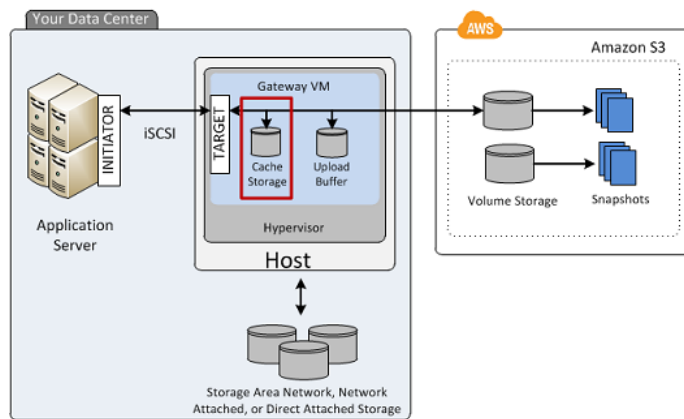
Adding Local Disks for Cache Storage (Gateway-Cached)

Topics

- [Sizing Cache Storage \(Gateway-Cached\) \(p. 94\)](#)
- [Adding a Virtual Disk for Cache Storage \(Gateway-Cached\) \(p. 95\)](#)

In the gateway-cached architecture, your gateway maintains cache storage on-premises for recently accessed data. The gateway persistently holds the data in cache storage that has not be uploaded to Amazon S3. You will need to allocate disks on-premises for cache storage.

The following diagram highlights cache storage in the larger picture of the AWS Storage Gateway architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).



Sizing Cache Storage (Gateway-Cached)

The gateway uses the cache storage to provide low-latency access to your recently accessed data. The cache storage acts as the on-premises durable store for data that is pending upload to Amazon S3 from the upload buffer. So cache storage should be larger than the upload buffer.

The total cache storage for a gateway can be up to 16 TiB.

To estimate the amount of cache storage your gateway needs, the formula depends on your use case:

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a VMware ESXi Host

- **Backup Use Case**—Use a cache storage capacity of 1.1 times the upload buffer capacity. For a backup use case, the cache is durable storage that holds data prior to upload to AWS, and it must be sized greater than the upload buffer to ensure that no data is lost in the event of a VM failure.
- **Other Use Cases**—Use the larger of the following two values: 20 percent of your existing on-premise storage or 1.1 times the upload buffer size.

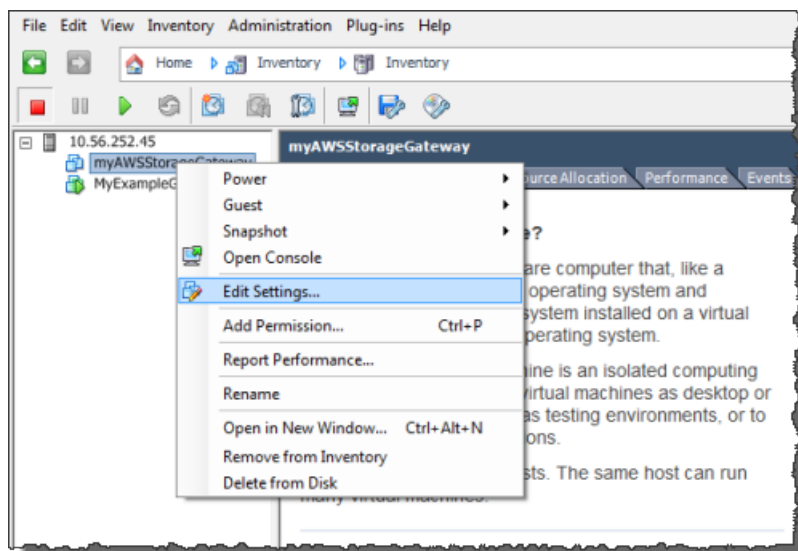
You can initially use this approximation to provision disks for the cache storage. You can then use Amazon CloudWatch operational metrics to monitor the cache storage usage and provision more storage as needed using the console. For using the metrics and setting up alarms, see [Monitoring Cache Storage](#) (p. 271).

Adding a Virtual Disk for Cache Storage (Gateway-Cached)

You can allocate virtual disks to the VM from either the direct-attached storage (DAS) disks or from the storage area network (SAN) disks available on your host. The following procedure provides instructions for adding a virtual disk from a DAS disk that is available on the host.

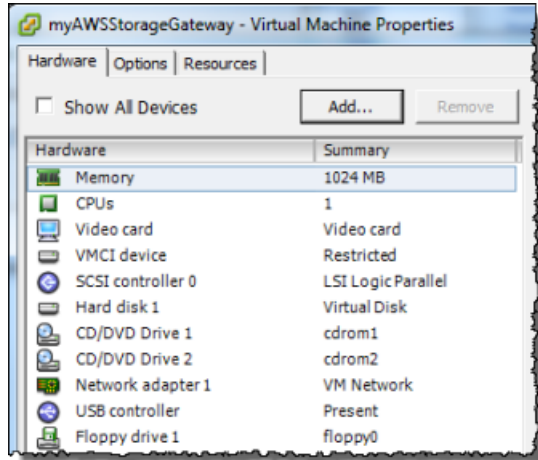
To allocate a new virtual disk to the VM for cache storage

1. Start the VMware vSphere client and connect to your host.
2. In the client, right-click the name of your gateway VM and click **Edit Settings....**




3. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, click **Add...** to add a device.

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a VMware ESXi
Host



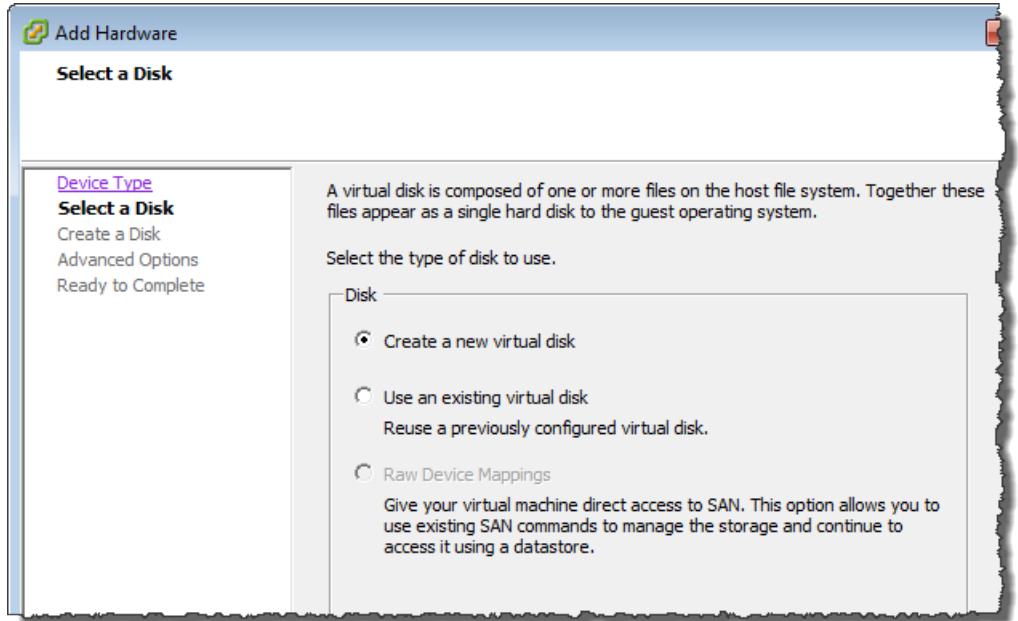
4. Follow the **Add Hardware** wizard to add a disk:
 - a. In the **Device Type** pane, click **Hard Disk** to add a disk, and click **Next**.



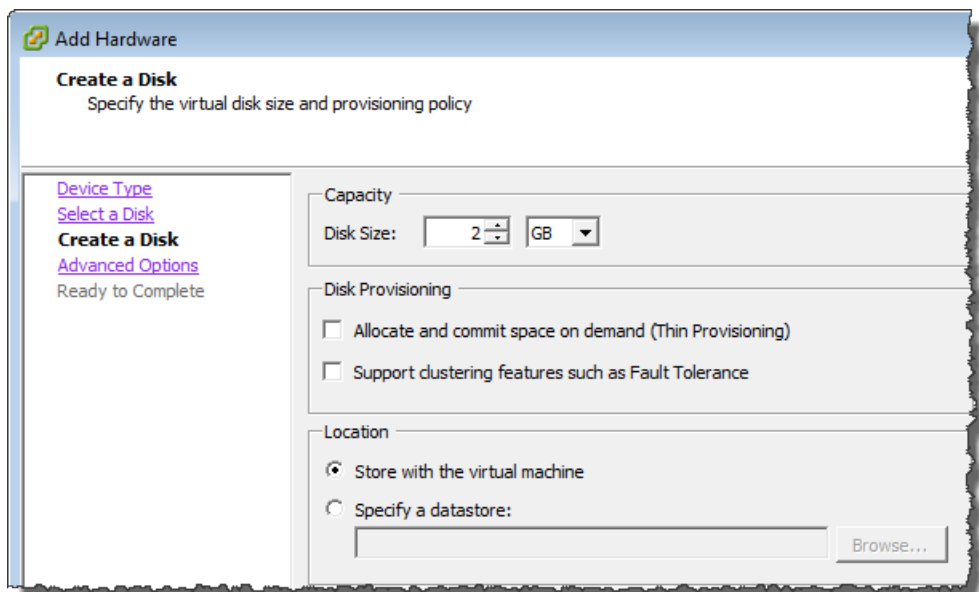
- b. In the **Select a Disk** pane, select **Create a new virtual disk** and click **Next**.

If the disk you are adding for your application storage contains pre-existing data that you want to preserve, select the **Use an existing virtual disk** option.

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a VMware ESXi
Host



- c. In the **Create a Disk** pane, specify the size of the disk and click **Next..**



- d. In the **Advanced Options** pane, click **Next**.
- e. In the **Ready to Complete** pane, click **Finish**.
5. If you have not already done so, you must configure your VM to use a paravirtualized controller for your local disks.

Important

Configuring your VM for paravirtualization is a critical task. If you do not configure paravirtualization, the AWS Storage Gateway console will not be able to communicate with the disks that you have allocated. For steps on configuring paravirtualization, see [Configure AWS Storage Gateway VM to Use Paravirtualization \(p. 110\)](#).

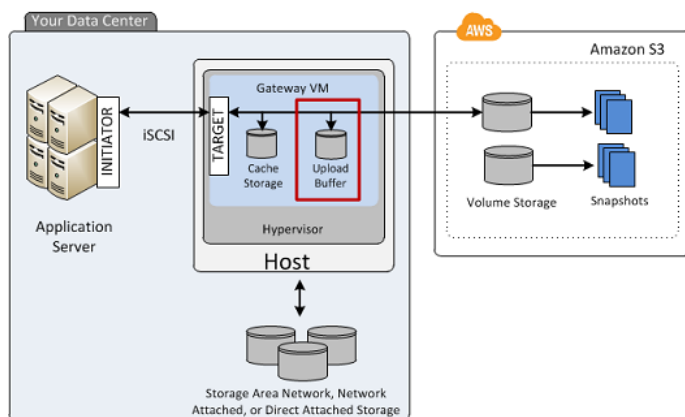
Adding Local Disks for the Upload Buffer (Gateway-Cached)

Topics

- [Sizing the Upload Buffer \(Gateway-Cached\)](#) (p. 98)
- [Adding a Virtual Disk for the Upload Buffer \(Gateway-Cached\)](#) (p. 99)

You must allocate disk(s) on your premises for the gateway to use as the upload buffer to temporarily buffer your data prior to uploading to AWS.

The following diagram highlights the upload buffer in the larger picture of the AWS Storage Gateway architecture (see [How AWS Storage Gateway Works](#) (p. 3)).



Sizing the Upload Buffer (Gateway-Cached)

You can determine the size of your upload buffer by using an upload buffer formula. We strongly recommend that you allocate at least 150 GiB of upload buffer. Therefore, if the formula returns a value less than 150 GiB, use 150 GiB as the amount you allocate to upload buffer. You can configure up to 2 TiB of upload buffer capacity per gateway.

Note

When the upload buffer reaches its capacity, your applications can continue to read from and write data to your storage volumes; however, the gateway will not write any of your volume data to its upload buffer and will not upload any of this data to AWS.

To estimate the amount of upload buffer more precisely, you can calculate the incoming and outgoing data rates and base an estimate on these rates.

- **Rate of Incoming Data**—This refers to the application throughput, the rate at which your on-premises applications are writing data to your gateway over some period of time.
- **Rate of Outgoing Data**—This refers to the network throughput, the rate at which your gateway is able to upload data to AWS. This depends on your network speed, utilization, and whether you've enabled bandwidth throttling. This rate should be adjusted for compression. When uploading data to AWS, the gateway applies data compression where possible. For example, if your application data is text-only, you might get effective compression ratio of about 2:1. However, if you are writing videos, the gateway might not be able to achieve any data compression, requiring more upload buffer for the gateway.

If your incoming rate is higher than the outgoing rate, you can use the following formula to determine the approximate size of the upload buffer your gateway needs.

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a VMware ESXi Host

$$\left(\begin{array}{l} \text{Application} \\ \text{Throughput} \\ \text{(MB/s)} \end{array} - \begin{array}{l} \text{Network} \\ \text{Throughput} \\ \text{to AWS (MB/s)} \end{array} \right) \times \begin{array}{l} \text{Compression} \\ \text{Factor} \end{array} \times \begin{array}{l} \text{Duration} \\ \text{of writes} \\ \text{(s)} \end{array} = \begin{array}{l} \text{Upload} \\ \text{Buffer} \\ \text{(MB)} \end{array}$$

For example, assume that your business applications will write text data to your gateway at a rate of 40 megabytes per second for 12 hours a day and your network throughput is 12 megabytes per second. Assuming a compression factor of 2:1 for the text data, you need to allocate approximately 690 GB of space for the upload buffer.

```
((40 MB/sec) - (12 MB/sec * 2)) * (12 hours * 3600 seconds/hour) = 691200 megabytes
```

Note that you can initially use this approximation to determine the disk size that you want to allocate to the gateway as upload buffer space. Add more upload buffer space as needed using the AWS Storage Gateway console. Also, you can use the Amazon CloudWatch operational metrics to monitor upload buffer usage and determine additional storage requirements. For information on metrics and setting the alarms, see [Monitoring the Upload Buffer \(p. 267\)](#).

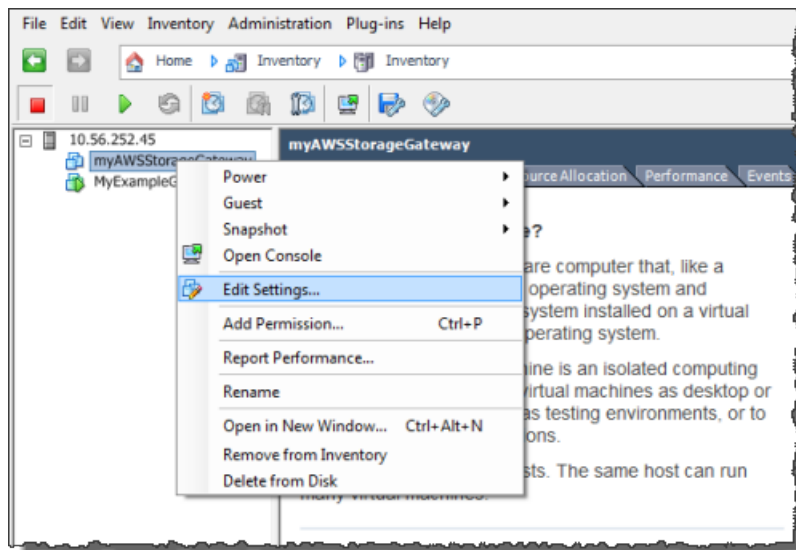
Adding a Virtual Disk for the Upload Buffer (Gateway-Cached)

In this section, you allocate a virtual disk to your VM that will be used as the upload buffer for your gateway.

You can allocate virtual disks to the VM from either the direct-attached storage (DAS) disks or from the storage area network (SAN) disks available on your host. The following procedure provides instructions for adding a virtual disk from a DAS disk available on the host.


To allocate a new virtual disk to the VM for the upload buffer

1. Start the VMware vSphere client and connect to your host.
2. In the client, right-click the name of your gateway VM and click **Edit Settings....**

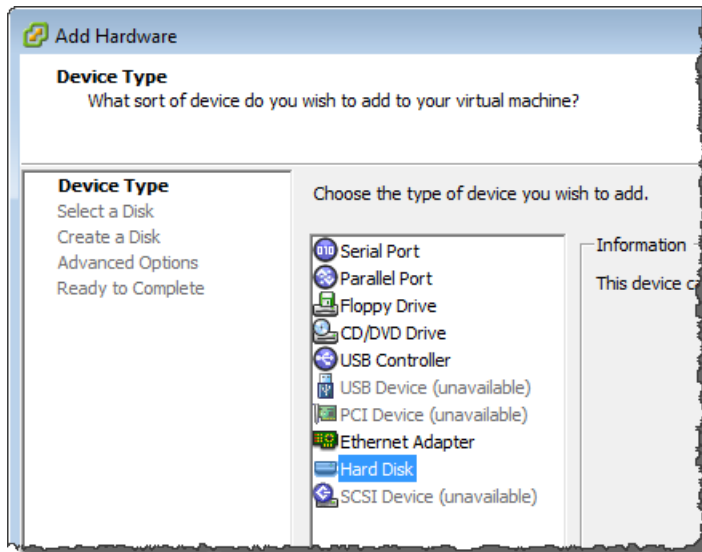


3. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, click **Add...** to add a device.

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a VMware ESXi
Host



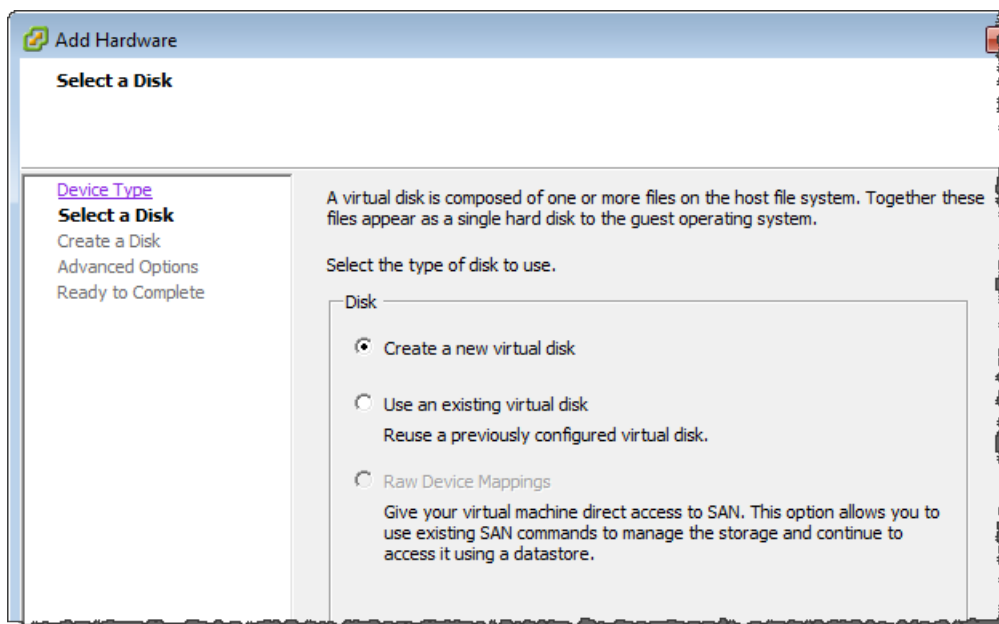
4. Follow the **Add Hardware** wizard to add a disk:
 - a. In the **Device Type** pane, click **Hard Disk** to add a disk, and click **Next**.



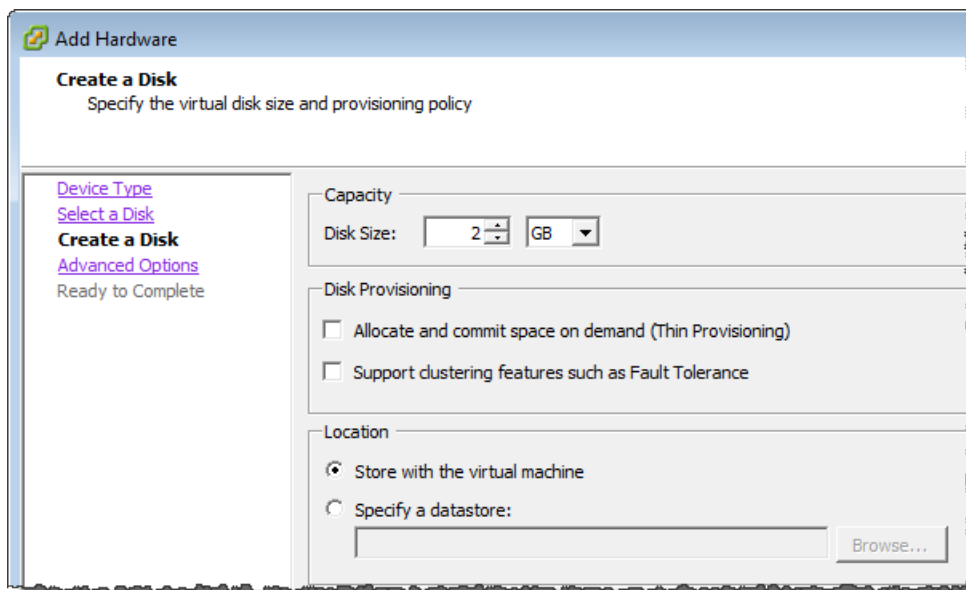
- b. In the **Select a Disk** pane, select **Create a new virtual disk** and click **Next**.

If the disk you are adding for your application storage contains pre-existing data that you want to preserve, select the **Use an existing virtual disk** option.

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a VMware ESXi
Host



- c. In the **Create a Disk** pane, specify the size of the disk and click **Next..**



- d. In the **Advanced Options** pane, click **Next**.
- e. In the **Ready to Complete** pane, click **Finish**.
5. If you have not already done so, you must configure your VM to use a paravirtualized controller for your local disks.

Important

Configuring your VM for paravirtualization is a critical task. If you do not configure paravirtualization, the AWS Storage Gateway console will not be able to communicate with the disks that you have allocated. For steps on configuring paravirtualization, see [Configure AWS Storage Gateway VM to Use Paravirtualization \(p. 110\)](#).

Provisioning Local Disks (Gateway-Stored)

Topics

- [Adding Local Disks for Volume Storage \(Gateway-Stored\) \(p. 102\)](#)
- [Adding Local Disks for Upload Buffer \(Gateway-Stored\) \(p. 105\)](#)

In the gateway-stored architecture, the gateway stores your volume data on your on-premises storage hardware. All your application data reside on your premises. You must provision disks to the gateway VM for the volume storage. You must also provision disks for the gateway's upload buffer. For more information about how the gateway works, see [How AWS Storage Gateway Works \(p. 3\)](#).

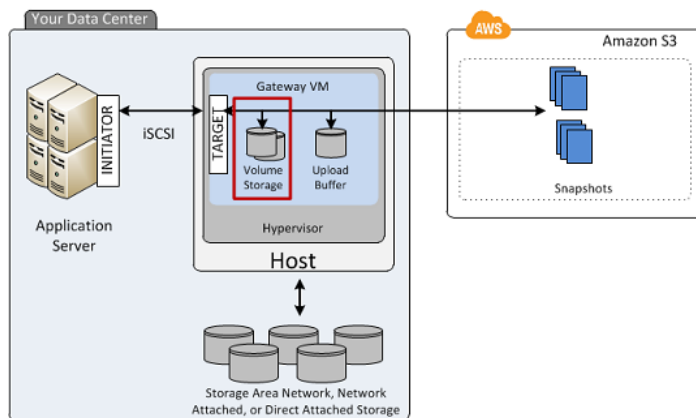
Note

When you provision disks, it is strongly recommended that you do not provision local disks for upload buffer and local application storage that use the same underlying physical storage resource (disk). Underlying physical storage resources are represented as a datastore in VMware. When you deploy the gateway VM, you choose a datastore on which to store the VM files. When you provision a local disk (e.g., to use as local application storage or upload buffer), you have the option to store the virtual disk in the same datastore as the VM or a different datastore. If you have more than one datastore, then it is strongly recommended that you choose one datastore for the local application storage and another for the upload buffer. A datastore that is backed by only one underlying physical disk, or that is backed by a less-performant RAID configuration such as RAID 1, may lead to poor performance in some situations when used to back both the local application storage and upload buffer.

Adding Local Disks for Volume Storage (Gateway-Stored)

In the gateway-stored architecture, your application data is stored locally. You will need to provision disks to the gateway VM to store your data.

The following diagram highlights storage volumes in the larger picture of the AWS Storage Gateway architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).



Each disk can be up to 1 TiB in size and must be rounded to the nearest GiB, where GiB is calculated using Base 2 (i.e., GiB = 1024³ bytes).

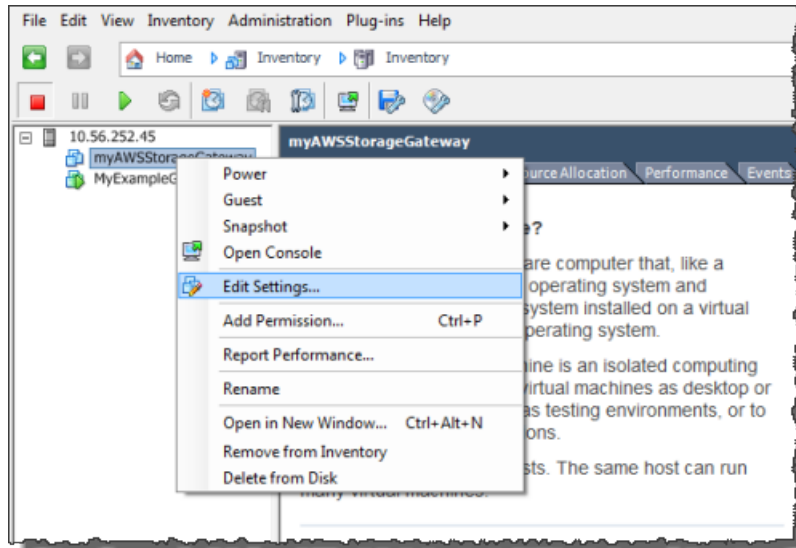
You can provision disks to the VM for volume storage from either the direct-attached storage (DAS) disks or from the storage area network (SAN) disks. For volume storage, the disk you allocate can have existing data. We preserve this data when creating your iSCSI storage volumes. The following procedure provides instructions for adding a virtual disk from a DAS disk.

AWS Storage Gateway User Guide

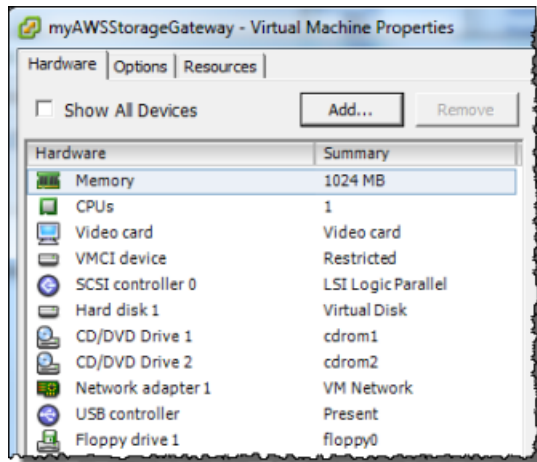
Deploying and Activating a Gateway on a VMware ESXi Host

To allocate a new virtual disk to the VM for application data

1. Start the VMware vSphere client and connect to your host.
2. In the client, right-click the name of your gateway VM and click **Edit Settings...**




3. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, click **Add...** to add a device.



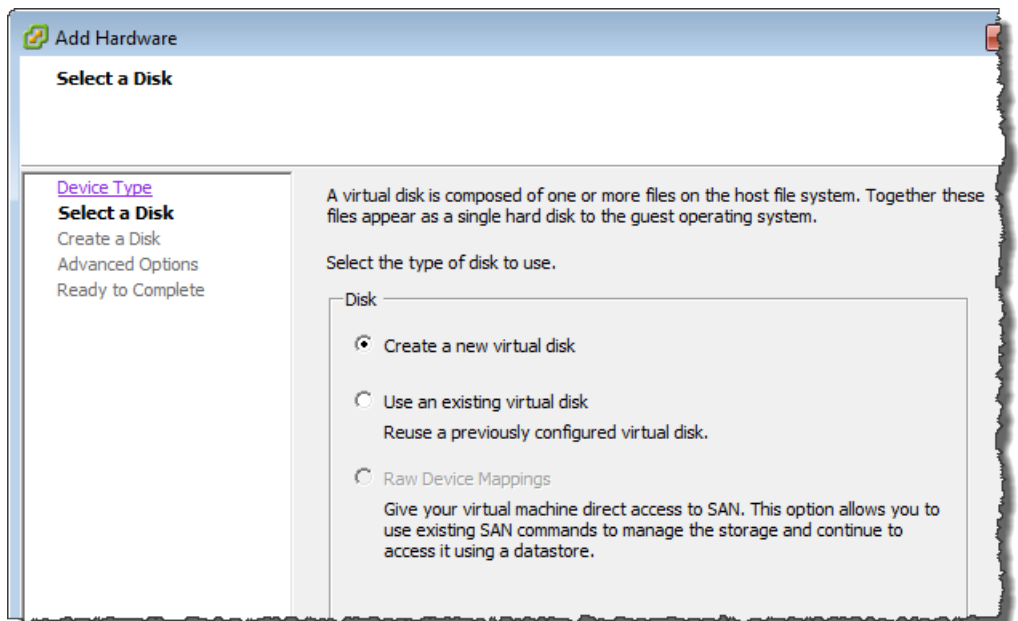
4. Follow the **Add Hardware** wizard to add a disk:
 - a. In the **Device Type** pane, click **Hard Disk** to add a disk, and click **Next**.

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a VMware ESXi Host



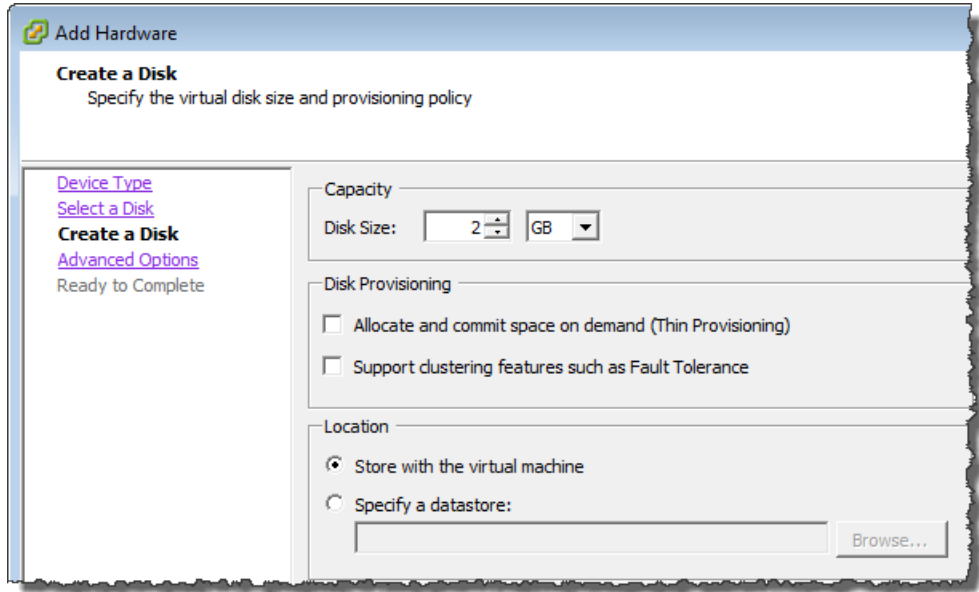
- b. In the **Select a Disk** pane, select **Create a new virtual disk** and click **Next**.

If the disk you are adding for your application storage contains pre-existing data that you want to preserve, select the **Use an existing virtual disk** option.



- c. In the **Create a Disk** pane, specify the size of the disk and click **Next**.

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a VMware ESXi
Host



- d. In the **Advanced Options** pane, click **Next**.
 - e. In the **Ready to Complete** pane, click **Finish**.
5. If you have not already done so, you must configure your VM to use a paravirtualized controller for your local disks.

Important

Configuring your VM for paravirtualization is a critical task. If you do not configure paravirtualization, the AWS Storage Gateway console will not be able to communicate with the disks that you have allocated. For steps on configuring paravirtualization, see [Configure AWS Storage Gateway VM to Use Paravirtualization](#) (p. 110).

Adding Local Disks for Upload Buffer (Gateway-Stored)

Topics

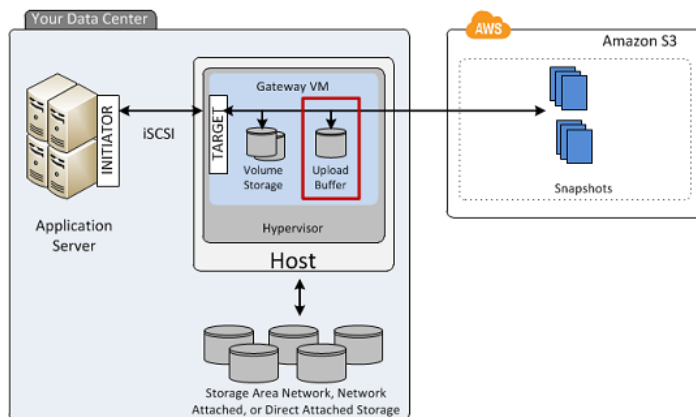
- [Sizing the Upload Buffer \(Gateway-Stored\)](#) (p. 106)
- [Adding a Virtual Disk for the Upload Buffer \(Gateway-Stored\)](#) (p. 107)

You must allocate disk(s) on your premises for the gateway to use as the upload buffer to temporarily buffer your data prior to uploading to AWS.

The following diagram highlights the upload buffer in the larger picture of the AWS Storage Gateway architecture (see [How AWS Storage Gateway Works](#) (p. 3)).

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a VMware ESXi Host



Sizing the Upload Buffer (Gateway-Stored)

You can determine the size of your upload buffer by using an upload buffer formula. It is strongly recommended that you allocate at least 150 GiB of upload buffer space. Therefore, if the formula returns a value less than 150 GiB, use 150 GiB as the amount you allocate to the upload buffer. You can configure up to 2 TiB of upload buffer capacity per gateway.

Note

When the upload buffer reaches its capacity, your applications can continue to read from and write data to your storage volumes; however, the gateway is not writing any of your volume data to its upload buffer and not uploading any of this data to AWS.

To estimate the amount of upload buffer space, calculate the incoming and outgoing data rates and base an estimate on these rates.

- **Rate of Incoming Data**—This refers to the application throughput, the rate at which your on-premises applications are writing data to your gateway over some period of time.
- **Rate of Outgoing Data**—This refers to the network throughput, the rate at which your gateway is able to upload data to AWS. This depends on your network speed, utilization, and whether you've enabled bandwidth throttling. This rate should be adjusted for compression. When uploading data to AWS, the gateway applies data compression where possible. For example, if your application data is text-only, you might get effective compression ratio of about 2:1. However, if you are writing videos, the gateway might not be able to achieve any data compression, requiring more upload buffer space for the gateway.

If your incoming rate is higher than the outgoing rate, you can use the following formula to determine the approximate size of the upload buffer your gateway needs.

$$\left(\begin{array}{c} \text{Application} \\ \text{Throughput} \\ \text{(MB/s)} \end{array} - \begin{array}{c} \text{Network} \\ \text{Throughput} \\ \text{to AWS (MB/s)} \end{array} \right) \times \begin{array}{c} \text{Compression} \\ \text{Factor} \end{array} \times \begin{array}{c} \text{Duration} \\ \text{of writes} \\ \text{(s)} \end{array} = \begin{array}{c} \text{Upload} \\ \text{Buffer} \\ \text{(MB)} \end{array}$$

For example, assume that your business applications will write text data to your gateway at a rate of 40 megabytes per second for 12 hours a day and your network throughput is 12 megabytes per second. Assuming a compression factor of 2:1 for the text data, you need to allocate approximately 690 GB of space for the upload buffer.

$$((40 \text{ MB/sec}) - (12 \text{ MB/sec} * 2)) * (12 \text{ hours} * 3600 \text{ seconds/hour}) = 691200 \text{ megabytes}$$

Note that you can initially use this approximation to determine the disk size that you want to allocate to the gateway as upload buffer space. Add more upload buffer space as needed using the AWS Storage Gateway console. Also, you can use the Amazon CloudWatch operational metrics to monitor upload buffer usage and determine additional storage requirements. For information on metrics and setting the alarms, see [Monitoring the Upload Buffer \(p. 267\)](#).

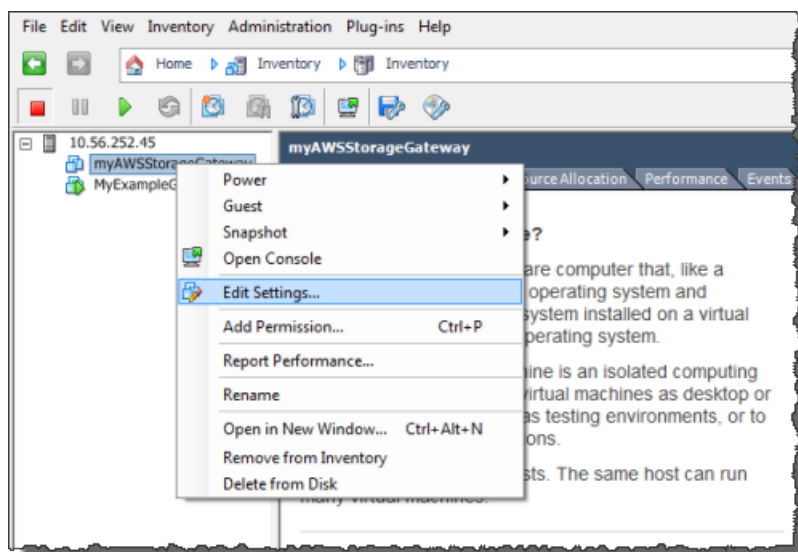
Adding a Virtual Disk for the Upload Buffer (Gateway-Stored)

In this section, you allocate a virtual disk to your VM that will be used as the upload buffer for your gateway. To estimate the upload buffer your gateway requires, see [Sizing the Upload Buffer \(Gateway-Stored\) \(p. 106\)](#).

You can allocate virtual disks to the VM from either the direct-attached storage (DAS) disks or from the storage area network (SAN) disks available on your host. The following procedure provides step-by-step instructions to add a virtual disk from a DAS disk available on the host.


To allocate a new virtual disk to the VM for the upload buffer

1. Start the VMware vSphere client and connect to your host.
2. In the client, right-click the name of your gateway VM and click **Edit Settings....**

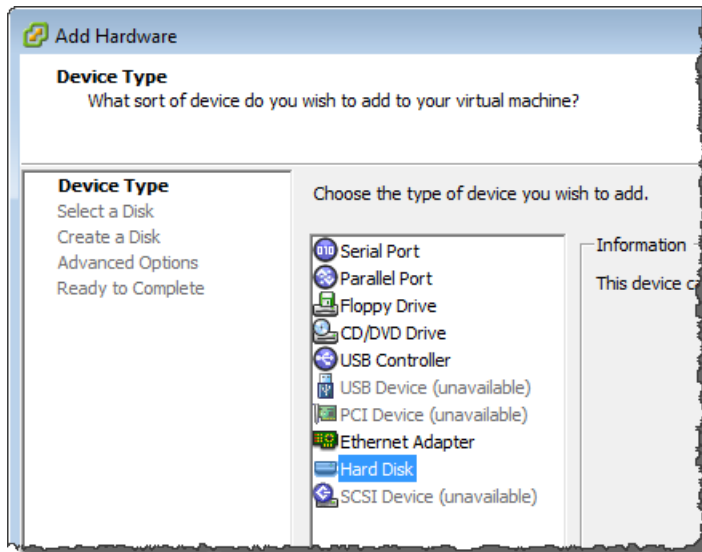


3. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, click **Add...** to add a device.

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a VMware ESXi
Host



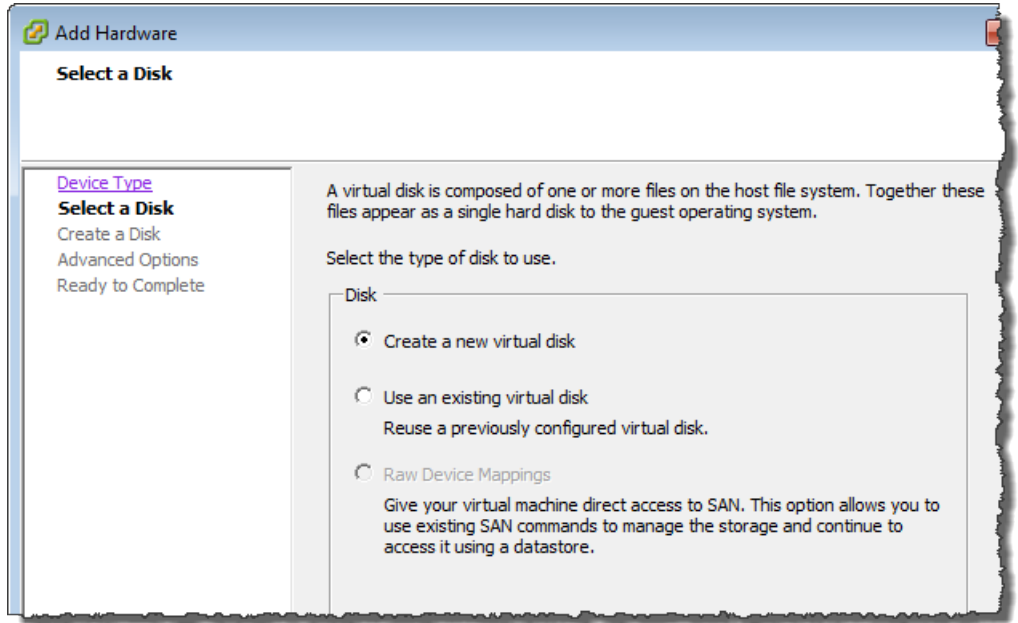
4. Follow the **Add Hardware** wizard to add a disk:
 - a. In the **Device Type** pane, click **Hard Disk** to add a disk, and click **Next**.



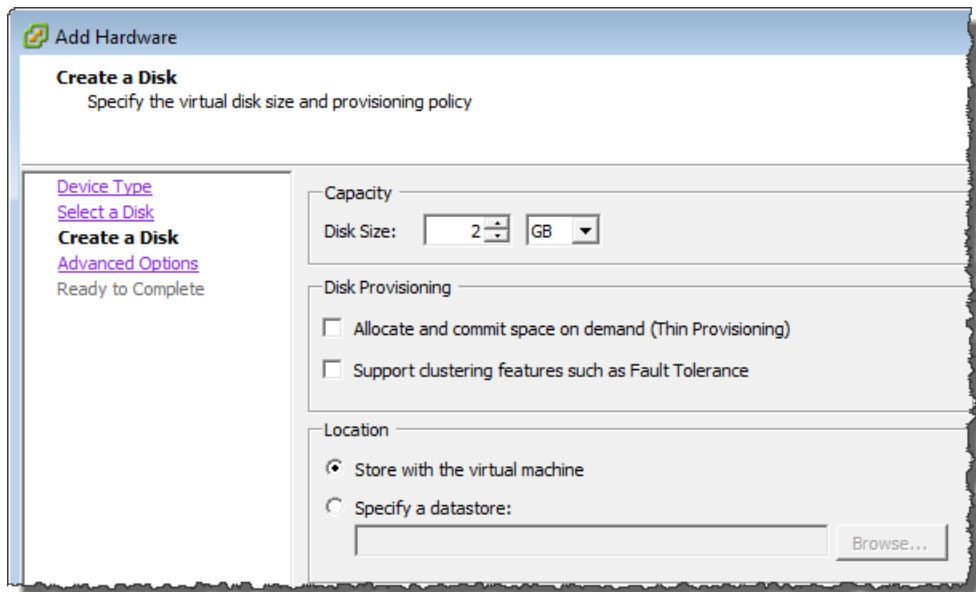
- b. In the **Select a Disk** pane, select **Create a new virtual disk** and click **Next**.

If the disk you are adding for your application storage contains pre-existing data that you want to preserve, select the **Use an existing virtual disk** option.

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a VMware ESXi
Host



- c. In the **Create a Disk** pane, specify the size of the disk and click **Next..**



- d. In the **Advanced Options** pane, click **Next**.
- e. In the **Ready to Complete** pane, click **Finish**.
5. If you have not already done so, you must configure your VM to use a paravirtualized controller for your local disks.

Important

Configuring your VM for paravirtualization is a critical task. If you do not configure paravirtualization, the AWS Storage Gateway console will not be able to communicate with the disks that you have allocated. For steps on configuring paravirtualization, see [Configure AWS Storage Gateway VM to Use Paravirtualization \(p. 110\)](#).

Configure AWS Storage Gateway VM to Use Paravirtualization

In order for the AWS Storage Gateway console to properly recognize your disks, you must configure your VM to use paravirtualized controllers for local disks. In practice, you will set paravirtualization during your initial set up of your gateway, that is, after you deployed the VM, added local disks, but before you power on the VM. To set paravirtualization, the VM must be powered off.

Note

You can only set the virtualization of an iSCSI controller if you have provisioned at least one SCSI disk to the VM. For more information, see [Provisioning Local Disk Storage for an AWS Storage Gateway VM \(p. 92\)](#).

To configure your VM to use paravirtualized controllers

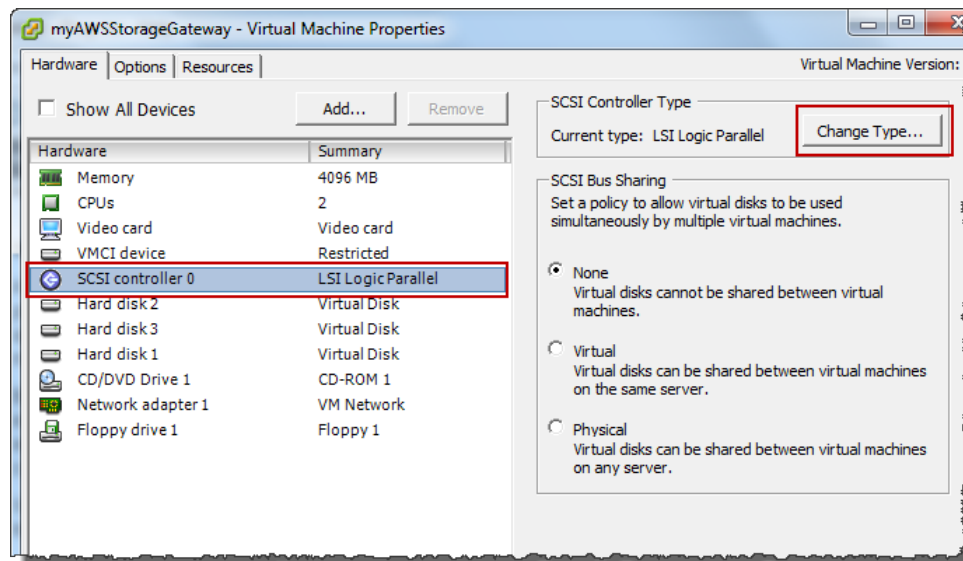
1. In the VMware vSphere client, right-click the name of your gateway virtual machine.

Verify that the VM is powered off. If not, power it off. For more information, see [Steps for Activating a Gateway \(p. 112\)](#). Before powering off the VM, make sure that the gateway is not in use.


2. Select **Edit Settings....**

The **Virtual Machine Properties** dialog box opens.

3. In the **Hardware** tab, select the **SCSI controller 0** setting in the **Hardware** column and click **Change Type....**



4. Select the **VMware ParaVirtual** SCSI controller type and click **OK**.



Activating AWS Storage Gateway

After you deploy the AWS Storage Gateway VM, you must activate the gateway using the AWS Storage Gateway console. The activation process associates your gateway with your AWS account. Once you establish this connection, you can manage almost all aspects of your gateway from the console. In the activation process, you specify the IP address of your gateway, name your gateway, identify the AWS region in which you want your snapshot backups stored, and specify the gateway timezone. After this activation, you begin incurring charges. For information about pricing, see [AWS Storage Gateway](#).

Pre-Activation Checklist

You can activate a gateway after you have completed the steps summarized in the following table. The console wizard walks you through these steps.

Step	Description
Download and Deploy the VM	In the AWS Storage Gateway console, download the latest virtual machine (VM) that is distributed as an .ova file and deploy this VM on your VMware host. For more information, see Downloading and Deploying AWS Storage Gateway VM (p. 91) .
Provision local disks to the VM	The provisioned VM has no disks. Depending on the gateway type you activated you either must add local disks for: <ul style="list-style-type: none"> cache storage and upload buffer for a gateway activated for cached volumes. For more information, see Provisioning Local Disks (Gateway-Cached) (p. 93). application data and upload buffer for a gateway activated for stored volumes. For more information, see Provisioning Local Disks (Gateway-Stored) (p. 102).

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a VMware ESXi
Host

Step	Description
Configure the VM to use paravirtualization	Configuring your VM for paravirtualization is a critical task. If you do not configure paravirtualization, the AWS Storage Gateway console will not be able to communicate with the disks that you have allocated. For more information, see Configure AWS Storage Gateway VM to Use Paravirtualization (p. 110) .

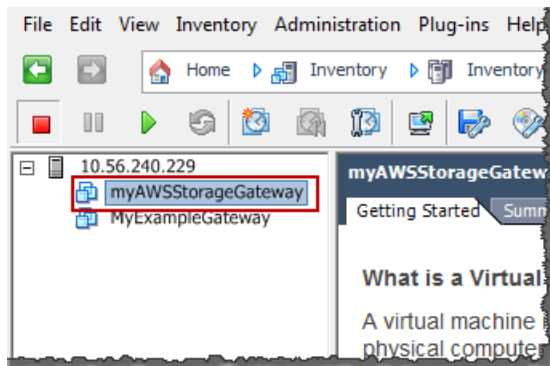
Steps for Activating a Gateway

You can activate the gateway using the AWS Storage Gateway console or the AWS Storage Gateway API (see [ActivateGateway \(p. 307\)](#)). To activate a gateway, you need to know the IP address of the gateway VM. Before starting the activation process, ensure that you have network access to the gateway from the computer that you will use to perform the activation.

The following procedure demonstrates how to activate a gateway using the vSphere client to get the IP address of the gateway VM and then how to use that IP address in the console **Setup and Activate Gateway** wizard.

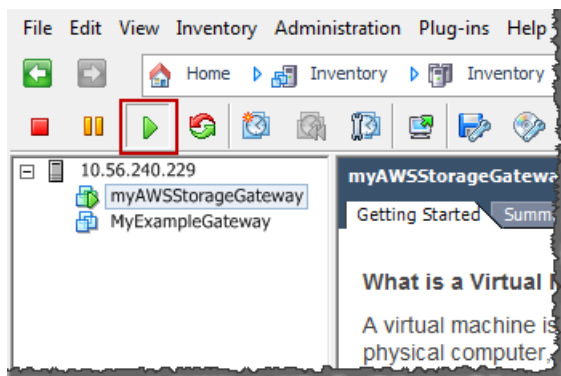
To activate your gateway using the console

1. Power on the VM if it is not already on.
 - a. In the vSphere client, select the gateway VM.



- b. Click the **Power On** icon on the **Toolbar** menu.

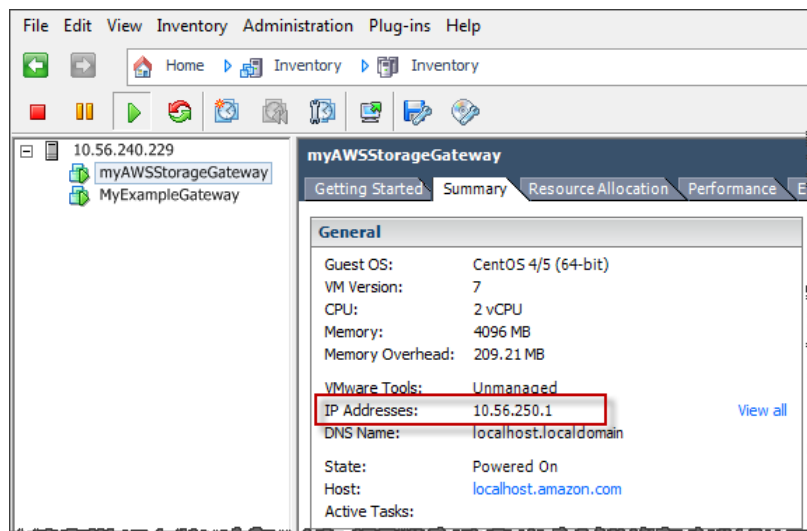
Your gateway VM icon now includes a green arrow icon indicating you have powered on the VM.



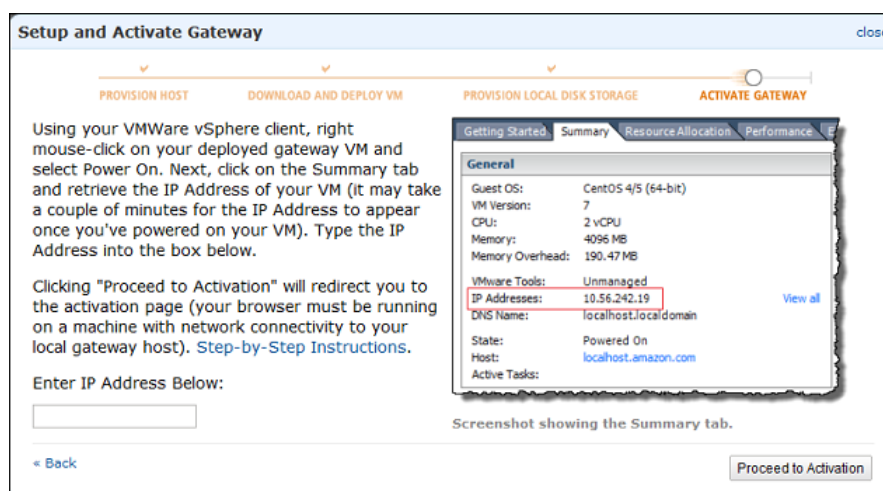
AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a VMware ESXi Host

2. Activate the gateway.
 - a. Obtain the IP address of your gateway. Note that, after powering on the VM, it might take a few minutes for the IP address to appear.
 - i. Using the vSphere client, log in to your host.
 - ii. Select the deployed gateway VM.
 - iii. Click the **Summary** tab for the IP address.



- b. Associate your gateway to your AWS account
 - i. Return to the console, open the **Setup and Activate Gateway** wizard if you haven't already, proceed to the **ACTIVATE GATEWAY** step, enter the IP address and click **Proceed to Activation**. Your browser must be running on a machine with network connectivity to your local gateway host.



Note

If activation fails, check that the IP address you entered is correct and try to activate again. If the IP address is correct, then confirm that the gateway can access the

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a VMware ESXi Host

Internet and, if needed, set up a proxy (see [Routing AWS Storage Gateway Through a Proxy](#) (p. 238)).

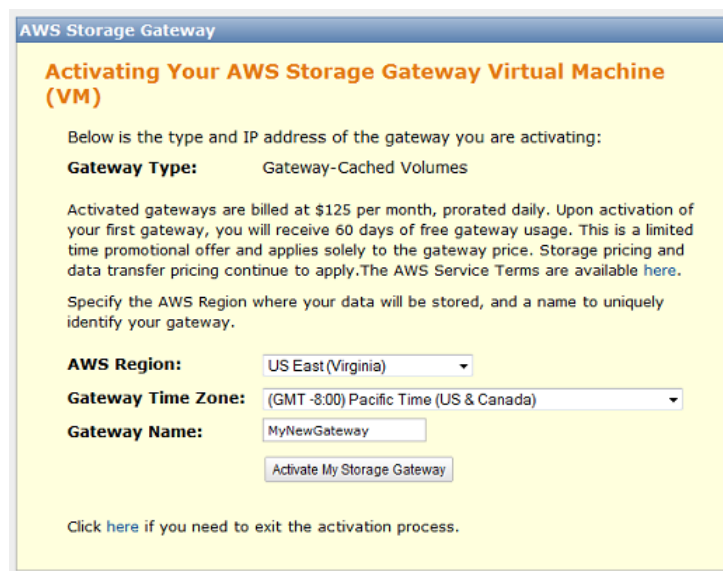
- ii. On the activation page, fill in the requested information to complete the activation process.

The **Gateway Type** specifies what type of gateway you are activating. You can activate a gateway for cached volumes or stored volumes. For more information, see [How AWS Storage Gateway Works](#) (p. 3).

The **AWS Region** determines where AWS stores your snapshots. If you choose to restore a snapshot to an Amazon EBS volume, then the Amazon EBS volume must be in the same region as the snapshot. You cannot change the region after the gateway is activated.

The **Gateway Time Zone** is the time zone used when displaying time-based information such as maintenance messages from AWS and snapshot scheduling. You can change the time zone post-activation.

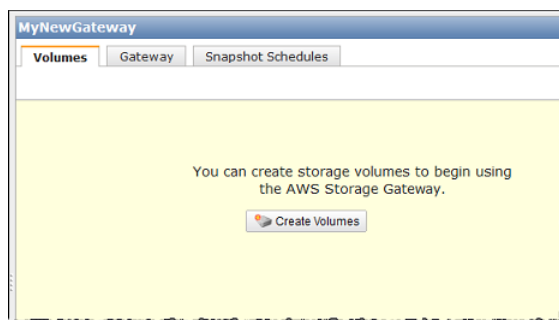
The **Gateway Name** identifies your gateway in the console. You use this name to manage your gateway in the console and you can change it post-activation.



The screenshot shows the 'Activating Your AWS Storage Gateway Virtual Machine (VM)' page. It includes a title bar 'AWS Storage Gateway' and a main heading 'Activating Your AWS Storage Gateway Virtual Machine (VM)'. Below the heading, it states: 'Below is the type and IP address of the gateway you are activating: Gateway Type: Gateway-Cached Volumes'. A paragraph explains billing: 'Activated gateways are billed at \$125 per month, prorated daily. Upon activation of your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies solely to the gateway price. Storage pricing and data transfer pricing continue to apply. The AWS Service Terms are available here.' It then asks to 'Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway.' The form contains three dropdown menus: 'AWS Region' set to 'US East (Virginia)', 'Gateway Time Zone' set to '(GMT -8:00) Pacific Time (US & Canada)', and 'Gateway Name' set to 'MyNewGateway'. There is an 'Activate My Storage Gateway' button and a link to 'Click here if you need to exit the activation process.'

- iii. Click **Activate My Storage Gateway**.

Upon successful activation, the **AWS Storage Gateway** console shows the activated gateway and link for you to create volumes.



Related Section


- [API Reference for AWS Storage Gateway \(p. 283\)](#)

Deploying and Activating AWS Storage Gateway On-Premises on a Microsoft Hyper-V Host

This section explains how you can use Microsoft Hyper-V to create an on-premises virtual machine to host AWS Storage Gateway. The tasks described here assume you have already provisioned a Microsoft Hyper-V host. If you have not already done so, see [Provision a Hyper-V Host to Deploy the AWS Storage Gateway VM \(p. 34\)](#) which is part of the getting started exercise.

Downloading and Deploying AWS Storage Gateway VM

After you provision a Microsoft Hyper-V host, the next steps in the **Setup and Activate Gateway** wizard are to select the Microsoft Hyper-V platform, download the VM software for this platform, and then deploy the VM.



For instructions, see [Download and Deploy the AWS Storage Gateway VM on Your Host \(p. 34\)](#).

Provisioning Local Disk Storage for an AWS Storage Gateway VM

Topics

- [About the Disk the Gateway VM Uses to Store System Data \(p. 116\)](#)
- [Provisioning Local Disks \(Gateway-Cached\) \(p. 117\)](#)
- [Provisioning Local Disks \(Gateway-Stored\) \(p. 124\)](#)

Before you provision local disk storage for the gateway VM you deployed, you should decide the type of iSCSI storage volumes you plan to use. You have the following options:

- **Use Gateway-Cached volumes** – In this case, the gateway stores your volume data in Amazon S3.

In this case, the gateway maintains a cache storage for recently accessed data to provide low-latency access. The gateway persistently holds the data that has not been uploaded to Amazon S3 in the cache storage; therefore, you must allocate disks on-premises for the cache storage. You must also allocate disks for the upload buffer to temporarily buffer your data prior to uploading to AWS. The cache storage should be larger than the upload buffer (see [How AWS Storage Gateway Works \(p. 3\)](#)).

- **Use Gateway-Stored volumes** – In this case, the gateway stores your volume data on your on-premises storage hardware.

You must allocate disks on-premises to hold all your data. The gateway then securely uploads data snapshots to Amazon S3 for cost-effective backup and rapid disaster recovery. You must also allocate disks for the gateway's upload buffer (see [How AWS Storage Gateway Works \(p. 3\)](#)).

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a Microsoft Hyper-V Host

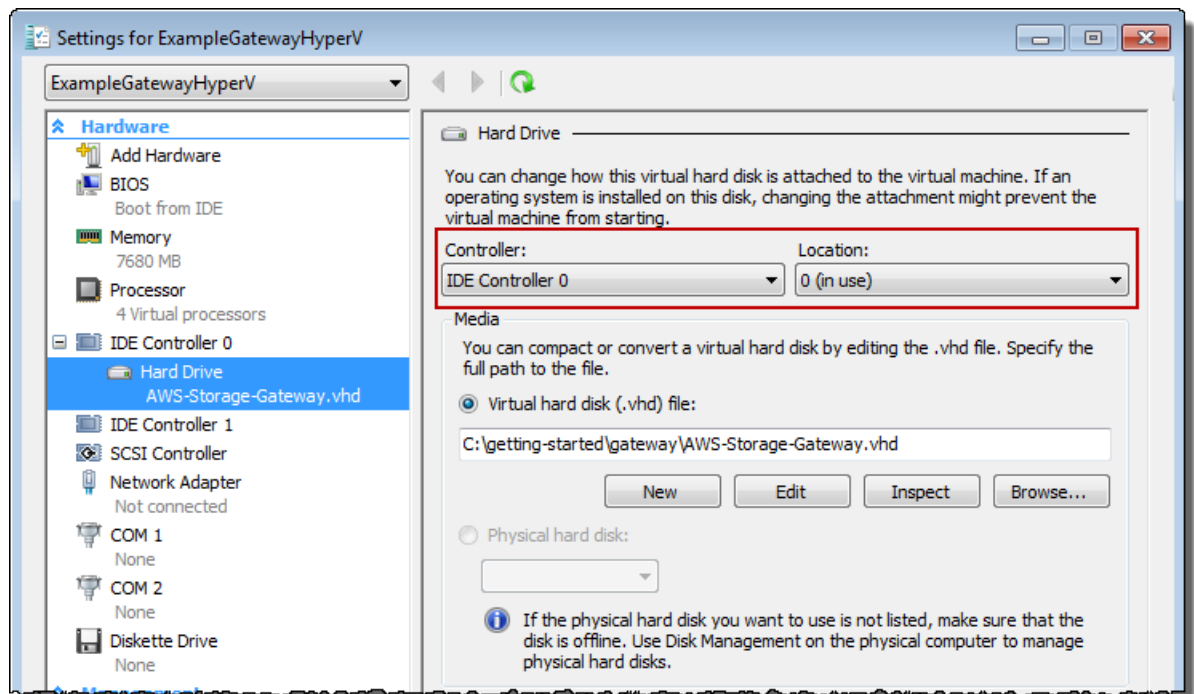
If you follow the **Setup and Activate Gateway** console wizard, the console shows the following prompt for you to choose the volume type.



After selecting the volume type, you must provision local disks to the gateway VM required to support the volume type you selected before activating the gateway.

About the Disk the Gateway VM Uses to Store System Data

After you deploy (import) the gateway VM, it includes preconfigured processors, memory, and an IDE disk with the VM infrastructure on it. This IDE disk appears as **AWS-Storage-Gateway.vhd** and controller IDE (0:0) in the **Settings** window in the Microsoft Hyper-V Manager, as shown in the following example screen shot. However, you cannot access or use this disk directly. The gateway uses it to store system data.



Provisioning Local Disks (Gateway-Cached)

Topics

- [Adding Local Disks for Cache Storage \(Gateway-Cached\) \(p. 117\)](#)
- [Adding Local Disks for the Upload Buffer \(Gateway-Cached\) \(p. 120\)](#)

In the gateway-cached architecture, the gateway stores your volume data in Amazon S3. However, you must provision disks to the gateway VM for cache storage and the upload buffer. For more information about how the gateway works, see [How AWS Storage Gateway Works \(p. 3\)](#).

Note

When you provision disks, it is strongly recommended that you do not provision local disks for upload buffer and cache storage that use the same underlying physical storage resource (disk). When you deploy the gateway VM, you choose a disk location to store the VM files. When you provision a local disk (e.g., to use as cache storage or upload buffer), you have the option to store the virtual disk in the same disk location as the VM or in a different location. If you have more than one disk, we strongly recommend that you choose one disk location for the cache storage and another for the upload buffer. One disk location that is backed by only one underlying physical disk, or that is backed by a less-performant RAID configuration such as RAID 1, may lead to poor performance in some situations when used to back both the cache storage and upload buffer.


Adding Local Disks for Cache Storage (Gateway-Cached)

Topics

- [Sizing Cache Storage \(Gateway-Cached\) \(p. 117\)](#)
- [Adding a Virtual Disk for Cache Storage \(Gateway-Cached\) \(p. 118\)](#)

In the gateway-cached architecture, your gateway maintains cache storage on-premises for recently accessed data. The gateway persistently holds the data in cache storage that has not be uploaded to Amazon S3. You will need to allocate disks on-premises for cache storage.

The following diagram highlights cache storage in the larger picture of the AWS Storage Gateway architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).



Sizing Cache Storage (Gateway-Cached)

The gateway uses the cache storage to provide low-latency access to your recently accessed data. The cache storage acts as the on-premises durable store for data that is pending upload to Amazon S3 from the upload buffer. So cache storage should be larger than the upload buffer.

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a Microsoft Hyper-V Host

The total cache storage for a gateway can be up to 16 TiB.

To estimate the amount of cache storage your gateway needs, the formula depends on your use-case:

- **Backup Use Case**—Use a cache storage capacity of 1.1 times the upload buffer capacity. For a backup use-case, the cache is durable storage that holds data prior to upload to AWS, and it must be sized greater than the upload buffer to ensure that no data is lost in the event of a VM failure.
- **Other Use Cases**—Use the larger of the following two values: 20 percent of your existing on-premise storage or 1.1 times the upload buffer size.

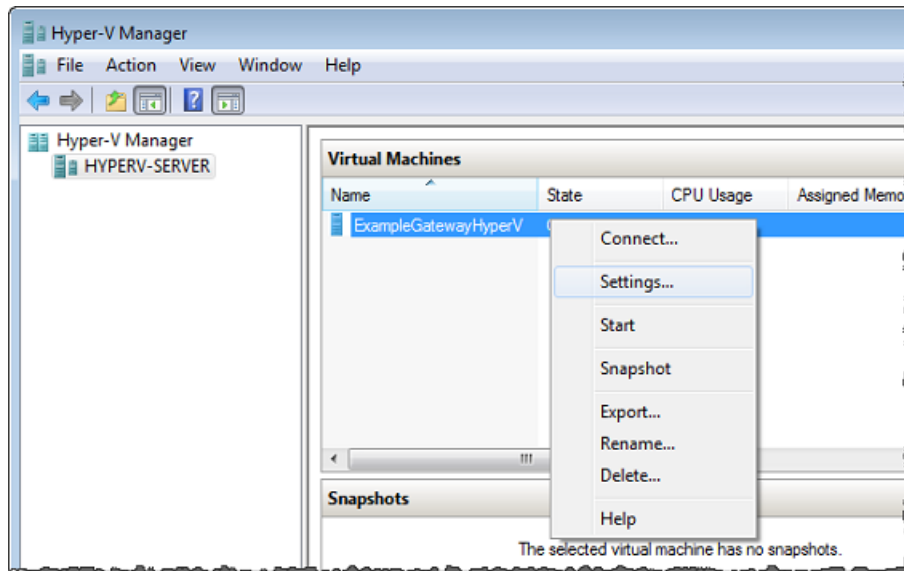
You can initially use this approximation to provision disks for the cache storage. You can then use Amazon CloudWatch operational metrics to monitor the cache storage usage and provision more storage as needed using the console. For using the metrics and setting up alarms, see [Monitoring Cache Storage](#) (p. 271).

Adding a Virtual Disk for Cache Storage (Gateway-Cached)

You can allocate virtual disks to the VM from either the direct-attached storage (DAS) disks or from the storage area network (SAN) disks available on your host. The following procedure provides instructions for adding a virtual disk from a DAS disk that is available on the host.

To allocate a new virtual disk to the VM for cache storage

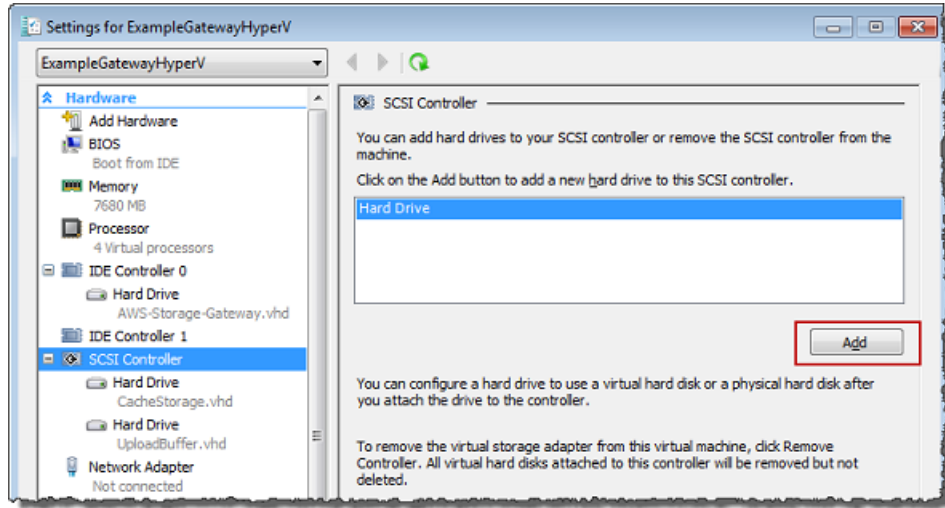
1. Start the Microsoft Hyper-V Manager and connect to your host.
2. In the client, right-click the name of your gateway VM and click **Settings...**



3. In the **Hardware** list in the left pane, click **SCSI Controller**.
4. In the **SCSI Controller** pane, Click **Add**.

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a Microsoft Hyper-V Host



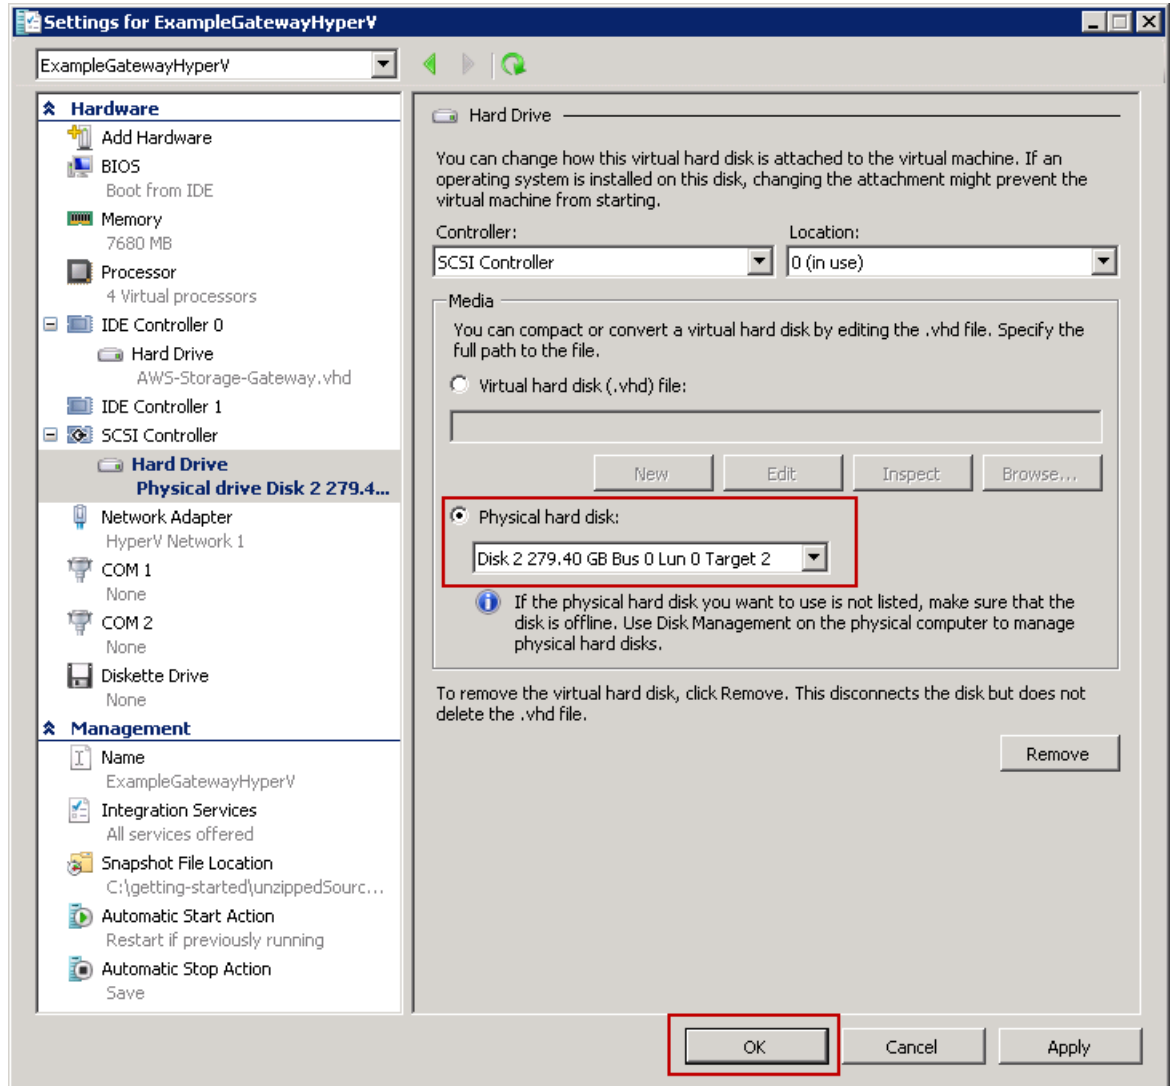
5. In the **Hard Drive** pane, in the **Media** section click **Physical hard disk** and select a disk from the box.

Note

For an example of using a virtual hard disk, see [Allocate a Local Disk for Cache Storage](#) (p. 44) in the getting started exercise.

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a Microsoft Hyper-V Host



6. Click **OK**.

Adding Local Disks for the Upload Buffer (Gateway-Cached)

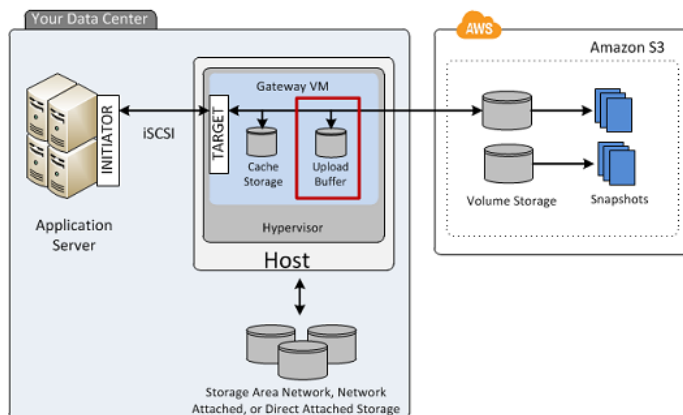
Topics

- [Sizing the Upload Buffer \(Gateway-Cached\) \(p. 121\)](#)
- [Adding a Virtual Disk for the Upload Buffer \(Gateway-Cached\) \(p. 122\)](#)

You must allocate disk(s) on your premises for the gateway to use as the upload buffer to temporarily buffer your data prior to uploading to AWS.

The following diagram highlights the upload buffer in the larger picture of the AWS Storage Gateway architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).

AWS Storage Gateway User Guide Deploying and Activating a Gateway on a Microsoft Hyper-V Host



Sizing the Upload Buffer (Gateway-Cached)

You can determine the size of your upload buffer by using an upload buffer formula. We strongly recommend that you allocate at least 150 GiB of upload buffer space. Therefore, if the formula returns a value less than 150 GiB, use 150 GiB as the amount you allocate to the upload buffer. You can configure up to 2 TiB of upload buffer capacity per gateway.

Note

When upload buffer reaches its capacity, your applications can continue to read from and write data to your storage volumes; however, the gateway will not writing any of your volume data to its upload buffer and will not upload any of this data to AWS.

To estimate the amount of upload buffer more precisely, you can calculate the incoming and outgoing data rates and base an estimate on these rates.

- **Rate of Incoming Data**—This refers to the application throughput, the rate at which your on-premises applications are writing data to your gateway over some period of time.
- **Rate of Outgoing Data**—This refers to the network throughput, the rate at which your gateway is able to upload data to AWS. This depends on your network speed, utilization, and whether you've enabled bandwidth throttling. This rate should be adjusted for compression. When uploading data to AWS, the gateway applies data compression where possible. For example, if your application data is text-only, you might get effective compression ratio of about 2:1. However, if you are writing videos, the gateway might not be able to achieve any data compression, requiring more upload buffer for the gateway.

If your incoming rate is higher than the outgoing rate, you can use the following formula to determine the approximate size of the upload buffer your gateway needs.

$$\left(\begin{array}{c} \text{Application} \\ \text{Throughput} \\ \text{(MB/s)} \end{array} - \begin{array}{c} \text{Network} \\ \text{Throughput} \\ \text{to AWS (MB/s)} \end{array} \times \begin{array}{c} \text{Compression} \\ \text{Factor} \end{array} \right) \times \begin{array}{c} \text{Duration} \\ \text{of writes} \\ \text{(s)} \end{array} = \begin{array}{c} \text{Upload} \\ \text{Buffer} \\ \text{(MB)} \end{array}$$

For example, assume that your business applications will write text data to your gateway at a rate of 40 megabytes per second for 12 hours a day and your network throughput is 12 megabytes per second. Assuming a compression factor of 2:1 for the text data, you need to allocate approximately 690 GB for the upload buffer.

$$((40 \text{ MB/sec}) - (12 \text{ MB/sec} * 2)) * (12 \text{ hours} * 3600 \text{ seconds/hour}) = 691200 \text{ megabytes}$$

Note that you can initially use this approximation to determine the disk size that you want to allocate to the gateway as upload buffer space. Add more upload buffer space as needed using the AWS Storage Gateway console. Also, you can use the Amazon CloudWatch operational metrics to monitor upload buffer usage and determine additional storage requirements. For information on metrics and setting the alarms, see [Monitoring the Upload Buffer \(p. 267\)](#).

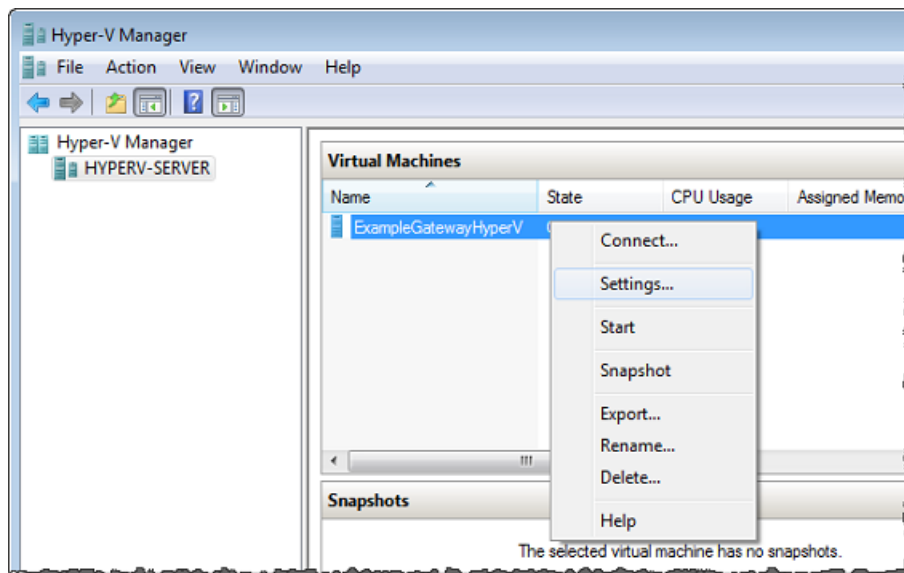
Adding a Virtual Disk for the Upload Buffer (Gateway-Cached)

In this section, you allocate a virtual disk to your VM that will be used as the upload buffer for your gateway.

You can allocate virtual disks to the VM from either the direct-attached storage (DAS) disks or from the storage area network (SAN) disks available on your host. The following procedure provides step-by-step instructions to add a virtual disk from a DAS disk available on the host.

To allocate a new virtual disk to the VM for the upload buffer

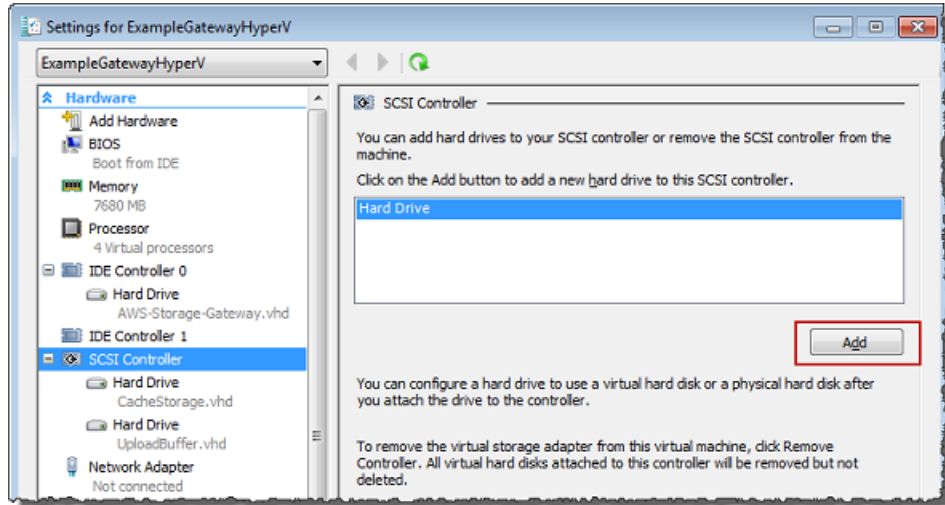
1. Start the Microsoft Hyper-V Manager and connect to your host.
2. In the client, right-click the name of your gateway VM and click **Settings....**



3. In the **Hardware** list in the left pane, click **SCSI Controller**.
4. In the **SCSI Controller** pane, Click **Add**.

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a Microsoft Hyper-V Host




5. In the **Hard Drive** pane, in the **Media** section click **Physical hard disk** and select a disk from the box.

Note

For an example of using a virtual hard disk, see [Allocate a Local Disk for Cache Storage](#) (p. 44) in the getting started exercise.

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a Microsoft Hyper-V Host



6. Click **OK**.

Provisioning Local Disks (Gateway-Stored)

Topics

- [Adding Local Disks for Volume Storage \(Gateway-Stored\) \(p. 125\)](#)
- [Adding Local Disks for Upload Buffer \(Gateway-Stored\) \(p. 127\)](#)

In the gateway-stored architecture, the gateway stores your volume data on your on-premises storage hardware. All your application data reside on your premises. You must provision disks to the gateway VM for the volume storage. You must also provision disks for the gateway's upload buffer. For more information about how the gateway works, see [How AWS Storage Gateway Works \(p. 3\)](#).

Note

When you provision disks, it is strongly recommended that you do not provision local disks for upload buffer and local application storage that use the same underlying physical storage resource (disk). In Microsoft Hyper-V, when you provision a local disk for gateway (e.g., to use it as local application storage or upload buffer), you can create it as a virtual hard disk (.vhd) file or from

AWS Storage Gateway User Guide

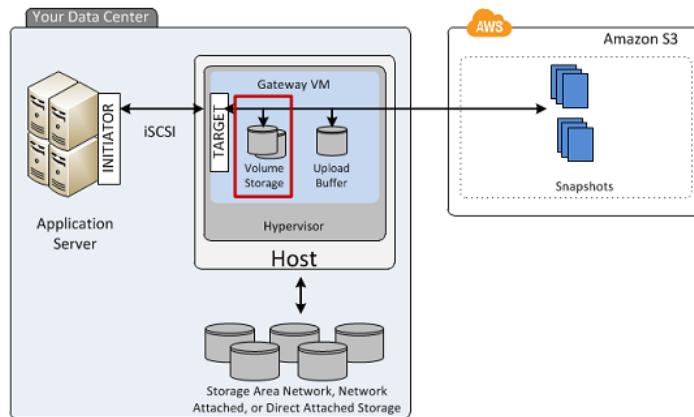
Deploying and Activating a Gateway on a Microsoft Hyper-V Host

a physical hard disk. Whether you choose a .vhd or a physical disk depends on your performance and portability requirements. Provisioning a local disk based on a physical disk is less portable than using a .vhd.

Adding Local Disks for Volume Storage (Gateway-Stored)

In the gateway-stored architecture, your application data is stored locally. You will need to provision disks to the gateway VM to store your data.

The following diagram highlights storage volumes in the larger picture of the AWS Storage Gateway architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).



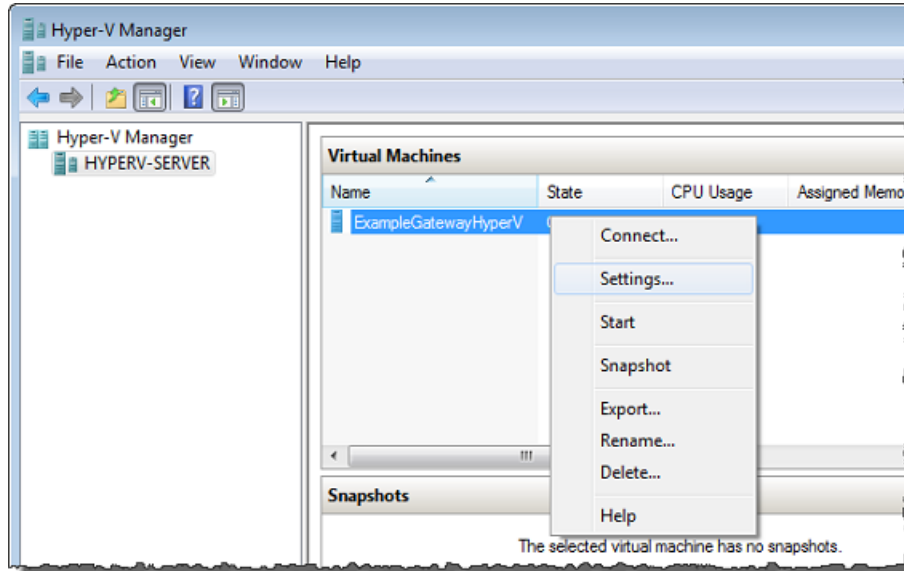
Each disk can be up to 1 TiB in size and must be rounded to the nearest GiB, where GiB is calculated using Base 2 (i.e., GiB = 1024³ bytes).

You can provision disks to the VM for volume storage from either the direct-attached storage (DAS) disks or from the storage area network (SAN) disks. For volume storage, the disk you allocate can have existing data. We preserve this data when creating your iSCSI storage volumes. The following procedure provides instructions for adding a virtual disk from a DAS disk.

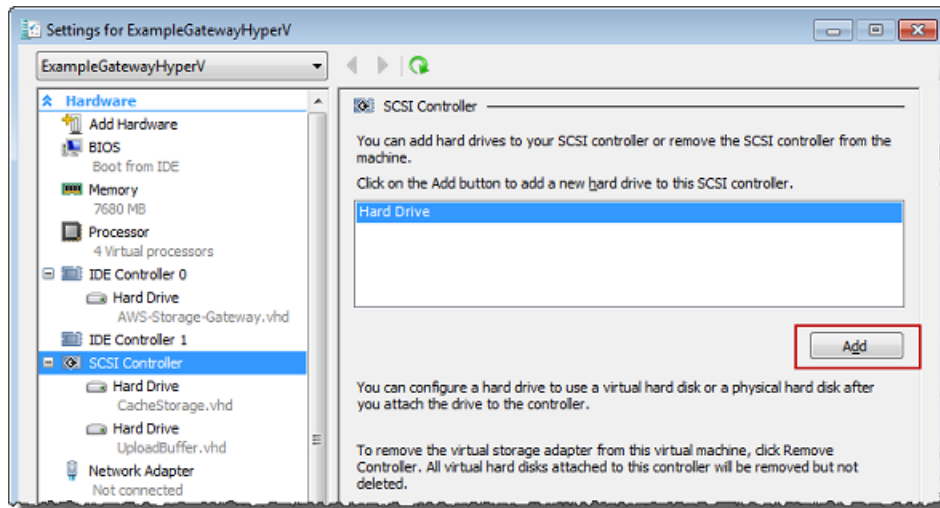
To allocate a new virtual disk to the VM for application data

1. Start the Microsoft Hyper-V Manager and connect to your host.
2. In the client, right-click the name of your gateway VM and click **Settings**....

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a Microsoft
Hyper-V Host



3. In the **Hardware** list in the left pane, click **SCSI Controller**.
4. In the **SCSI Controller** pane, Click **Add**.




5. In the **Hard Drive** pane, in the **Media** section click **Physical hard disk** and select a disk from the box.

Note

For an example of using a virtual hard disk, see [Allocate a Local Disk for Cache Storage](#) (p. 44) in the getting started exercise.

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a Microsoft Hyper-V Host



6. Click **OK**.

Adding Local Disks for Upload Buffer (Gateway-Stored)

Topics

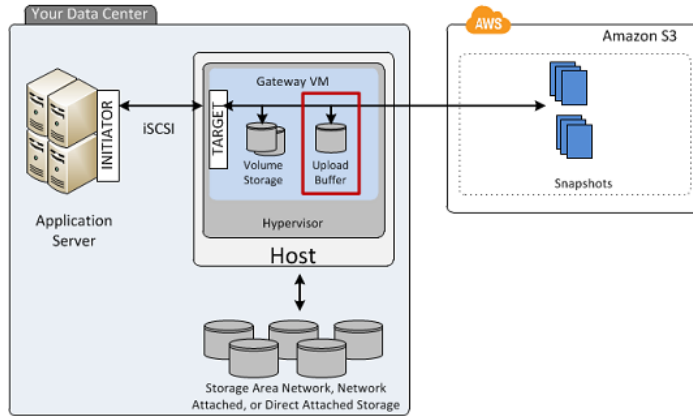
- [Sizing the Upload Buffer \(Gateway-Stored\) \(p. 128\)](#)
- [Adding a Virtual Disk for the Upload Buffer \(Gateway-Stored\) \(p. 129\)](#)

You must allocate disk(s) on your premises for the gateway to use as the upload buffer to temporarily buffer your data prior to uploading to AWS.

The following diagram highlights the upload buffer in the larger picture of the AWS Storage Gateway architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a Microsoft Hyper-V Host



Sizing the Upload Buffer (Gateway-Stored)

You can determine the size of your upload buffer by using an upload buffer formula. We strongly recommend that you allocate at least 150 GiB of upload buffer space. Therefore, if the formula returns a value less than 150 GiB, use 150 GiB as the amount you allocate to the upload buffer. You can configure up to 2 TiB of upload buffer capacity per gateway.

Note

When the upload buffer reaches its capacity, your applications can continue to read from and write data to your storage volumes; however, the gateway is not writing any of your volume data to its upload buffer and not uploading any of this data to AWS.

To estimate the amount of upload buffer space, calculate the incoming and outgoing data rates and base an estimate on these rates.

- **Rate of Incoming Data**—This refers to the application throughput, the rate at which your on-premises applications are writing data to your gateway over some period of time.
- **Rate of Outgoing Data**—This refers to the network throughput, the rate at which your gateway is able to upload data to AWS. This depends on your network speed, utilization, and whether you've enabled bandwidth throttling. This rate should be adjusted for compression. When uploading data to AWS, the gateway applies data compression where possible. For example, if your application data is text-only, you might get effective compression ratio of about 2:1. However, if you are writing videos, the gateway might not be able to achieve any data compression, requiring more upload buffer space for the gateway.

If your incoming rate is higher than the outgoing rate, you can use the following formula to determine the approximate size of the upload buffer your gateway needs.

$$\left(\begin{array}{c} \text{Application} \\ \text{Throughput} \\ \text{(MB/s)} \end{array} - \begin{array}{c} \text{Network} \\ \text{Throughput} \\ \text{to AWS (MB/s)} \end{array} \right) \times \begin{array}{c} \text{Compression} \\ \text{Factor} \end{array} \times \begin{array}{c} \text{Duration} \\ \text{of writes} \\ \text{(s)} \end{array} = \begin{array}{c} \text{Upload} \\ \text{Buffer} \\ \text{(MB)} \end{array}$$

For example, assume that your business applications will write text data to your gateway at a rate of 40 megabytes per second for 12 hours a day and your network throughput is 12 megabytes per second. Assuming a compression factor of 2:1 for the text data, you need to allocate approximately 690 GB for the upload buffer.

$$((40 \text{ MB/sec}) - (12 \text{ MB/sec} * 2)) * (12 \text{ hours} * 3600 \text{ seconds/hour}) = 691200 \text{ megabytes}$$

Note that you can initially use this approximation to determine the disk size that you want to allocate to the gateway as upload buffer space. Add more upload buffer space as needed using the AWS Storage Gateway console. Also, you can use the Amazon CloudWatch operational metrics to monitor upload buffer usage and determine additional storage requirements. For information on metrics and setting the alarms, see [Monitoring the Upload Buffer \(p. 267\)](#).

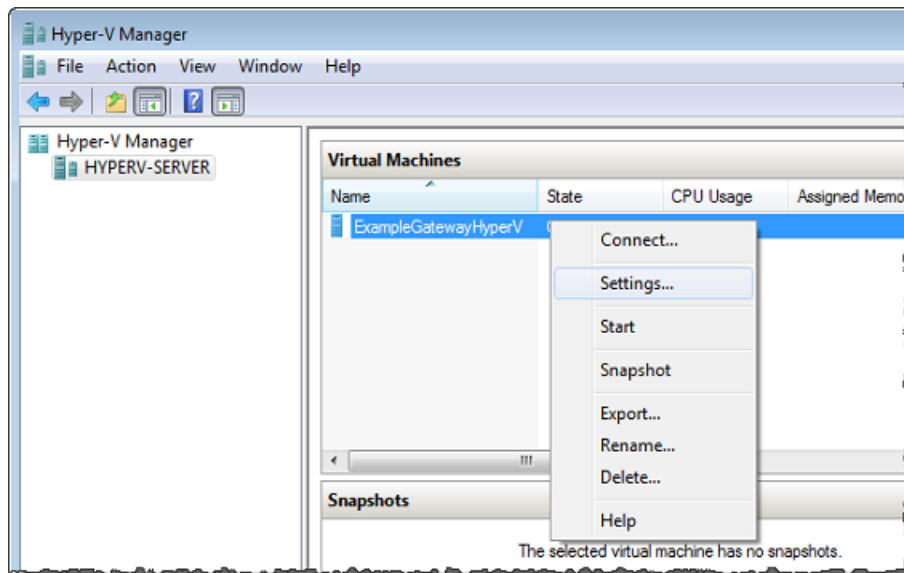
Adding a Virtual Disk for the Upload Buffer (Gateway-Stored)

In this section, you allocate a virtual disk to your VM that will be used as the upload buffer for your gateway. To estimate the upload buffer your gateway requires, see [Sizing the Upload Buffer \(Gateway-Stored\) \(p. 106\)](#).

You can allocate virtual disks to the VM from either the direct-attached storage (DAS) disks or from the storage area network (SAN) disks available on your host. The following procedure provides step-by-step instructions to add a virtual disk from a DAS disk available on the host.

To allocate a new virtual disk to the VM for the upload buffer


1. Start the Microsoft Hyper-V Manager and connect to your host.
2. In the client, right-click the name of your gateway VM and click **Settings....**



3. In the **Hardware** list in the left pane, click **SCSI Controller**.
4. In the **SCSI Controller** pane, Click **Add**.

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a Microsoft Hyper-V Host




5. In the **Hard Drive** pane, in the **Media** section click **Physical hard disk** and select a disk from the box.

Note

For an example of using a virtual hard disk, see [Allocate a Local Disk for Cache Storage](#) (p. 44) in the getting started exercise.

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a Microsoft Hyper-V Host



6. Click **OK**.

Activating AWS Storage Gateway

After you deploy the AWS Storage Gateway VM, you must activate the gateway using the AWS Storage Gateway console. The activation process associates your gateway with your AWS account. Once you establish this connection, you can manage almost all aspects of your gateway from the console. In the activation process, you specify the IP address of your gateway, name your gateway, identify the AWS region in which you want your snapshot backups stored, and specify the gateway timezone. After this activation, you begin incurring charges. For information about pricing, see [AWS Storage Gateway](#).

Pre-Activation Checklist

You can activate a gateway after you have completed the steps summarized in the following table. The console wizard walks you through these steps.

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a Microsoft
Hyper-V Host

Step	Description
Download and Deploy the VM	In the AWS Storage Gateway console, download the latest virtual machine (VM) that is distributed as an .zip file and deploy (import) this VM on your Microsoft Hyper-V host. For more information, see Downloading and Deploying AWS Storage Gateway VM (p. 115) .
Provision local disks to the VM	The provisioned VM has no disks. Depending on the gateway type you activated, you either must add local disks for: <ul style="list-style-type: none">• cache storage and upload buffer for a gateway activated for cached volumes. For more information, see Provisioning Local Disks (Gateway-Cached) (p. 117).• application data and upload buffer for a gateway activated for stored volumes. For more information, see Provisioning Local Disks (Gateway-Stored) (p. 124).

Steps for Activating a Gateway

You can activate the gateway using the AWS Storage Gateway console or the AWS Storage Gateway API (see [ActivateGateway \(p. 307\)](#)). To activate a gateway, you need to know the IP address of the gateway VM. Before starting the activation process, ensure that you have network access to the gateway from the computer that you will use to perform the activation.


The following procedure demonstrates how to activate a gateway using the Microsoft Hyper-V Manager to get the IP address of the gateway VM and then how to use that IP address in the console **Setup and Activate Gateway** wizard.

To activate your gateway using the console


1. Power on the VM if it is not already on.
 - a. Start the Microsoft Hyper-V Manager and connect to the hypervisor.
 - b. In the **Virtual Machines** list pane, select the virtual machine **ExampleGatewayHyperV**.
 - c. In the **Actions** pane, select **Start**.

The **Virtual Machine Connection** window appears.

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a Microsoft
Hyper-V Host



- d. If an authentication window appears, enter the user name and password provided to you by the hypervisor administrator.



- e. After a few moments, the virtual machine is ready for you to log in.

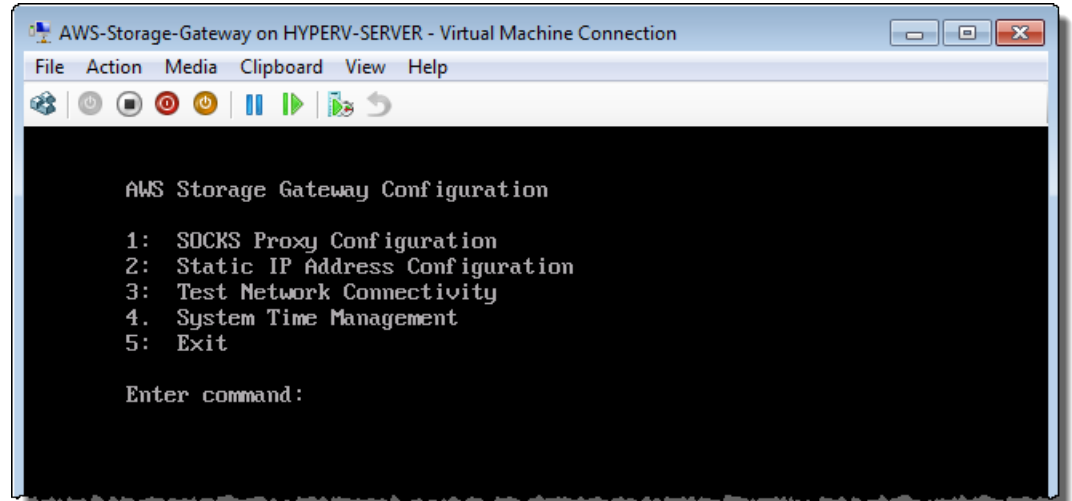
2. Activate the gateway.

- a. Obtain the IP address of your gateway. Note that, after powering on the VM, it might take a few minutes for the gateway to be ready for you to log in and get the IP address.
- i. In the Microsoft Hyper-V Manager, select the deployed gateway VM.
 - ii. In the **Virtual Machines** list pane, select the virtual machine **ExampleGatewayHyperV**.
 - iii. In the **Actions** pane, select **Connect**.

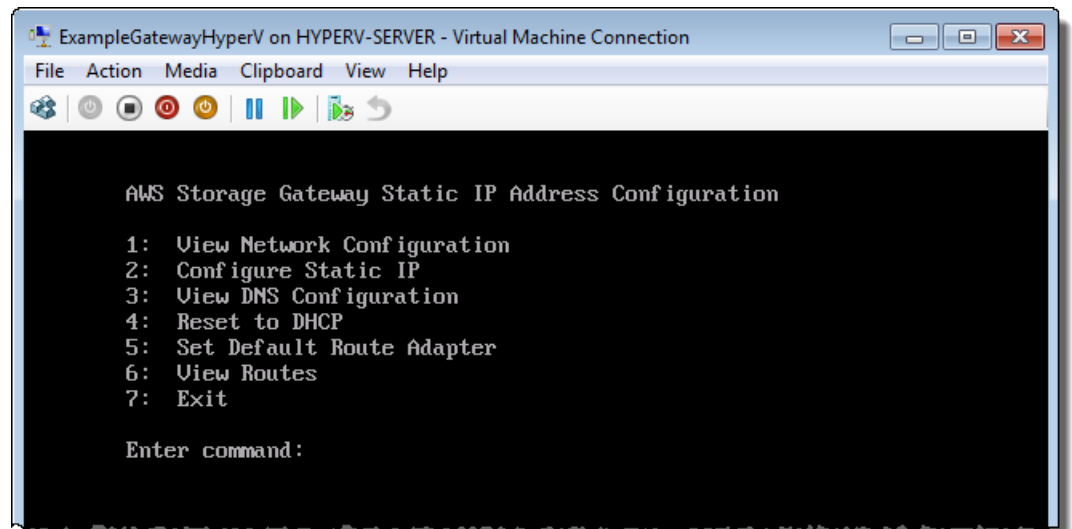
The **Virtual Machine Connection** window appears.

- iv. At the login prompt, enter the user name `sguser`, and the password `sgpassword`.
- v. In the **AWS Storage Gateway Configuration** menu, select option 2, **Static IP Address Configuration**.

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a Microsoft
Hyper-V Host



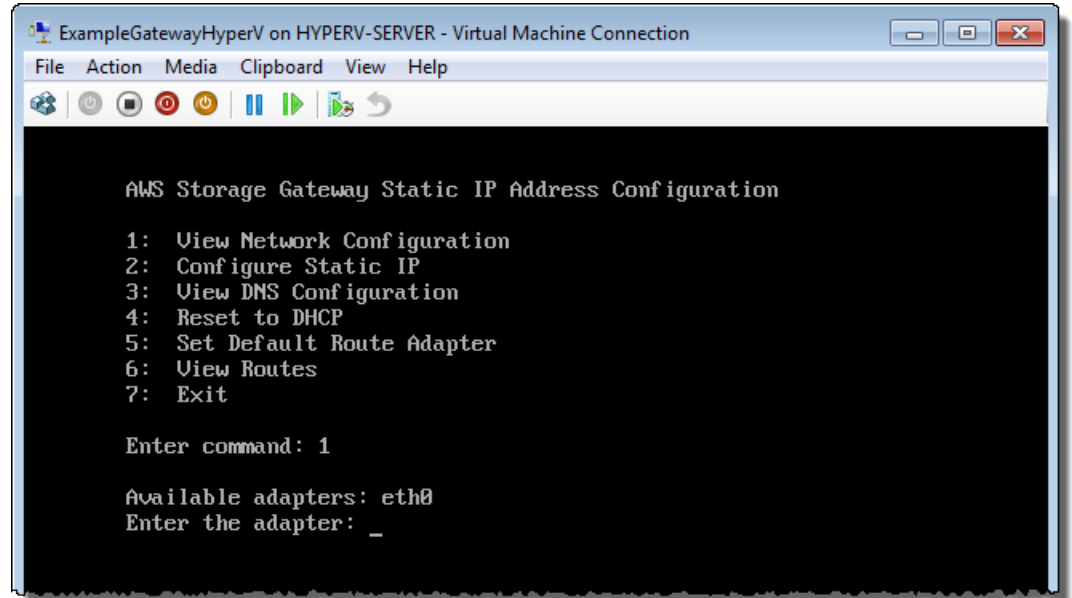
- vi. In the **AWS Storage Gateway Static IP Address Configuration** menu, select option 1, **View Network Configuration**.



- vii. Type the identifier of the adapter.

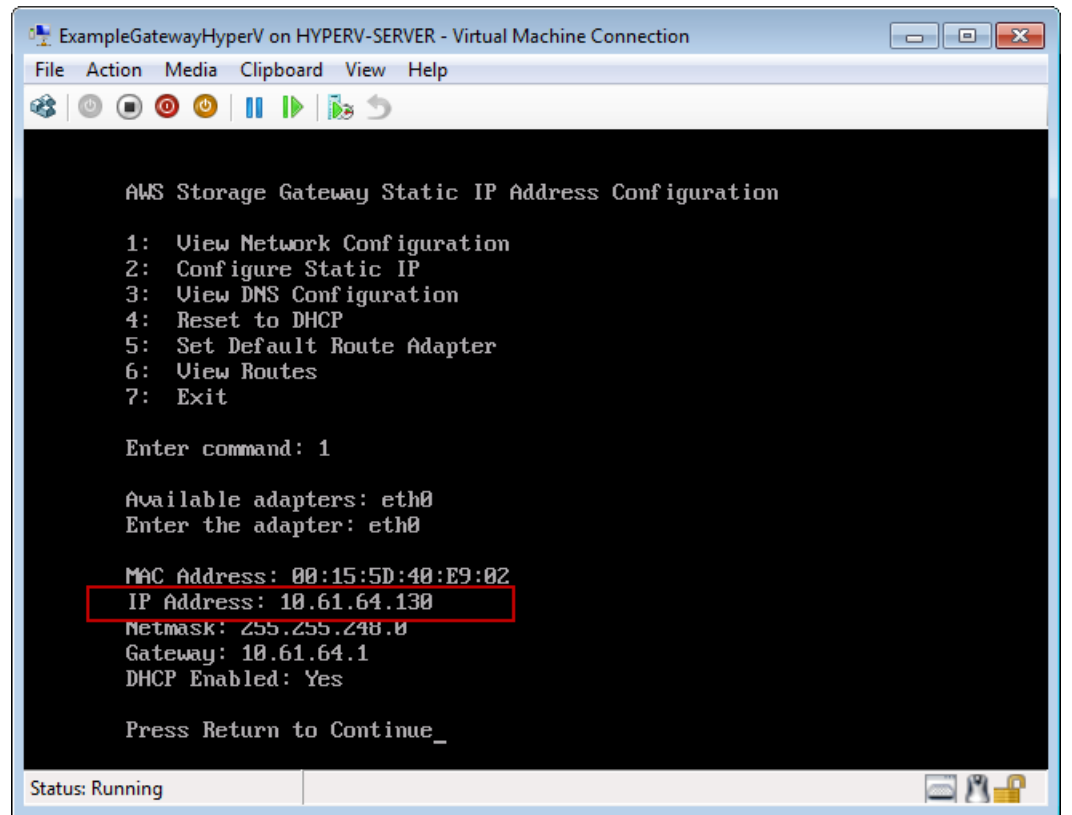
In most scenarios, `eth0` will be the adapter identifier.

AWS Storage Gateway User Guide
Deploying and Activating a Gateway on a Microsoft
Hyper-V Host



viii. Get the IP address from the adapter information.

In the example below, the IP address is 10.61.64.130. Your gateway's IP address will be different.

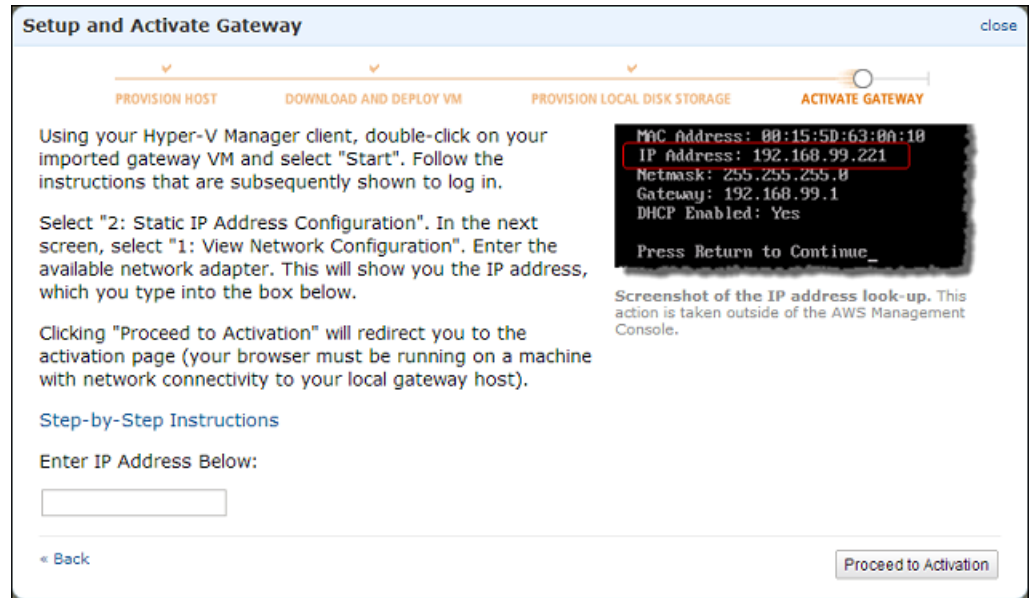


ix. Press **Return**, and follow the prompts to exit the configuration menu.

AWS Storage Gateway User Guide

Deploying and Activating a Gateway on a Microsoft Hyper-V Host

- b. Associate your gateway to your AWS account
 - i. Return to the console, open the **Setup and Activate Gateway** wizard if you haven't already, proceed to the **ACTIVATE GATEWAY** step, enter the IP address and click **Proceed to Activation**. Your browser must be running on a machine with network connectivity to your local gateway host.



Note

If activation fails, check that the IP address you entered is correct and try to activate again. If the IP address is correct, then confirm that the gateway can access the Internet and, if needed, set up a proxy (see [Routing AWS Storage Gateway Through a Proxy \(p. 238\)](#)).

- ii. On the activation page, fill in the requested information to complete the activation process.

The **Gateway Type** specifies what type of gateway you are activating. You can activate a gateway for cached volumes or stored volumes. For more information, see [How AWS Storage Gateway Works \(p. 3\)](#).

The **AWS Region** determines where AWS stores your snapshots. If you choose to restore a snapshot to an Amazon EBS volume, then the Amazon EBS volume must be in the same region as the snapshot. You cannot change the region after the gateway is activated.

The **Gateway Time Zone** is the time zone used when displaying time-based information such as maintenance messages from AWS and snapshot scheduling. You can change the time zone post-activation.

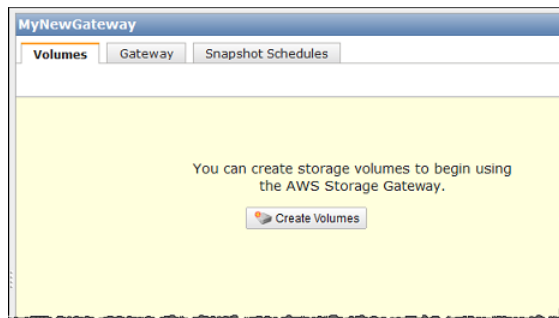
The **Gateway Name** identifies your gateway in the console. You use this name to manage your gateway in the console and you can change it post-activation.

The screenshot shows the 'Activating Your AWS Storage Gateway Virtual Machine (VM)' page. It includes the following information and form fields:

- Gateway Type:** Gateway-Cached Volumes
- Text:** Activated gateways are billed at \$125 per month, prorated daily. Upon activation of your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies solely to the gateway price. Storage pricing and data transfer pricing continue to apply. The AWS Service Terms are available [here](#).
- Text:** Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway.
- Form Fields:**
 - AWS Region:** US East (Virginia)
 - Gateway Time Zone:** (GMT -8:00) Pacific Time (US & Canada)
 - Gateway Name:** MyNewGateway
- Button:** Activate My Storage Gateway
- Text:** Click [here](#) if you need to exit the activation process.

- iii. Click **Activate My Storage Gateway**.

Upon successful activation, the **AWS Storage Gateway** console shows the activated gateway and link for you to create volumes.



Related Section

- [API Reference for AWS Storage Gateway \(p. 283\)](#)

Deploying and Activating AWS Storage Gateway on Amazon EC2

Topics

- [Comparison of an Amazon EC2 Gateway with an On-Premises Gateway \(p. 138\)](#)
- [Launching and Activating an Amazon EC2 Gateway AMI \(p. 139\)](#)
- [Managing Your Amazon EC2 Gateway \(p. 146\)](#)

AWS Storage Gateway User Guide

Comparison of an Amazon EC2 Gateway with an On-Premises Gateway

In this section, we discuss how to deploy a gateway on Amazon Elastic Compute Cloud (Amazon EC2). After you deploy the gateway and add disk storage, you activate and work with the gateway just as you would for an on-premises gateway.

The AWS Storage Gateway for Amazon EC2 gateway is an Amazon Machine Image (AMI) from which you can create an instance. The AMI is available in AWS Marketplace as [AWS Storage Gateway](#) or can be accessed from the AWS Storage Gateway console. This section describes how to use the AMI to create an Amazon EC2–hosted gateway. For detailed information about working with AMIs, go to [Amazon Machine Images \(AMI\)](#) in the *Amazon EC2 User Guide*.

A gateway hosted in an Amazon EC2 instance is recommended for the gateway cached-volume architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)). Gateways deployed on Amazon EC2 can support the creation of volumes up to 32 TiB. To get started with creating an Amazon EC2–hosted gateway, go to [Launching and Activating an Amazon EC2 Gateway AMI \(p. 139\)](#).

Comparison of an Amazon EC2 Gateway with an On-Premises Gateway

There are a few differences between a gateway hosted in an Amazon EC2 instance in the cloud and a gateway hosted in a virtualization environment like VMware ESXi or Microsoft Hyper-V that is hosted on-premises. The following table summarizes these differences. Outside of these differences, you work with an Amazon EC2–hosted gateway exactly the same way you do with an on-premises gateway in terms of creating storage volumes (see [Creating a Storage Volume \(Gateway-Cached\) \(p. 157\)](#)) and exposing these volumes as iSCSI targets so that your client applications can connect to them (see [Configuring Your Application Access to Storage Volumes \(p. 161\)](#)).

Functionality	On-Premises Gateway	EC2 Gateway
Gateway availability	The on-premises gateway is available as an OVA file that can be downloaded from the AWS Storage Gateway console. For more information about deploying the OVA file, see Downloading and Deploying AWS Storage Gateway VM (p. 91) .	The Amazon EC2 gateway is available as an AMI from which you create an EC2 instance. The AMI is available in the AWS Marketplace or you can click the deployment link in the AWS Storage Gateway console, which will take you to AWS Marketplace. For more information about deploying an Amazon EC2 AMI as a gateway, see Launching and Activating an Amazon EC2 Gateway AMI (p. 139) .
Gateway architecture	An on-premises gateway is recommended for both stored and cached gateways. For more information about gateway types, see How AWS Storage Gateway Works (p. 3) .	A gateway hosted in an Amazon EC2 instance is recommended for cached gateways. You can configure storage volumes up to 32 TiB for an Amazon EC2–hosted gateway.

Functionality	On-Premises Gateway	EC2 Gateway
Logging on to the gateway	You can log in to your on-premises gateway to perform maintenance tasks such as routing your gateway through a proxy, configuring your gateway to use a static IP address, and testing your gateway's connection to the Internet. For more information, see Logging Into Your AWS Storage Gateway Local Console (p. 234).	You use the <code>sguser</code> user and your private key to grant or revoke access for AWS support to your gateway. By default, access to the gateway by AWS support is not enabled. You can enable support access after you launch your instance. Be sure to keep the key pair that you used during for the creation of instance so that you can connect to the instance later. For more information, see Enabling and Disabling AWS Support Access (p. 149).
Multiple network adapters	You can configure multiple network adapters for an on-premises gateway. For more information, see Configuring AWS Storage Gateway for Multiple Network Adapters (NICs) (p. 245).	You cannot configure multiple adapters for an EC2-hosted gateway.
Maintaining your gateway	For information about maintaining your on-premises gateway, see Configuring AWS Storage Gateway for Multiple Network Adapters (NICs) (p. 245).	For information about maintaining your gateway deployed on Amazon EC2, see Managing Your Amazon EC2 Gateway (p. 146).

Launching and Activating an Amazon EC2 Gateway AMI

This section describes how to launch and activate a gateway deployed on Amazon Elastic Compute Cloud (Amazon EC2). The steps you take are summarized as follows:

1. In the AWS Storage Gateway console, launch the setup wizard, which takes you to AWS Marketplace where you select the AMI.
2. In the Amazon EC2 console, configure the AMI and launch an instance.
3. In the Amazon EC2 console, get the IP address of the instance.
4. In the AWS Storage Gateway console, activate the gateway.

Launching the AWS Storage Gateway AMI

You can access the AWS Storage Gateway AMI using the AWS Storage Gateway console as shown below or by going directly to it (as [AWS Storage Gateway](#)) in AWS Marketplace.

Important

Regardless of how you access the AMI, we strongly recommend that you choose the **Launch with EC2 Console** option in AWS Marketplace for launching your instance. The steps for doing this are documented below. If you choose to use the **1-Click Launch** functionality to launch an instance, you will need to add Amazon EBS volumes to your instances as a separate step after the instance is launched (see [Adding and Removing Amazon EBS Volumes from Your Instance](#) (p. 147)).

To launch the AWS Storage Gateway AMI

1. In the AWS Storage Gateway console, click **Deploy a new Gateway on Amazon EC2**.

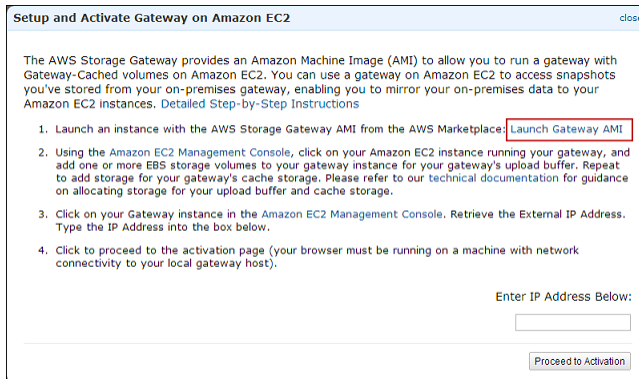
AWS Storage Gateway User Guide

Launching and Activating an Amazon EC2 Gateway AMI



2. In the **Setup and Activate Gateway on Amazon EC2** wizard, in step 1, click **Launch Gateway AMI**.

This will launch in a new browser tab.



3. On the AWS Marketplace page for the AMI, click **Continue**.

AWS Storage Gateway
Sold by: Amazon Web Services

The AWS Storage Gateway is a service connecting a software appliance with cloud-based storage to provide seamless and secure integration between an organization's EC2 instances and AWS's storage infrastructure. The service enables you to securely store data to the AWS cloud for scalable and cost-effective storage. The AWS Storage Gateway supports industry-standard storage protocols that work with your existing applications. It provides low-latency performance by maintaining frequently accessed data locally while securely storing all of your data encrypted in the Amazon Simple Storage Service ... [Read more](#)

Customer Rating: [Be the first to review this product](#)

Latest Version: 1.0

Base Operating System: Linux/Unix, Amazon Linux 2012.09

Delivery Method: 64-bit Amazon Machine Image (AMI) ([Learn more](#))

Support: [See details below](#)

AWS Services Required: Amazon EC2, Amazon EBS

Highlights

- Secure: The AWS Storage Gateway securely transfers your data to AWS over SSL and stores data encrypted at rest in Amazon S3 using Advanced Encryption Standard (AES) 256, a secure symmetric-key encryption standard using 256-bit encryption keys.
- Durably backed by Amazon S3: The AWS Storage Gateway durably stores your application data by uploading it to Amazon S3. Amazon S3 stores data in multiple facilities and on multiple devices within each facility.

Pricing Details

Hourly Fees
Total hourly fees will vary by instance type and EC2 region.


For region: **US East (Virginia)**

EC2 Instance Type	Software	EC2	Total*
Standard XL (m1.xlarge)	\$0.00/hr	\$0.52/hr	\$0.52/hr
High-Memory XL (m2.xlarge)	\$0.00/hr	\$0.45/hr	\$0.45/hr
High-Memory 2XL (m2.2xlarge)	\$0.00/hr	\$0.90/hr	\$0.90/hr
High-Memory 4XL (m2.4xlarge)	\$0.00/hr	\$1.80/hr	\$1.80/hr
High-CPU XL (c1.xlarge)	\$0.00/hr	\$0.66/hr	\$0.66/hr

*EBS fees and data transfer fees not included. Assumes On-Demand EC2 pricing; prices for Reserved and Spot Instances will be lower. [See details](#)


[Learn about instance types](#)

4. On the launch page for the AMI, select the **Launch with EC2 Console** tab.



5. If this is your first time using the AWS Storage Gateway AMI, click **Accept Terms**; otherwise, skip to the next step.

Keep the browser page open. Within a few moments, a subscription confirmation email is sent to the email address of the account with which you logged into AWS Marketplace.



- In the **Region** list, select the region you want to launch the instance in by clicking the **Launch with EC2 Console** link next to the Region.

Region	ID	
US East (Virginia)	ami-200c6949	Launch with EC2 Console
US West (Oregon)	ami-4a7aee7a	Launch with EC2 Console
US West (Northern California)	ami-ee96bbab	Launch with EC2 Console
EU West (Ireland)	ami-6804111c	Launch with EC2 Console
Asia Pacific (Singapore)	ami-60b0fc32	Launch with EC2 Console
Asia Pacific (Sydney)	ami-aad84890	Launch with EC2 Console
Asia Pacific (Tokyo)	ami-9867e499	Launch with EC2 Console
South America (Sao Paulo)	ami-87974d9a	Launch with EC2 Console

When you click a link, you are directed to the Amazon EC2 console.

Configuring the AMI and Launching an Instance

At this point, you have found the [AWS Storage Gateway](#) in AWS Marketplace and selected **Launch with EC2 Console**. The procedure in this section explains how to finish configuring the instance and launch it. There are three things you should keep in mind as you configure the instance:


- The instance type must be one of the types described on the AWS Marketplace page for [AWS Storage Gateway](#) or the instance will not launch. For example, the type must be at least a Standard XL (m1.xlarge).
- The instance comes provisioned with two storage devices, a ROOT device and one EBS volume. Do not remove the EBS volume. You will need to add additional EBS volumes that you will later allocate as local storage for the gateway to use. For a cached gateway, you need to add at least two more EBS volumes, one for cache storage and one for upload buffer. Follow the guidelines for sizing these two storage types as discussed in [Sizing the Upload Buffer \(Gateway-Cached\)](#) (p. 98) and [Sizing Cache Storage \(Gateway-Cached\)](#) (p. 94).
- After you select the AMI from AWS Marketplace and begin to configure the instance, you must assign the instance to one or more security groups. A security group controls traffic to your gateway instance. At least one security that the gateway is assigned to must allow port (80) for activation to occur. To allow connections to iSCSI storage targets of the gateway, you will need to allow port 3260 traffic as well. You might want to check your existing security groups or create a new security group for your gateway instance before you launch your instance. For more information about the security group requirements, see [Configuring Security Groups for Your Amazon EC2 Gateway Instance](#) (p. 149).

To configure and launch an instance

- In the **Request Instances Wizard**, click **Continue**.

AWS Storage Gateway User Guide

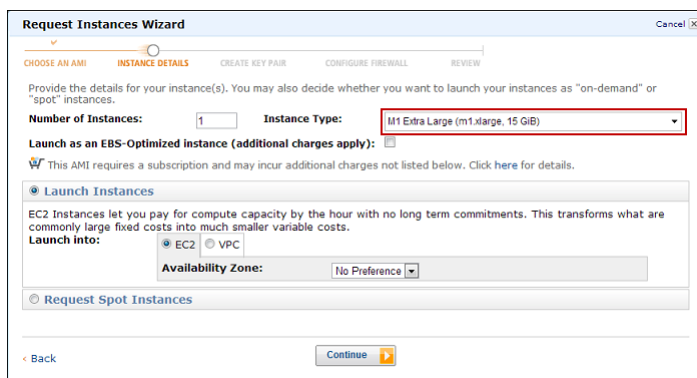
Launching and Activating an Amazon EC2 Gateway AMI



2. In the **INSTANCE DETAILS** step of the wizard, configure the details of the instance.
 - a. Configure the instance type and then click **Continue**.

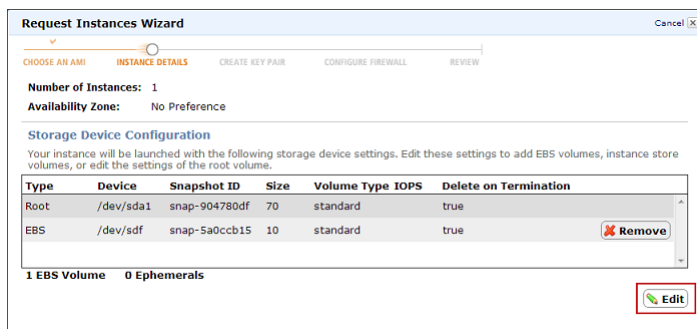
Important

You must specify at least a Standard XL (m1.xlarge) instance type or the instance will not launch. Review the instance types you can launch from this AMI on the [AWS Storage Gateway AMI page](#) in AWS Marketplace.



- b. Accept defaults for **Advanced Instance** options and click **Continue**.
 - c. Configure the storage device settings for the instance.

The instance comes provisioned with two storage devices, a ROOT device and one EBS volume. Do not remove the EBS volume. You must add more storage devices so that you can configure them as upload buffer and cache storage later.



- i. Click **Edit** to add a storage device.
 - ii. Select the **EBS Volumes** radio button, specify the details of the volume, and click **Add**.

AWS Storage Gateway User Guide

Launching and Activating an Amazon EC2 Gateway AMI

Storage Device Configuration
Your instance will be launched with the following storage device settings. Edit these settings to add EBS volumes, instance store volumes, or edit the settings of the root volume.

Root Volume
 EBS Volumes
 Instance Store Volumes

Create and map an EBS volume to the specified device. Increasing EBS Performance.

Snapshot:

Volume Size: GIB Volume Type: IOPS:

Device: Delete on Termination:

Type	Device	Snapshot ID	Size	Volume Type	IOPS	Delete on Termination
Root	/dev/sda1	snap-904780df	70	standard		true
EBS	/dev/sdb	snap-5a0ccb15	10	standard		true <input type="button" value="Remove"/>

iii. Continue to use the **Add** button to add more volumes as needed.

Add at least two storage devices so that you can later configure one storage device as upload buffer and one as cache storage. For a gateway-cached setup, you can add up to a total of 18 TiB of storage where up to 2 TiB can be allocated to upload buffer and up to 16 TiB allocated to cache storage.

iv. When you have added all the volumes you need, click **Continue**.

Storage Device Configuration
Your instance will be launched with the following storage device settings. Edit these settings to add EBS volumes, instance store volumes, or edit the settings of the root volume.

Root Volume
 EBS Volumes
 Instance Store Volumes

Optionally, edit the root volume of your instance and then click Save.

Volume Size: GIB Volume Type: IOPS:

Device: Delete on Termination:

Type	Device	Snapshot ID	Size	Volume Type	IOPS	Delete on Termination
Root	/dev/sda1	snap-904780df	70	standard		true
EBS	/dev/sdb		20	standard		true <input type="button" value="Remove"/>
EBS	/dev/sdc		10	standard		true <input type="button" value="Remove"/>

3 EBS Volumes 0 Ephemerals

d. (Optional) In the **INSTANCE DETAILS** step create tags for the instance.

For example, you might give a value to the **Name** key so that you can later easily recognize the gateway in a list of instances.

Request Instances Wizard Cancel X

CHOOSE AN AMI **INSTANCE DETAILS** CREATE KEY PAIR CONFIGURE FIREWALL REVIEW

Add tags to your instance to simplify the administration of your EC2 infrastructure. A form of metadata, tags consist of a case-sensitive key/value pair, are stored in the cloud and are private to your account. You can create user-friendly names that help you organize, search, and browse your resources. For example, you could define a tag with key = Name and value = Webserver. You can add up to 10 unique keys to each instance along with an optional value for each key. For more information, go to [Using Tags in the EC2 User Guide](#).

Key (127 characters maximum)	Value (255 characters maximum)	Remove
Name	ec2 cached gateway	<input type="button" value="Remove"/>
		<input type="button" value="Remove"/>

Add another Tag. (Maximum of 10)

3. In the **CREATE KEY PAIR** step, choose a key pair, and click **Continue**.

Important

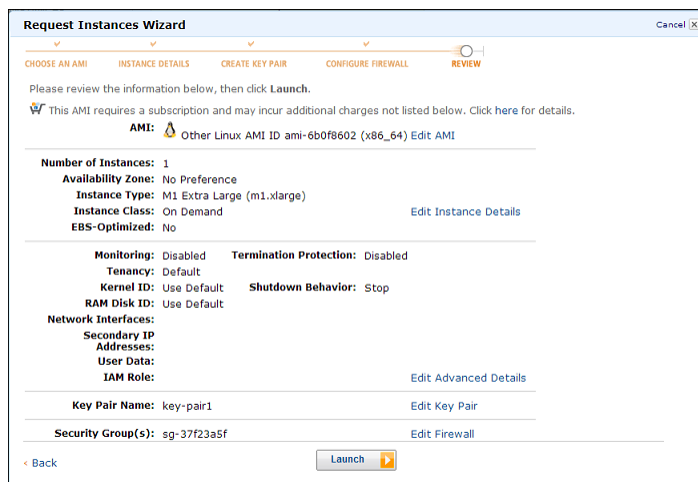
Save the private key of the pair so that you can later enable AWS Support access to the gateway. For more information, see [Enabling and Disabling AWS Support Access \(p. 149\)](#).

4. In the **CONFIGURE FIREWALL** step, assign your instance to one or more security groups, and click **Continue**.

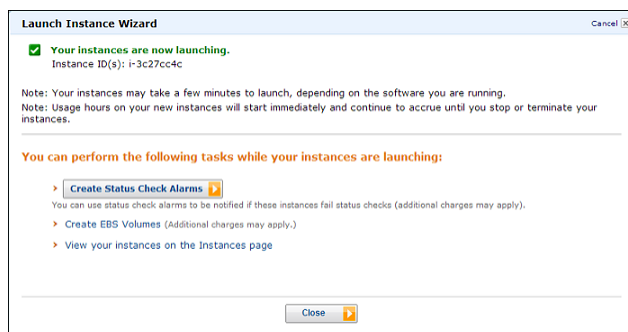
Important

At least one security group must have port (80) allowed so you can activate your gateway. For more information about configuring security groups for your gateway instance such as using a custom security group, see [Configuring Security Groups for Your Amazon EC2 Gateway Instance](#) (p. 149)

5. Review the information for the creation of the instance and then click **Launch**.



6. In the final page of the **Launch Instance Wizard**, click **Close**.



Finding the IP Address of the Amazon EC2 Instance

After you launch an instance based on the [AWS Storage Gateway](#) in AWS Marketplace, it may take several minutes for the instance to become available.

To find the IP address of an Amazon EC2 gateway instance

1. In the Amazon EC2 console, click **Instances** in the left navigation pane
2. In the instances list, find and select the gateway instance.

To confirm you have selected the correct instance, check that the AMI field in the **Description** tab of the instance starts with this string: "aws-storage-gateway-ami-v". This indicates that the instance is based on the AWS Storage Gateway AMI.

3. In the **Description** tab of the instance, find the **Public DNS** value.

The IP address of the instance is part of the DNS name of the instance.

In the following example, the Public DNS value is `ec2-174-129-175-69.compute-1.amazonaws.com`, from which you can infer the IP address as `174.129.175.69`. This IP address is what you'll use to activate the EC2 gateway. Your gateway's IP address will be different.



Warning


Stopping the instance may cause the IP address to change when the instance is restarted. In this case, initiators previously connected to the gateway volumes will not be able to reconnect. However, the IP address of the instance will not change when rebooting the instance. Rebooting the instance may be needed for some maintenance operations on the gateway.

Activating the Gateway

At this point, you created an Amazon EC2 instance that hosts a gateway and you have obtained the IP address of the instance. You are ready to activate the gateway.

To activate the gateway

1. In the AWS Storage Gateway console, start the **Setup and Activate Gateway on Amazon EC2** wizard if it isn't already started.
2. Enter the IP address of the gateway in the **Enter IP Address** box, and then click **Proceed to Activation**.



3. On the activation page, fill in the requested information to complete the activation process.

The **Gateway Type** specifies what type of gateway you are activating. In this case, you are activating a gateway for cached volumes.

The **AWS Region** determines where AWS stores your snapshots. If you choose to restore a snapshot to an Amazon EBS volume, then the Amazon EBS volume must be in the same region as the snapshot. You cannot change the region after the gateway is activated.

Note

If you choose a region in the activation wizard that is different from the one you launched the AMI in (from AWS Marketplace), then additional charges may apply.

The **Gateway Time Zone** is the time zone used when displaying time-based information such as maintenance messages from AWS and snapshot scheduling. You can change the time zone post-activation.

The **Gateway Name** identifies your gateway in the console. You use this name to manage your gateway in the console, and you can change it post-activation.

4. Click **Activate My Storage Gateway**.

If activation does not occur in a few moments check the troubleshooting steps in [Troubleshooting Amazon EC2 Gateway Issues](#) (p. 255),

Managing Your Amazon EC2 Gateway

Topics

- [Adding and Removing Amazon EBS Volumes from Your Instance](#) (p. 147)
- [Enabling and Disabling AWS Support Access](#) (p. 149)
- [Configuring Security Groups for Your Amazon EC2 Gateway Instance](#) (p. 149)
- [Cleaning Up Resources After Deleting a Gateway Deployed on Amazon EC2](#) (p. 150)

In this section, we review how you can manage your gateway deployed on Amazon Elastic Compute Cloud (Amazon EC2). Management tasks you will perform with your gateway include adding and removing Amazon EBS volumes, enabling and disabling AWS Support access to your gateway, configuring security groups, and cleaning up your AWS resources after you are done working with a gateway.

For information about managing a gateway deployed on-premises, see [Managing Your Activated Gateway](#) (p. 176).

Adding and Removing Amazon EBS Volumes from Your Instance


After your gateway is deployed on Amazon Elastic Compute Cloud (Amazon EC2) and activated (see [Launching and Activating an Amazon EC2 Gateway AMI \(p. 139\)](#)), you might need to configure additional storage to be used as upload buffer and cache storage capacity as your application needs change. Or, you might also need to reduce upload buffer or cache storage. In either case, you work with Amazon Elastic Block Store (Amazon EBS) storage, either adding more block storage or reducing it. For more information about Amazon EBS, go to [Amazon Elastic Block Store \(Amazon EBS\)](#) in the *Amazon EC2 User Guide*.

Before you add more storage to the gateway, you should review how to size your upload buffer and cache storage based on your application needs for a cached gateway ([Sizing the Upload Buffer \(Gateway-Cached\) \(p. 98\)](#) and [Sizing Cache Storage \(Gateway-Cached\) \(p. 94\)](#)). Once you configure additional local storage, you work with it just as you would with an on-premises gateway (for example, see [Configuring Upload Buffer and Cache Storage \(p. 150\)](#)). For a gateway-cached setup, you can have up to 18 TiB of storage consisting of up to 2 TiB allocated to upload buffer and up to 16 TiB allocated to cache storage.


The following tasks demonstrates how to add and remove Amazon EBS volumes from an instance. While this is straightforward task for Amazon EC2 instances, you need to take a little extra care when the instance is hosting AWS Storage Gateway. These procedures assumes that you already have a deployed and activated gateway.

To add an Amazon EBS volume to your Amazon EC2–hosted gateway

1. In the Amazon EC2 console, in the navigation pane, click **Volumes**.




2. Click **Create Volume**.



3. In the **Create Volume** dialog box, specify the size of the volume, select an Availability Zone, and then click **Yes, Create**.

Important

Create the Amazon EBS volume in the same availability zone as your gateway; otherwise, you will not be able to attach it to the gateway instance.



4. In the Amazon EC2 console, in the navigation pane under **Elastic Block Store**, click **Volumes**.
5. Find the volume that was created, right-click it, and select **Attach Volume**.
6. In the **Attach Volume** dialog box, specify a gateway instance, and then click **Yes, Attach**.

If you get an error that the device is already in use, choose a different device attachment point. For example, if `/dev/sdg` is in use, try `/dev/sdh`. For more information, see [Attaching a Volume to an Instance](#) in the *Amazon EC2 User Guide*.



7. In the Amazon EC2 console, in the navigation pane, click **Instances**, and select the gateway instance to show its details.

Confirm in the **Block Devices** section of the instance details that a new device was added.



8. In the AWS Storage Gateway console, configure the Amazon EBS volume you added as either upload buffer or cache storage.

To remove an Amazon EBS volume from your Amazon EC2–hosted gateway

1. Shut down the gateway by following the steps in the [Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console](#) (p. 224) procedure.

Note

Before shutting down the gateway, ensure that it is not in use by an application that is writing data to it and that no snapshots are progress. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the console. For more information, see [Editing a Snapshot Schedule](#) (p. 208).

2. In the Amazon EC2 console, in the navigation pane, click **Instances**.
3. Find the instance running the gateway.
4. Note the block devices attached to the instance and find the device you want to remove.

Note that the root device and the swap device are X and Y and should not be removed.

5. In the Amazon EC2 console, in the navigation pane, click **Volumes**.

6. Find the volume you want to remove from the gateway, select it, right-click, and select **Force Detach**.

Enabling and Disabling AWS Support Access

After you have deployed a gateway on Amazon EC2 you can optionally enable AWS Support access to the gateway to help troubleshoot issues. By default, AWS Support access is disabled.

To enable AWS support access to a gateway deployed on Amazon EC2

1. If the security group you specified when you launched the instance does not contain a rule allowing SSH (port 22) access, add it.

For more information about security groups and how to add a security group rule, go to [Amazon EC2 Security Groups](#) in the *Amazon EC2 User Guide*

Note

If you are adding a new rule to an existing security group, you should understand the implications for all instances that use that security group.

2. To enable AWS Support access use the following command.

```
ssh -i IDENTITY_FILE sguser@INSTANCE_IP_ADDRESS grant-aws-support-access
```

Where *IDENTITY_FILE* is the .pem private key file of the key pair you used when you created the instance, and *INSTANCE_IP_ADDRESS* is the IP address of the gateway.

To disable AWS support access to a gateway deployed on Amazon EC2

- To disable AWS Support access use the following command.

```
ssh -i IDENTITY_FILE sguser@INSTANCE_IP_ADDRESS revoke-aws-support-access
```

Where *IDENTITY_FILE* is the .pem private key file of the key pair you used when you created the instance, and *INSTANCE_IP_ADDRESS* is the IP address of the gateway.

Configuring Security Groups for Your Amazon EC2 Gateway Instance

A security group controls traffic to your Amazon EC2 gateway instance. When you create an instance from the AWS Storage Gateway AMI from AWS Marketplace, you have two choices for launching the instance. In [Launching the AWS Storage Gateway AMI \(p. 139\)](#) we showed how to launch the instance by using the **Launch with an EC2 Console** feature of AWS Marketplace. This is the recommended approach. You can also launch an instance by using the **1-Click Launch** feature in AWS Marketplace. In this case, an auto-generated security group `AWS Storage Gateway-1-0-AutogenByAWSMP-` is created. This security group has the correct port (80) rule to allow you to activate your gateway. For more information about security groups, go to [Security Group Concepts](#) in the *Amazon Elastic Compute Cloud User Guide*.

Regardless of the security group that you use, we recommend that:

- The security group should not allow incoming connections from the outside Internet; it should allow only instances within the appliance's security group to talk to the appliance. If you need to allow instances to connect to the appliance from outside the appliance's security group, we recommend that you allow connections only on ports 3260 (for iSCSI) and 80 (for activation).

- You allow port 22 access only if you are using AWS Support for troubleshooting purposes. For more information, see [Enabling and Disabling AWS Support Access \(p. 149\)](#).
- If you wish to activate your appliance from a host outside the appliance's security group, you will need to allow incoming connections on port 80 from the IP address of that host. If you cannot determine the activating host's IP address, you can open up port 80, activate your gateway, and then close access on port 80 after completing activation.

If you are using an Amazon EC2 instance as an initiator, that is, to connect to the iSCSI targets on the gateway you deployed on Amazon EC2, then you have two options. You can put the initiator instance in the same security group as the gateway or you will need to configure access so the initiator can communicate with the gateway.

Cleaning Up Resources After Deleting a Gateway Deployed on Amazon EC2

If you are done using a gateway you deployed on Amazon EC2, it is recommended that you clean up the AWS resources that were used for the gateway, specifically the Amazon EC2 instance and any Amazon EBS volumes. Doing so helps avoid unintended usage charges. We suggest that you take the following actions for deleting your gateway and cleaning up its resources:

- In the AWS Storage Gateway console, delete the gateway as shown in [Deleting a Gateway Using the AWS Storage Gateway Console \(p. 233\)](#).
- In the Amazon EC2 console, stop the instance if you plan on using the gateway again. Terminate the instance if you do not plan on using the instance again. Before terminating the instance, note the block devices and their identifiers that attached to the instance if you plan on deleting volumes.
- In the Amazon EC2 console, remove any Amazon EBS volumes that were attached to the instance if you do not plan on using them again.

Configuring Upload Buffer and Cache Storage

Topics

- [Configuring Upload Buffer \(Gateway-Cached\) \(p. 150\)](#)
- [Configuring Cache Storage \(Gateway-Cached\) \(p. 152\)](#)
- [Configuring Upload Buffer \(Gateway-Stored\) \(p. 154\)](#)


To provide seamless integration between your on-premises environment and AWS's storage infrastructure, each gateway requires some local storage that it uses to buffer and cache data. This section discusses how to configure local disk storage for a gateway. For a gateway-cached volume setup, you configure local disk storage as an upload buffer and cache storage. For a gateway-stored volume setup, you configure local disk storage as an upload buffer. For more information about the different architectures you can choose for your gateway, see [How AWS Storage Gateway Works \(p. 3\)](#).

Configuring an upload buffer and cache storage for a gateway is required before you can create a volume for your applications to use. The upload buffer and cache storage are created from local disks you provisioned for your gateway VM (see [Provisioning Local Disk Storage for an AWS Storage Gateway VM \(p. 92\)](#)).

Configuring Upload Buffer (Gateway-Cached)

This section describes how to configure your gateway's upload buffer. Your gateway requires an upload buffer to temporarily buffer your volume data prior to uploading it to AWS. The following diagram highlights

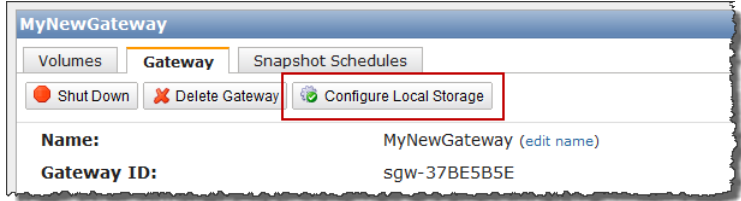
the upload buffer in the larger picture of the AWS Storage Gateway gateway-cached architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).



To configure upload buffer, you need to make sure you have local disks on the gateway VM that are available for use. For instructions about adding more local disks to your VM, see [Provisioning Local Disks \(Gateway-Cached\) \(p. 93\)](#).

To configure a local disk as an upload buffer for your gateway

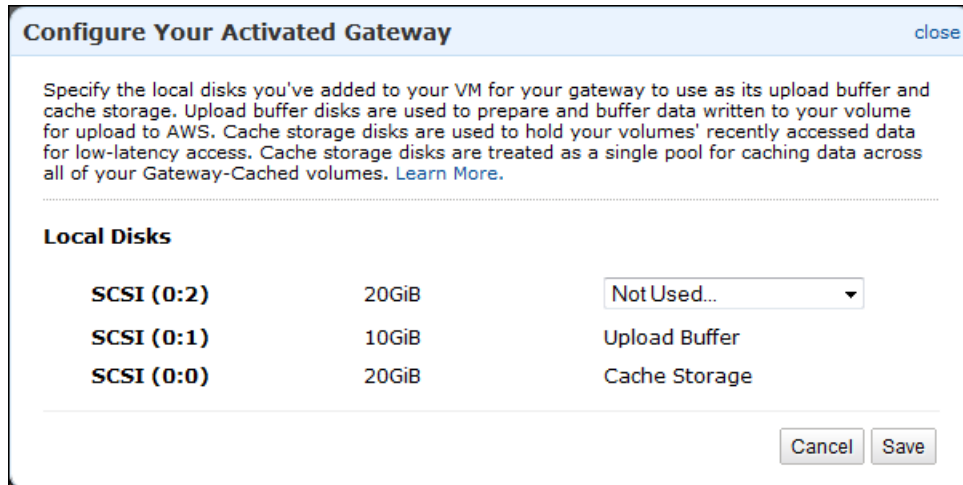
1. In the AWS Storage Gateway console:

If...	Then...
You are configuring a volume on the gateway for the first time.	You are directed to the Configure Your Activated Gateway wizard automatically. Go to Step 2.
Your gateway already has volumes defined.	<p>Open the Configure Your Activated Gateway wizard.</p> <ol style="list-style-type: none"> a. Click the gateway in the Navigation pane. b. Select the Gateway tab. c. Click Configure Local Storage. 

2. In the **Configure Your Activated Gateway** wizard, verify that there are local disks available to configure as an upload buffer.

The wizard shows a list of available disks on your local VM. If there are no local disks available, you must first add a local disk to your gateway VM. For more information, see [Adding a Virtual Disk for the Upload Buffer \(Gateway-Cached\) \(p. 99\)](#).

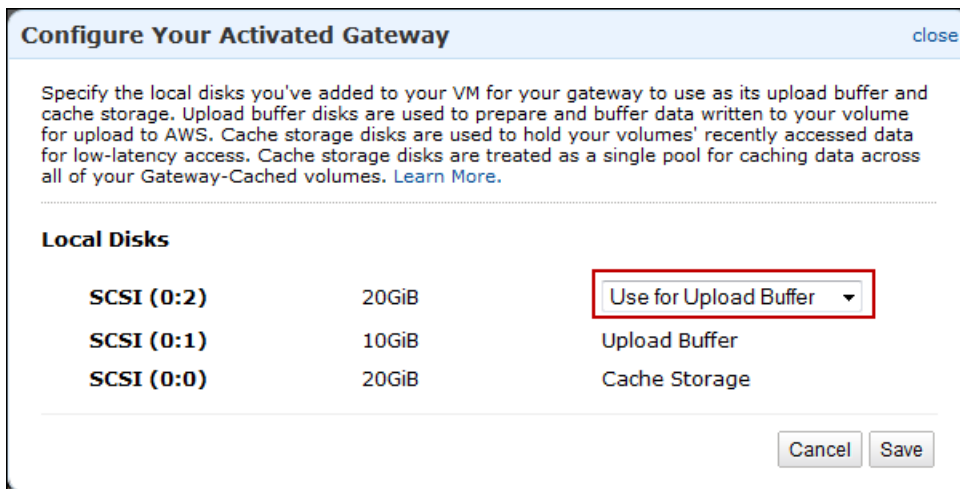
In the following example, the **SCSI (0:2)** disk is available to be configured as an upload buffer.



3. If there are disks available to configure as an upload buffer, then configure the gateway to use them.
 - a. Select the drop-down next to the disks that you want to allocate to the gateway as upload buffer storage.

Important

After configuring a disk as upload buffer storage, you lose any pre-existing data on the disk.

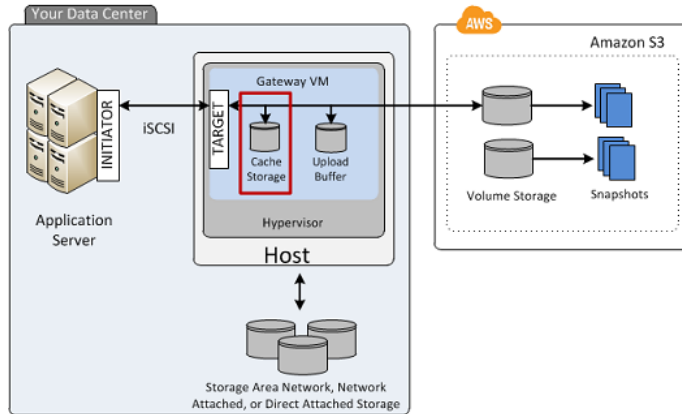


- b. Click **Save**.

This allocates the disk as an upload buffer for the gateway.

Configuring Cache Storage (Gateway-Cached)


This section describes how to configure your gateway's cache storage. Your gateway requires cache storage to cache recently accessed application data. The following diagram highlights the cache storage in the larger picture of the AWS Storage Gateway gateway-cached architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).



To configure a local disk as cache storage for your gateway

1. In the AWS Storage Gateway console:

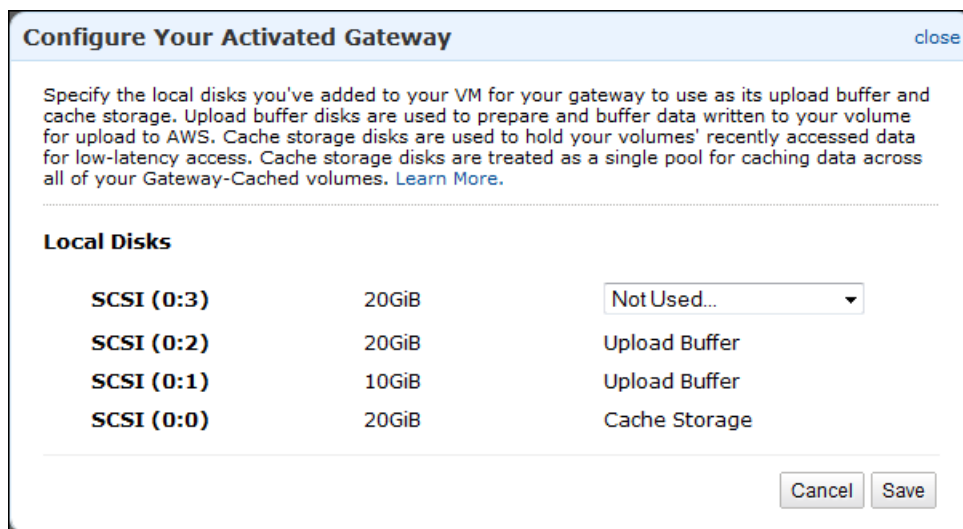
If...	Then...
You are configuring a volume on the gateway for the first time.	You are directed to the Configure Your Activated Gateway wizard automatically. Go to Step 2.
Your gateway already has volumes defined.	<p>Open the Configure Your Activated Gateway wizard.</p> <ol style="list-style-type: none"> a. Click the gateway in the Navigation pane. b. Select the Gateway tab. c. Click Configure Local Storage.



2. In the **Configure Your Activated Gateway** wizard, verify that there are local disks available to configure as cache storage.

The wizard shows a list of available disks on your local VM. If there are no local disks available to configure as cache storage, then you must first add a local disk to your gateway VM. For more information, see [Adding a Virtual Disk for Cache Storage \(Gateway-Cached\)](#) (p. 95).

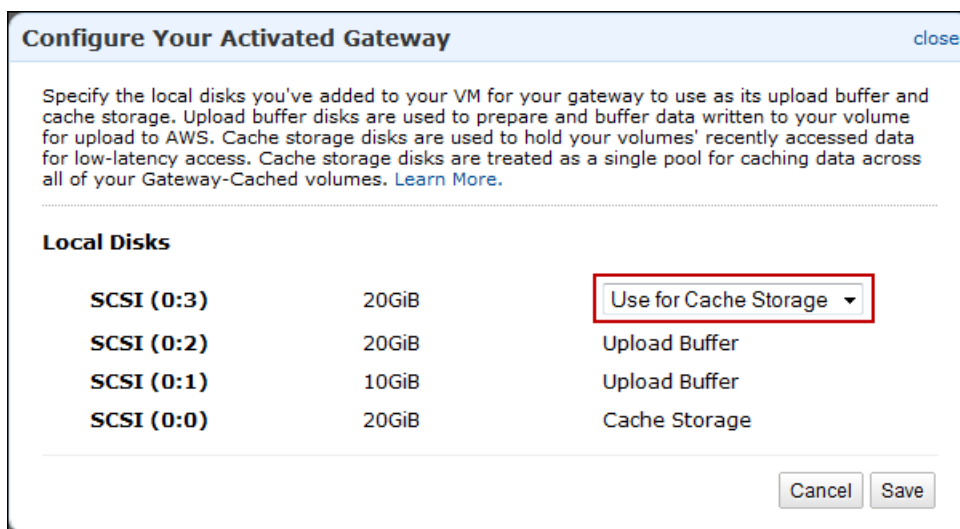
In following example, the **SCSI (0:3)** disk is available to be configured as cache storage.



3. If there are disks available to configure as cache storage, then configure the gateway to use them.
 - a. Select the drop-down next to the disks that you want to allocate to the gateway as cache storage.

Important

After configuring a disk as cache storage, you lose any pre-existing data on the disk.




- b. Click **Save**.

This allocates the disk as cache storage for the gateway.

Configuring Upload Buffer (Gateway-Stored)


This section describes how to configure your gateway's upload buffer. Your gateway requires an upload buffer to temporarily buffer your volume data prior to uploading it to AWS. The following diagram highlights the upload buffer in the larger picture of the AWS Storage gateway-stored architecture (see [How AWS Storage Gateway Works](#) (p. 3)).



To configure a local disk as an upload buffer for your gateway

1. In the AWS Storage Gateway console:

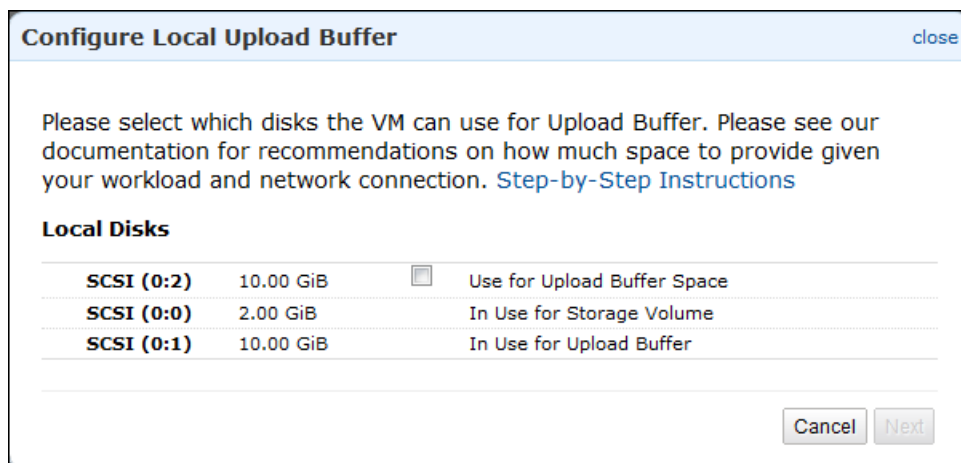
If...	Then...
You are configuring a volume on the gateway for the first time.	You are directed to the Configure Your Activated Gateway wizard automatically. Go to step 2.
Your gateway already has volumes defined.	<p>Open the Configure Your Activated Gateway wizard.</p> <ol style="list-style-type: none"> a. Click the gateway in the Navigation pane. b. Select the Gateway tab. c. Click Configure Local Storage.



2. In the **Configure Local Upload Buffer** wizard, verify that there are local disks available to configure as an upload buffer.

The wizard shows a list of available disks on your local VM. If there are no local disks available, you must first add a local disk to your gateway VM. For more information, see [Adding a Virtual Disk for the Upload Buffer \(Gateway-Stored\)](#) (p. 107).

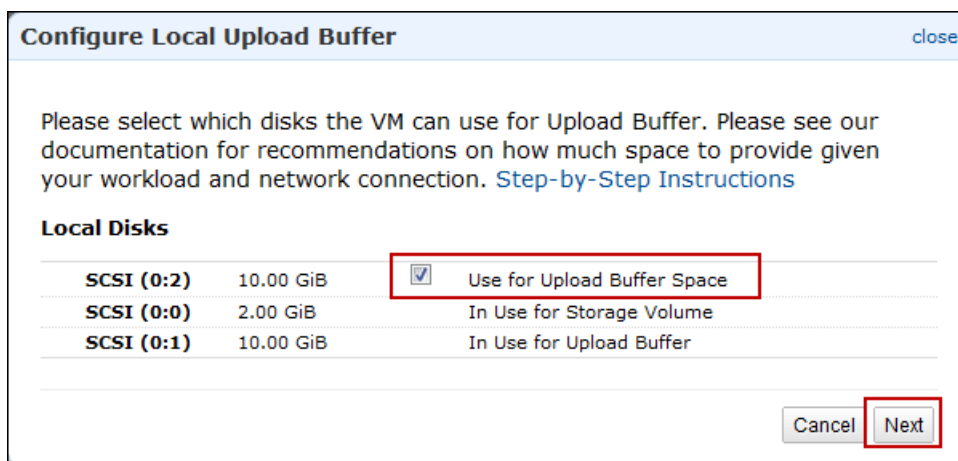
In following example, the **SCSI (0:2)** disk is available to be configured as upload buffer space.



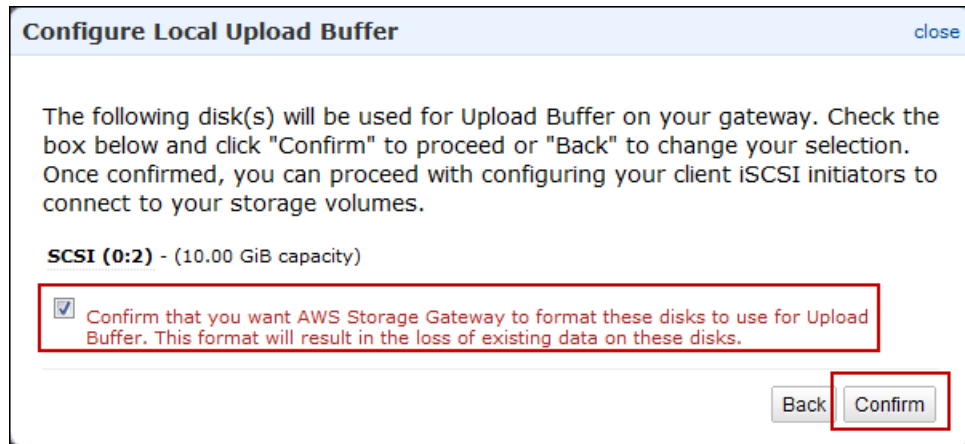
3. If there are disks available to configure as an upload buffer, configure the gateway to use them.
 - a. Select the check box next to the disks that you want to allocate to the gateway as the upload buffer, and then click **Next**. The **Next** button is enabled only if you select at least one disk.

Important

After configuring a disk as upload buffer storage, you lose any pre-existing data on the disk.



- b. In the confirmation dialog box, read and select the confirmation check box and click **Confirm**. This allocates the disk as upload buffer for the gateway.



Creating Storage Volumes

Topics

- [Creating a Storage Volume \(Gateway-Cached\) \(p. 157\)](#)
- [Creating a Storage Volume \(Gateway-Stored\) \(p. 159\)](#)

Your application data is stored on storage volumes. In this section, you learn about creating a storage volume for either the gateway-cached volume architecture or the gateway-stored volume architecture. For more information on the different AWS Storage Gateway architectures, see [How AWS Storage Gateway Works \(p. 3\)](#).

Creating a Storage Volume (Gateway-Cached)

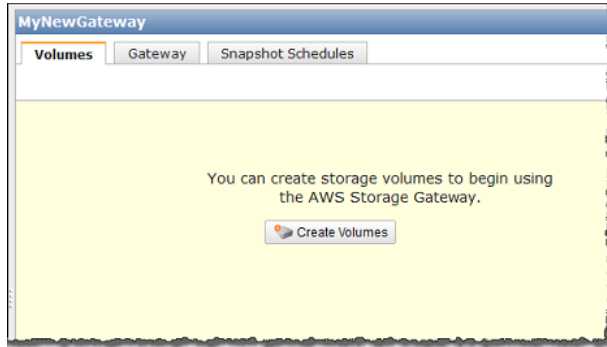
You can create a gateway-cached volume using the AWS Storage Gateway API (see [CreateCachediSCSIVolume \(p. 318\)](#)) or the AWS Storage Gateway console. The following task demonstrates creating a volume using the console. It assumes that you already have deployed and activated your gateway.

To create a storage volume using the console

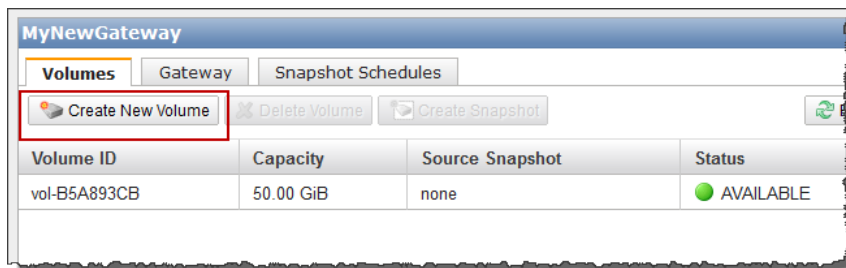
1. In the AWS Storage Gateway console, select the gateway in the **Navigation** pane.
2. If the gateway is activated but has no storage volumes.
 - a. Click **Create Volumes**.

AWS Storage Gateway User Guide

Creating a Storage Volume (Gateway-Cached)



- b. Follow the steps in [Step A: Create Cache Storage and an Upload Buffer on Your Local Disks and Configure Optional Alarms](#) (p. 62) in the Getting Started section to specify the cache storage and upload buffer.
3. If the gateway already has one or more storage volumes.
 - In the **Volumes** tab, click **Create New Volume**.



4. In the **Configure Your Activated Gateway** wizard, configure the volume.

Configure Your Activated Gateway close

Create an iSCSI storage volume up to 32 TBs in size. This volume will be stored in Amazon S3, with only a cache of recently accessed data kept locally. Your client applications will connect to this volume over an iSCSI interface. [Learn More](#).

Capacity: TBs (Max: 32 TBs)

iSCSI Target Name: iqn.1997-05.com.amazon:

Based on Snapshot ID:

Host IP:

Port:

- a. In the **Capacity** field, enter the size of the storage volume to create.

Note

Resizing a storage volume after it is created is not supported. To change the size of a volume later, you will need to create a snapshot of the volume, and create a new cached volume from the snapshot. For more information, see [Managing Storage Volumes \(Gateway-Cached\)](#) (p. 181).

- b. In the size drop-down list next to the **Capacity** field, select the appropriate size of the volume, GiBs or TiBs.
- c. Enter a name in the **iSCSI Target Name** field.

The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name in the **Targets** tab of the **iSCSI Microsoft Initiator** UI after discovery. For example, a name `target1` would appear as `iqn.1997-05.com.amazon:target1`. Ensure that the target name is globally unique within your SAN network.

- d. Specify the **Based on Snapshot ID** field if you are creating a volume from a snapshot.

You can specify the ID of an existing AWS Storage Gateway or Amazon EBS snapshot you previously created. In this case, the gateway creates the storage volume and downloads data to the local cache storage only on first access of the data. To learn about how to find a snapshot you want to use, see [Finding a Snapshot](#) (p. 200).

- e. The IP address shown in the **Host IP** field shows your gateway IP address.

If you've configured your local gateway host with multiple Network Interface Cards (NICs), you can specify which IP address you want to use for this storage volume.

- f. Note that the **Port** field shows the port to map an iSCSI target.

AWS Storage Gateway supports only port 3260.

- g. Click **Create Volume**.

This creates a storage volume and makes your disk available as an iSCSI target for your applications to connect to and store data on. For information on connecting to the iSCSI target, see [Configuring Your Application Access to Storage Volumes](#) (p. 161).

Note

If you want snapshots for this volume, you can either take an ad-hoc snapshot or set up a snapshot schedule for the volume. For more information, see [Editing a Snapshot Schedule](#) (p. 208).

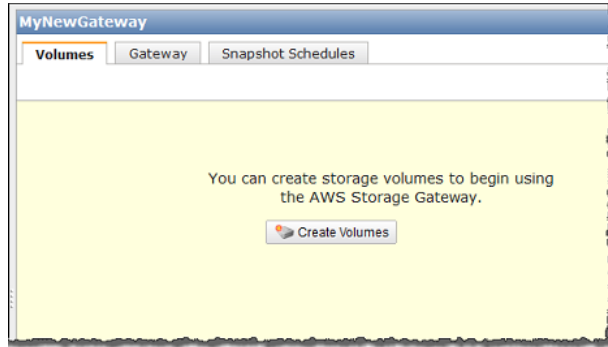
Creating a Storage Volume (Gateway-Stored)

You can create a gateway-stored volume using the AWS Storage Gateway API (see [CreateStorediSCSIVolume](#) (p. 327)) or the AWS Storage Gateway console. The following task demonstrates creating a volume using the console. It assumes that you already have deployed and activated your gateway. Furthermore, it is assumed that there is at least one locally provisioned disk of the gateway that is not used that you will configure as a gateway-stored volume. To provision a local disk for application storage, see [Provisioning Local Disks \(Gateway-Stored\)](#) (p. 102).

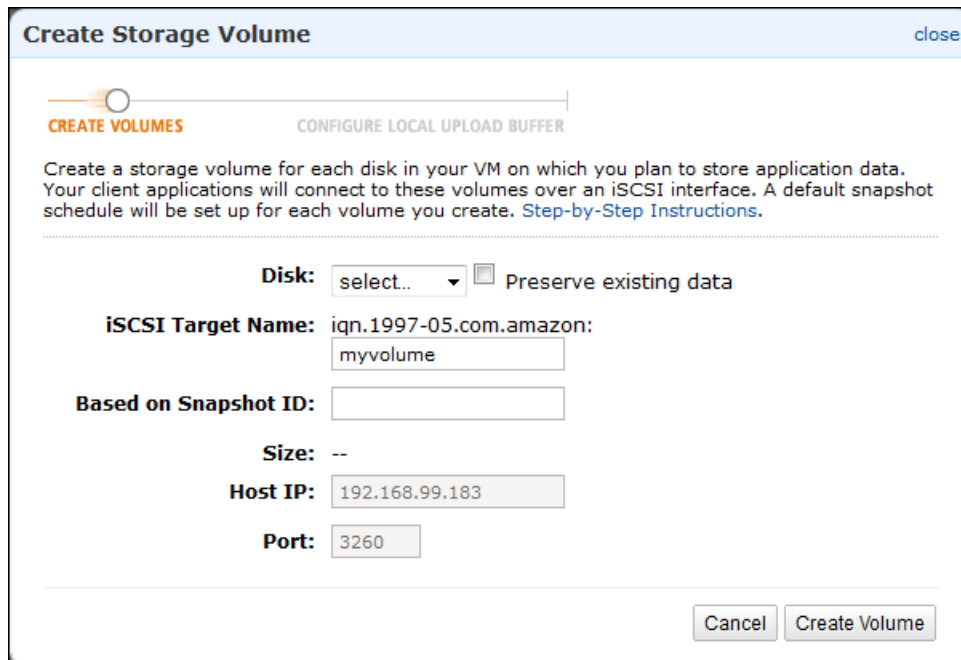
To create a storage volume using the console

1. In the AWS Storage Gateway console, select the gateway in the **Navigation** pane.

The console shows gateway specific information. If the gateway is activated but has no storage volumes, then the console shows the following page with the **Create Volumes** button.



2. Click **Create New Volume**.



3. In the **Create Storage Volume** wizard, enter the following information:
 - a. In the **Disk** drop-down list, select a local virtual disk that you provisioned for the gateway.
For information about provisioning disks, see [Provisioning Local Disk Storage for an AWS Storage Gateway VM](#) (p. 92).
 - b. Select the **Preserve existing data** check box if you want to preserve data on the disk.
AWS Storage Gateway bootstraps your volume upon creation, preserving and uploading your volume's existing data to AWS.
 - c. Enter a name in the **iSCSI Target Name** field.
The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name in the **Targets** tab of the **iSCSI Microsoft Initiator** UI after discovery. For example, a name `target1` would appear as `iqn.1997-05.com.amazon:target1`. Ensure that the target name is globally unique within your SAN network.
 - d. Specify the **Based on Snapshot ID** field if you are creating a volume from a snapshot.

You can specify the ID of an existing AWS Storage Gateway or Amazon EBS snapshot you previously created. This is a useful scenario if you want to restore a snapshot of another storage volume. In this case, the gateway creates the storage volume and downloads your existing snapshot data to the volume. However, there is no need to wait for all of the data to transfer from Amazon S3 to your volume before your application can start accessing the volume and all of its data. To learn more about snapshots, see [Working with Snapshots \(p. 199\)](#).

When you create a volume from an existing snapshot, any existing data on the disk is not preserved, the **Preserve existing data** check box must be unchecked.

- e. The IP address shown in the **Host IP** field shows your gateway IP address.

If you've configured your local gateway host with multiple Network Interface Cards (NICs), you can specify which IP address you want to use for this storage volume.

- f. Note that the **Port** field shows the port to map an iSCSI target.

AWS Storage Gateway supports only port 3260.

- g. Click **Create Volume**.

This creates a storage volume and makes your disk available as an iSCSI target for your applications to connect and store data.

Clicking this button also creates a snapshot schedule for your new volume. By default, AWS Storage Gateway takes snapshots once a day. You can change both the time the snapshot occurs each day, as well as the frequency (every 1, 2, 4, 8, 12, or 24 hours). For more information, see [Editing a Snapshot Schedule \(p. 208\)](#).

Note

Snapshots are incremental, compressed backups. For a given storage volume, the gateway saves only the blocks that have changed since the last snapshot. This minimizes the amount of storage that is used for your backups. To ensure that your gateway can keep up with the rate of incoming writes, it's important that you take snapshots at least once a day.

Configuring Your Application Access to Storage Volumes

Topics

- [Connecting from a Windows Client to Your Storage Volume \(p. 163\)](#)
- [Connecting from a Red Hat Client to Your Storage Volume \(p. 165\)](#)
- [Configuring CHAP Authentication for Your Storage Volume \(p. 167\)](#)

After you add local disks to your VM and create storage volumes, the gateway exposes these disks as iSCSI targets. Your client applications connect to these iSCSI targets to store data. Connect only one application to each iSCSI target. AWS Storage Gateway supports Red Hat and Windows client iSCSI initiators that enable you to connect to the targets. To learn more about adding local disks to your VM, see [Provisioning Local Disk Storage for an AWS Storage Gateway VM \(p. 92\)](#). To learn more about creating storage volumes, see [Managing Storage Volumes in AWS Storage Gateway \(p. 176\)](#).

Note

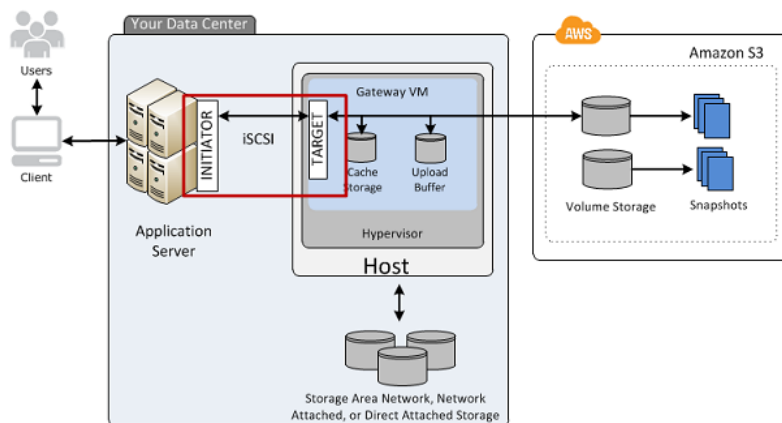
To connect to your storage volume target, your gateway must have an upload buffer configured. If an upload buffer is not configured for your gateway, then the status of your storage volumes is displayed as UPLOAD BUFFER NOT CONFIGURED. To configure an upload buffer for a

AWS Storage Gateway User Guide

Configuring Your Application Access to Storage Volumes

stored-volume gateway, see [To configure an upload buffer for your gateway using the console \(p. 195\)](#). To configure an upload buffer for a cached-volume gateway, see [To configure a local disk as upload buffer space for your gateway using the console \(p. 190\)](#).

The following diagram highlights the iSCSI target in the larger picture of the AWS Storage Gateway architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).



The Internet Small Computer System Interface (iSCSI), is an Internet Protocol (IP)-based storage networking standard for initiating and managing connections between IP-based storage devices, and clients.

The following table shows some of the iSCSI nomenclature that is used to describe the connection and the components involved.

Term	Description
iSCSI initiator	The client component of an iSCSI network. The initiator sends requests to the iSCSI target. Initiators can be implemented in software or hardware. The AWS Storage Gateway only supports software initiators.
iSCSI target	The server component of the iSCSI network that receives and responds to requests from initiators. Each of your storage volumes is exposed as an iSCSI target. Connect only one iSCSI initiator to each iSCSI target.
Microsoft iSCSI initiator	The software program on Windows computers that enables you to connect a client computer (e.g., the computer running the application whose data you want to write to the gateway) to an external iSCSI-based array (i.e., the gateway) using the host computer's Ethernet network adapter card. The Microsoft iSCSI initiator is implemented in software. Microsoft iSCSI initiator is already installed on Windows Server 2008 R2, Windows 7, Windows Server 2008, and Windows Vista. On these operating systems, you do not need to install the initiator.
Red Hat iSCSI initiator	The <code>iscsi-initiator-utils</code> Resource Package Manager (RPM) package provides you with an iSCSI initiator implemented in software for Red Hat. The package includes a server daemon for the iSCSI protocol.

You can connect to your storage volume from either a Windows or Red Hat client. You can optionally configure Challenge-Handshake Authentication Protocol (CHAP) for either client type.

To...	See...
Connect to your storage volume from Windows.	Step 2.3: Access Your AWS Storage Gateway Volumes (p. 72) in the Getting Started tutorial
Connect to your storage volume from Red Hat Linux.	Connecting from a Red Hat Client to Your Storage Volume (p. 165)
Configure CHAP Authentication for Windows and Red Hat Linux.	Configuring CHAP Authentication for Your Storage Volume (p. 167)

Connecting from a Windows Client to Your Storage Volume

When using a Windows client, you use the Microsoft iSCSI initiator to connect to your gateway storage volume.

The Getting Started exercise provides instructions about how to connect to your storage volumes. For more information, see [Step 2.3: Access Your AWS Storage Gateway Volumes \(p. 72\)](#).

Customizing Your Windows iSCSI Settings

After setting up your initiator, we highly recommend that you customize your iSCSI settings to prevent the initiator from disconnecting from targets. By increasing the iSCSI timeout values as shown in the following steps, you improve the ability of your application to deal with writes that take a long time and other transient issues such as network interruptions.

Note

Before making changes to the registry, you should make a backup copy. For information on making a backup copy and other best practices to follow when working with the registry, see [Registry best practices](#) in the *Windows Server TechCenter*.

To customize your Windows iSCSI settings

1. Increase the maximum time for which requests are queued.
 - a. Start Registry Editor (Regedit.exe).
 - b. Navigate to the device class Globally Unique Identifier (GUID) key that contains iSCSI controller settings.

Warning

Make sure you are working in the **CurrentControlSet** subkey and not another control set like **ControlSet001** or **ControlSet002**.

```
HK_Local_Machine\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}
```

- c. Find the subkey for the Microsoft iSCSI Initiator.

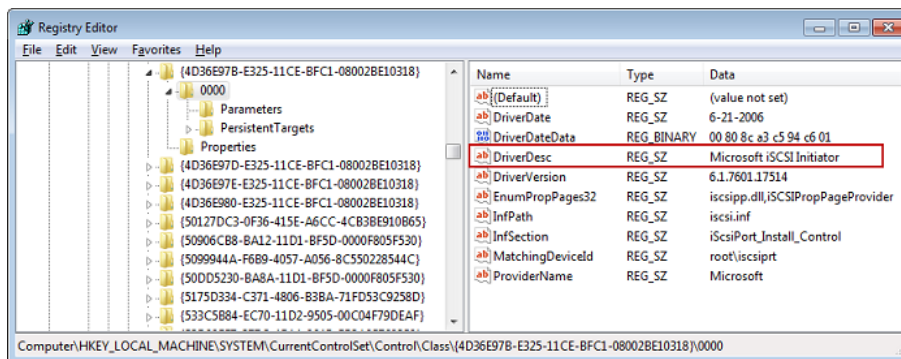
The key will be represented by a four-digit number such as 0000 or 00001.

```
HK_Local_Machine\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>
```

AWS Storage Gateway User Guide


Connecting from a Windows Client to Your Storage Volume

Depending on what is installed on your computer, the Microsoft iSCSI initiator may not be the subkey 0000. You can ensure that you have selected the correct subkey by verifying that the string `DriverDesc` has the value `Microsoft iSCSI Initiator` as shown in the following example.



- d. Click the **Parameters** subkey to show the iSCSI settings.
- e. Right-click the **MaxRequestHoldTime** DWORD (32-bit) value, select modify, and change its value to 600.

This value represents a hold time of 600 seconds. The example below shows the **MaxRequestHoldTime** string value with a value of 600.



2. Increase the disk timeout value.
 - a. Start Registry Editor (Regedit.exe).
 - b. Navigate to the **Disk** subkey in the **Services** subkey of the **CurrentControlSet**.

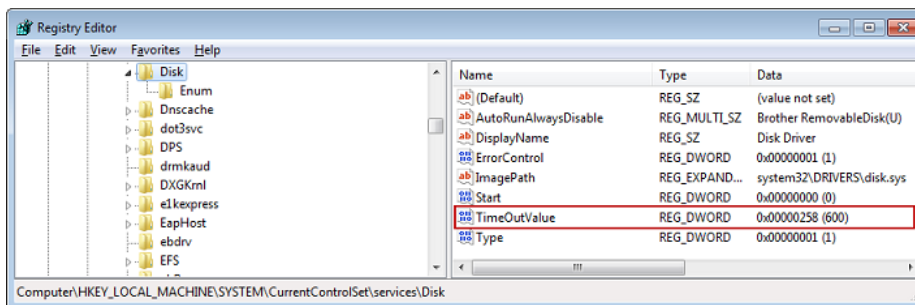
```
HK_Local_Machine\SYSTEM\CurrentControlSet\Services\Disk
```

- c. Right-click the **TimeoutValue** DWORD (32-bit) value, click modify, and change its value to 600.

This value represents a timeout value of 600 seconds.

AWS Storage Gateway User Guide

Connecting from a Red Hat Client to Your Storage Volume



- Restart your system to ensure that the new configuration values take effect.

Before restarting, you must make sure that all writes to storage volumes are flushed. To do this, take any mapped storage volume disks offline before restarting.

Connecting from a Red Hat Client to Your Storage Volume

When using Red Hat Linux, you use the `iscsi-initiator-utils` RPM package to connect to your gateway storage volume.

To connect a Linux client to the storage volume

- Install the `iscsi-initiator-utils` RPM package if it isn't already installed on your client.

You can use the following command to install the package.

```
sudo yum install iscsi-initiator-utils
```

- Ensure that the iSCSI daemon is running.
 - Verify that the iSCSI daemon is running using the following command.

```
sudo /etc/init.d/iscsi status
```

- If the status command does not return a status of *running*, then start the daemon using the following command.

```
sudo /etc/init.d/iscsi start
```

- Discover the storage volume targets defined for a gateway.

Use the following discovery command to list the targets of a gateway.

```
sudo /sbin/iscsiadm --mode discovery --type sendtargets --portal GATEWAY_IP:3260
```

Substitute your gateway's IP address for the `GATEWAY_IP` variable in the preceding command. You can find the gateway IP in the **iSCSI Target Info** properties of a volume in the AWS Storage Gateway console.

The output of the discovery command will look like this example output.

```
GATEWAY_IP:3260, 1 iqn.1997-05.com.amazon:myvolume
```

Your iSCSI Qualified Name (IQN) will be different than what is shown above because IQN values are unique to an organization. The name of the target is the name that you specified when you created the storage volume. You can find this target name as well on the **iSCSI Target Info** properties pane when you select a storage volume in the AWS Storage Gateway console.

4. Connect to a target.

You need to specify the correct `GATEWAY_IP` and IQN in the connect command.

```
sudo /sbin/iscsiadm --mode node --targetname iqn.1997-05.com.amazon:myvolume  
--portal GATEWAY_IP:3260,1 --login
```

5. Verify that the volume is attached to the client machine (initiator).

```
ls -l /dev/disk/by-path
```

After setting up your initiator, we highly recommend that you customize your iSCSI settings as discussed in [Customizing Your Linux iSCSI Settings \(p. 166\)](#).

Customizing Your Linux iSCSI Settings

After setting up your initiator, we highly recommend that you customize your iSCSI settings to prevent the initiator from disconnecting from targets. By increasing the iSCSI timeout values as shown below, you improve the ability of your application to deal with writes that take a long time and other transient issues such as network interruptions.

To customize your Linux iSCSI settings

1. Increase the maximum time for which requests are queued.
 - a. Open the `/etc/iscsi/iscsid.conf` file and find the following lines.

```
node.session.timeo.replacement_timeout = [replacement_timeout_value]  
node.conn[0].timeo.noop_out_interval = [noop_out_interval_value]  
node.conn[0].timeo.noop_out_timeout = [noop_out_timeout_value]
```

- b. Set the `replacement_timeout_value` value to 600.

Set the `noop_out_interval_value` value to 60.

Set the `noop_out_timeout_value` value to 600.

All three values are in seconds.

Note

The `iscsid.conf` settings must be made before discovering the gateway. If you have already discovered your gateway and/or logged in to the target, you can delete the

AWS Storage Gateway User Guide

Configuring CHAP Authentication for Your Storage Volume

entry from the discovery database using the following command, and then rediscover/login to pick up the new configuration.

```
iscsiadm -m discoverydb -t sendtargets -p gateway_ip:3260 -o delete
```

2. Increase the disk timeout value.
 - a. Open the `/etc/udev/rules.d/50-udev.rules` file and find the following line.

```
ACTION=="add", SUBSYSTEM=="scsi", SYSFS{type}=="0|7|14", \
RUN+="/bin/sh -c 'echo [timeout] > /sys$$DEVPATH/timeout' "
```

- b. Set the `timeout` value to 600.

This value represents a timeout value of 600 seconds.

3. Restart your system to ensure that the new configuration values take effect.

Before restarting, you must make sure that all writes to your storage volumes are flushed. To do this, unmount storage volumes before restarting.

Configuring CHAP Authentication for Your Storage Volume

AWS Storage Gateway supports authentication between your gateway and iSCSI initiators via CHAP (Challenge-Handshake Authentication Protocol). CHAP provides protection against playback attacks by periodically verifying the identity of an iSCSI initiator as authenticated to access a storage volume target. To set up CHAP, you must configure it in both the AWS Storage Gateway console and in the iSCSI initiator software you use to connect to the target.

This section discusses mutual CHAP, which is when the initiator authenticates the target and the target authenticates the initiator. To use mutual CHAP, you follow two steps:


- First, configure CHAP in the AWS Storage Gateway console.
 - [To configure CHAP in the AWS Storage Gateway console \(p. 167\)](#)
- Second, in your client initiator software, complete the CHAP configuration.
 - [To configure mutual CHAP on a Windows client \(p. 169\)](#)
 - [To configure mutual CHAP on a Red Hat Linux client \(p. 174\)](#)

To configure CHAP in the AWS Storage Gateway console


In this procedure, you specify two secret keys that are used to read and write to a storage volume. These same keys are used in the procedure to configure the client initiator.

1. Go to the **iSCSI Target Info** tab of the volume for which you want to configure CHAP.
2. Click the **Configure CHAP Authentication** link.

AWS Storage Gateway User Guide
Configuring CHAP Authentication for Your Storage Volume



3. Configure CHAP in the **Configure CHAP Authentication** dialog box.



- a. Check the **Enabled** box.
- b. Specify the **Initiator Name**.

The initiator name can be found using your iSCSI initiator software. For example, for Windows clients, the name is the value in the **Configuration** tab of the iSCSI initiator. For more information, see [To configure mutual CHAP on a Windows client \(p. 169\)](#).

Note

To change an initiator name, you must first disable CHAP, change the initiator name in your iSCSI initiator software, and then enable CHAP with the new name.

- c. Specify the **Secret used to Authenticate Initiator** field.

This secret must be at least 12 characters long. It is the secret key that the initiator (e.g., Windows client) must know to participate in CHAP with the target.

- d. Specify a secret in the **Secret used to Authenticate Target (Mutual CHAP)** field.

This secret must be at least 12 characters long. It is the secret key that the target must know to participate in CHAP with the initiator.

Note

The secret used to authenticate the target must be different than the secret to authenticate the initiator.

- e. Click **Save**.
- f. Click **Close** in the confirmation dialog box.


The **iSCSI Target Info** tab indicates that CHAP authentication is used.

4. Confirm that CHAP is enabled.

The **iSCSI Target Info** tab indicates that CHAP authentication is used.

AWS Storage Gateway User Guide

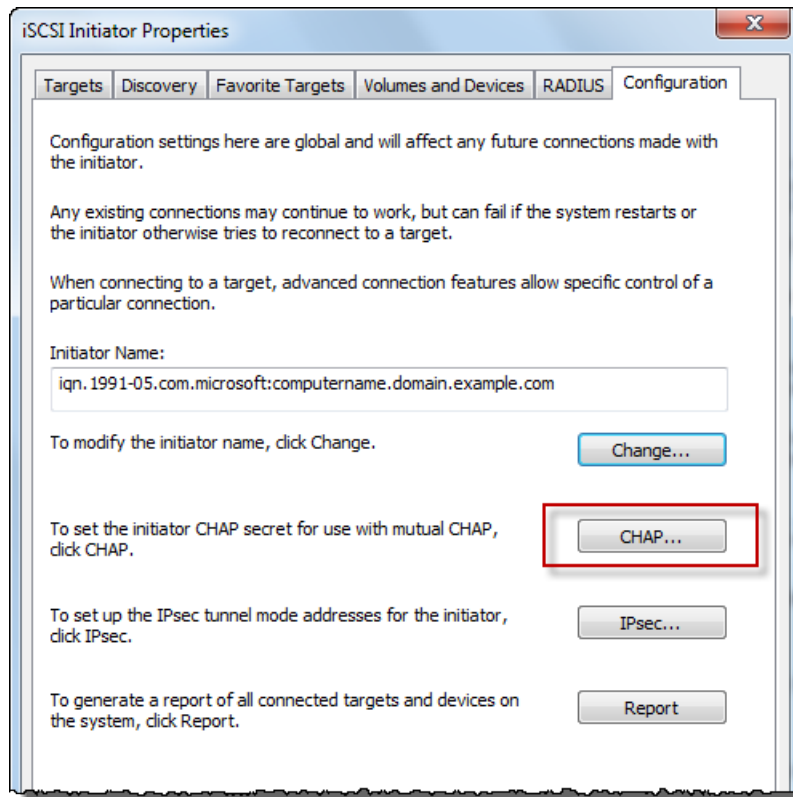
Configuring CHAP Authentication for Your Storage Volume



To configure mutual CHAP on a Windows client

In this procedure, you configure CHAP in the Microsoft iSCSI initiator using the same keys that you used to configure CHAP for the storage volume in the console.

1. If the iSCSI initiator is not already started, then in the **Start** menu of your Windows client computer, type `iscsicpl.exe` and run the program.
2. Configure the initiator's (the Windows client) mutual CHAP configuration.
 - a. Click the **Configuration** tab.

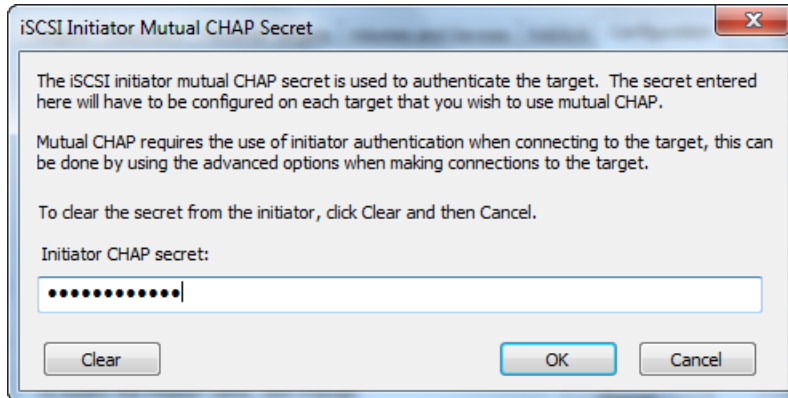


- b. Note that the **Initiator Name** field will be unique to your initiator and company. The name shown here is the value that you used in the **Configure CHAP Authentication** dialog box of the AWS Storage Gateway console.

The name shown in the example image is for demonstration purposes only.

AWS Storage Gateway User Guide
Configuring CHAP Authentication for Your Storage Volume


- c. Click the **CHAP** button.
- d. In the **iSCSI Initiator Mutual Chap Secret** dialog box, enter the mutual CHAP secret value.



In this dialog box, you are entering the secret that the initiator (Windows client) uses to authenticate the target (storage volume). This secret allows the target to read and write to the initiator. This secret maps to the **Secret used to Authenticate Target (Mutual CHAP)** field in the **Configure CHAP Authentication** dialog box. For more information see, [Configuring CHAP Authentication for Your Storage Volume \(p. 167\)](#).

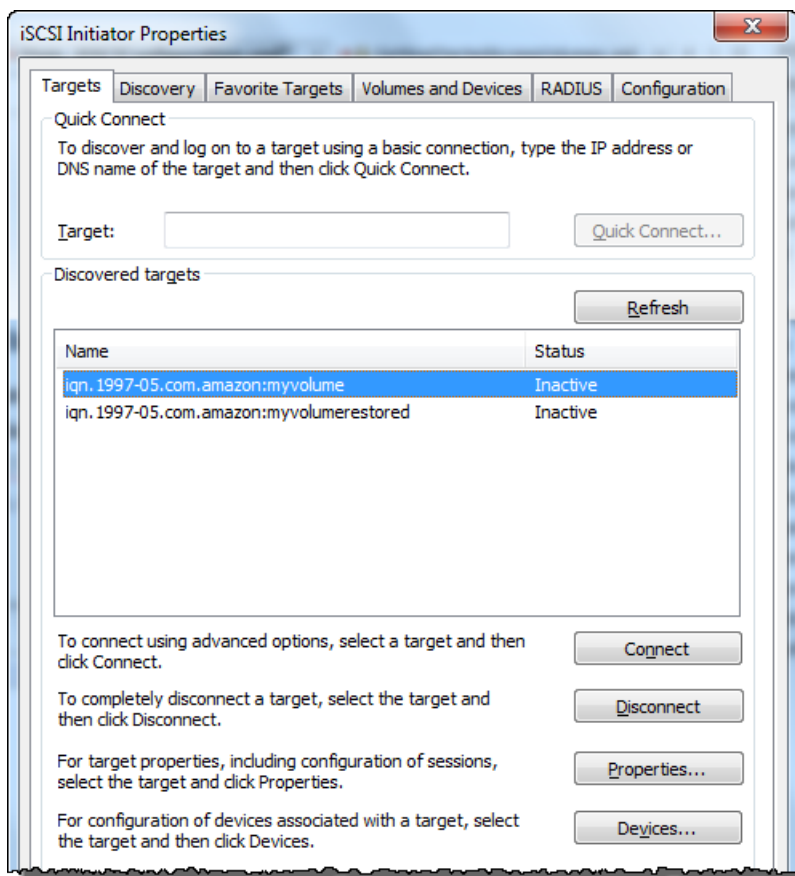
- e. If the key that you enter is less than 12 characters, an **Initiator CHAP secret** error dialog box appears.

Click **OK** and try entering the key again.




3. Configure the target with the initiator's secret to complete the mutual CHAP configuration.
 - a. Click the **Targets** tab.

AWS Storage Gateway User Guide
Configuring CHAP Authentication for Your Storage Volume




- b. Disconnect the target that you want to configure for CHAP if it is currently connected by selecting the target and clicking **Disconnect**.
- c. Select the target that you want to configure for CHAP, and click **Connect**.

AWS Storage Gateway User Guide
Configuring CHAP Authentication for Your Storage Volume

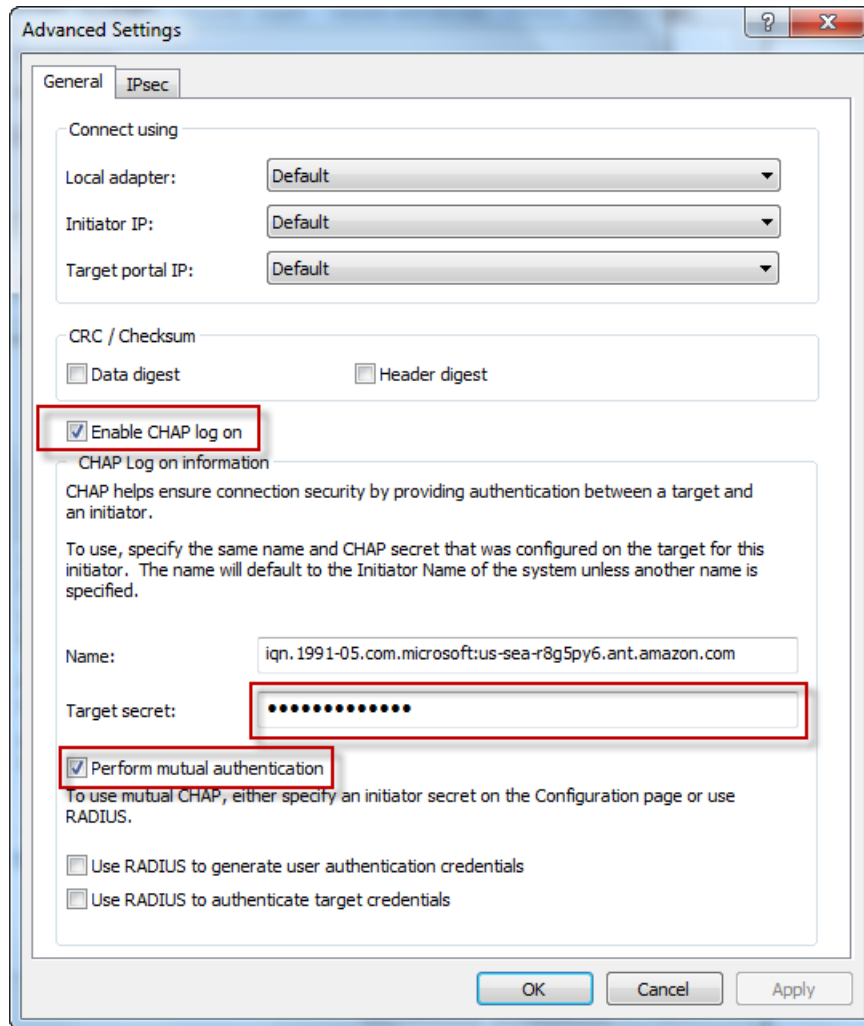


- d. In the **Connect to Target** dialog box, click **Advanced**.



- e. In the **Advanced Settings** dialog box, configure CHAP.

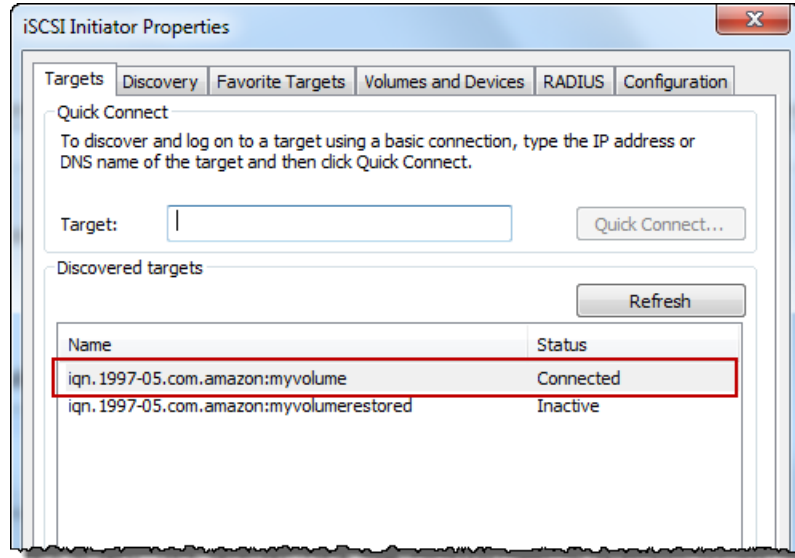
AWS Storage Gateway User Guide
Configuring CHAP Authentication for Your Storage Volume



- i. Select **Enable CHAP log on**.
 - ii. Enter the secret that is required to authenticate the initiator. This secret maps to the **Secret used to Authenticate Initiator** field in the **Configure CHAP Authentication** dialog box. For more information, see [Configuring CHAP Authentication for Your Storage Volume \(p. 167\)](#).
 - iii. Select **Perform mutual authentication**.
 - iv. Click **OK** to apply the changes.
- f. In the **Connect to Target** dialog box, click **OK**.
4. If you provided the correct secret key, the target will show a status of **Connected**.

AWS Storage Gateway User Guide

Configuring CHAP Authentication for Your Storage Volume



The following procedure assumes that the iSCSI daemon is running and that you have already connected to a target. If you have not completed these two tasks, see [Connecting from a Red Hat Client to Your Storage Volume](#) (p. 165).

To configure mutual CHAP on a Red Hat Linux client

In this procedure, you configure CHAP in the Linux iSCSI initiator using the same keys that you used to configure CHAP for the storage volume in the console.

1. Disconnect and remove any existing configuration for the target for which you are about to configure CHAP.
 - a. List the saved configurations to find the target name and ensure it is a defined configuration.

```
sudo /sbin/iscsiadm --mode node
```

- b. Disconnect from the target.

The following command disconnects from the target named *myvolume* that is defined on the Amazon IQN. Change the target name and IQN as required for your situation.

```
sudo /sbin/iscsiadm --mode node --logout GATEWAY_IP:3260,1 iqn.1997-05.com.amazon:myvolume
```

- c. Remove the configuration for the target.

The following command removes the configuration for the *myvolume* target.

```
sudo /sbin/iscsiadm --mode node --op delete --targetname iqn.1997-05.com.amazon:myvolume
```

2. Edit the iSCSI configuration file to enable CHAP.

AWS Storage Gateway User Guide

Configuring CHAP Authentication for Your Storage Volume

- a. Get the name of the initiator (the client you are using).

The following command gets the initiator name from the `/etc/iscsi/initiatorname.iscsi` file.

```
sudo cat /etc/iscsi/initiatorname.iscsi
```

The output from this command will look like this:

```
InitiatorName=iqn.1994-05.com.redhat:8e89b27b5b8
```

- b. Open the `/etc/iscsi/iscsid.conf` file.
- c. Uncomment the following lines in the file and specify the correct username and passwords (secret keys).

```
node.session.auth.authmethod = CHAP
node.session.auth.username = username
node.session.auth.password = password
node.session.auth.username_in = username_in
node.session.auth.password_in = password_in
```

Fill in the preceding items above using the following table as guidance.

Configuration Setting	Value
username	Use the initiator name that you found in a previous step in this procedure. The value will start with "iqn". For example, <code>iqn.1994-05.com.redhat:8e89b27b5b8</code> is a valid <code>username</code> .
password	This is the secret key used to authenticate the initiator (the client you are using) when it communicates with the storage volume.
username_in	Use the IQN of the target storage volume. The value will start with "iqn" and end with the target name. For example, <code>iqn.1997-05.com.amazon:myvolume</code> is a valid <code>username_in</code> .
password_in	This is the secret key used to authenticate the target (the storage volume) when it communicates to the initiator.

- d. Save the changes in the configuration file and close the file.
3. Discover and log into the target.

You can follow the steps provided in [Connecting from a Red Hat Client to Your Storage Volume \(p. 165\)](#) to discover and log into the target.

Managing Your Activated Gateway

Topics

- [Managing Storage Volumes in AWS Storage Gateway](#) (p. 176)
- [Managing the Upload Buffer and Cache Storage \(Gateway-Cached\)](#) (p. 187)
- [Configuring the Upload Buffer \(Gateway-Stored\)](#) (p. 193)
- [Working with Snapshots](#) (p. 199)
- [Performing Maintenance Tasks in AWS Storage Gateway](#) (p. 222)
- [Troubleshooting in AWS Storage Gateway](#) (p. 252)
- [Optimizing AWS Storage Gateway Performance](#) (p. 260)
- [Monitoring Your AWS Storage Gateway](#) (p. 261)
- [Related Section](#) (p. 276)

In this section, we review how you can manage your AWS Storage Gateway after you have deployed and activated it. Management tasks you will perform with your gateway include configuring storage volumes and upload buffer space, working with snapshots, general maintenance, troubleshooting, and monitoring your gateway. If you have not set up a gateway, see [Setting Up AWS Storage Gateway](#) (p. 90).

Managing Storage Volumes in AWS Storage Gateway

Topics

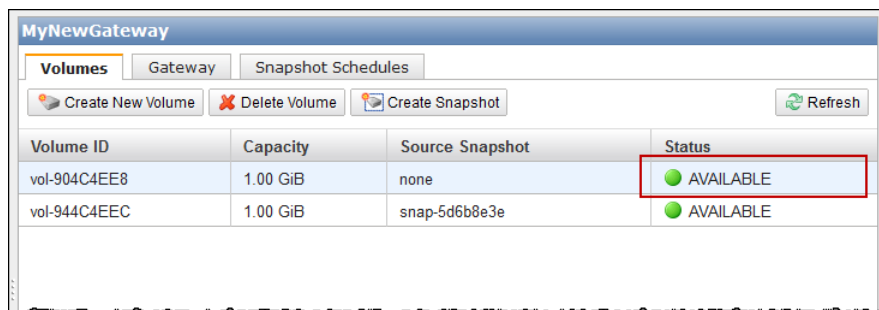
- [Understanding Storage Volume Status](#) (p. 177)
- [Managing Storage Volumes \(Gateway-Cached\)](#) (p. 181)
- [Managing Storage Volumes \(Gateway-Stored\)](#) (p. 183)

This section discusses how to manage your existing storage volumes including viewing the status of a volume, adding new volumes, and removing existing volumes.

Understanding Storage Volume Status

Each storage volume has an associated status that tells you at a glance, the health of the volume. Most of the time the status will indicate that the volume is functioning normally and no action is needed on your part. In a few cases, the status indicates a problem with the volume that may or may not require action on your part. This section helps you decide when you need to take action.

Volume status can be viewed in the console or by using one of the AWS Storage Gateway API operations, for example, see [DescribeCachediSCSIVolumes \(p. 348\)](#) or [DescribeStorediSCSIVolumes \(p. 363\)](#). The following example shows volume status in the AWS Storage Gateway console. Volume status appears in the **Status** field for each storage volume on your gateway. In the example, the volume highlighted is functioning normally because the status is AVAILABLE.



The following table describes each storage volume status and if, and when, you should take action. The AVAILABLE status is the normal status of a volume, and a volume should have this status all or the majority of the time it is in use.

Status	Description
AVAILABLE	The normal running status for a volume.
BOOTSTRAPPING	<p>This status means that the gateway is synchronizing data locally with a copy of the data stored in AWS. You typically do not need to take any action for this status as the storage volume will automatically go to AVAILABLE in most cases.</p> <p>The following are three scenarios when a volume status is BOOTSTRAPPING.</p> <ul style="list-style-type: none"> • A storage volume can be bootstrapping after an unexpected gateway shutdown, when a gateway's upload buffer capacity is exceeded, or when you create a gateway-stored volume and choose to preserve local disk data. • A gateway's upload buffer being exceeded. In this scenario, bootstrapping occurs when your volume is in PASS THROUGH and the amount of free upload buffer increases sufficiently. You can provide additional upload buffer as one way to increase the percentage of free upload buffer space. In this particular scenario, the storage volume goes from PASS THROUGH to BOOTSTRAPPING to AVAILABLE. You can continue to use this volume during this bootstrapping period; however, you cannot take snapshots. • You are creating a gateway-stored volume and preserving existing local disk data. In this scenario, your gateway starts uploading all of the data up to AWS and the volume remains in BOOTSTRAPPING until all of the data from the local disk is copied to AWS. You can use this volume during this bootstrapping period; however, you cannot take snapshots.

Status	Description
CREATING	The volume is currently being created and is not ready to be used. This is a transitional status; no action is required.
DELETING	The volume is currently being deleted. No action is required.
IRRECOVERABLE	An error occurred from which the volume cannot recover. For information on taking action in this situation, see Troubleshooting Storage Volume Issues (p. 256) .
PASS THROUGH	<p>This status means that data maintained locally is out of sync with data stored in AWS. This volume status can occur for several reasons.</p> <ul style="list-style-type: none">• One reason that can cause PASS THROUGH status is if your gateway has run out of upload buffer. Your applications can continue to read from and write data to your storage volumes while they are in PASS THROUGH; however, the gateway is not writing any of your volume data to its upload buffer and not uploading any of this data to AWS. The gateway will continue to upload any data written to the volume before entering the PASS THROUGH status. Any pending or scheduled snapshots of storage volume will fail while it's in PASS THROUGH mode. For information about what action to take when your storage volume is in PASS THROUGH because upload buffer is exceeded, see Troubleshooting Storage Volume Issues (p. 256).• Another reason for a volume to indicate the PASS THROUGH status is because there is more than one storage volume bootstrapping at once. Only one gateway storage volume can bootstrap at a time. For example, if you create two storage volumes and choose to preserve existing data on both of them, then the second storage volume will have the PASS THROUGH status until the first storage volume finishes bootstrapping. In this scenario, you do not need to take action. Each storage volumes will change to the AVAILABLE status automatically when it is finished being created. You can read and write to the storage volume while it is in PASS THROUGH or BOOTSTRAPPING.• Infrequently, the PASS THROUGH status can indicate that a disk allocated for upload buffer disk failed. For information about what action to take in this scenario, see Troubleshooting Storage Volume Issues (p. 256).
RESTORING	<p>The volume is being restored from an existing snapshot. This status applies only for stored-volumes in the gateway-stored volume setup (see How AWS Storage Gateway Works (p. 3)).</p> <p>If you restore two storage volumes at the same time, both storage volumes will show RESTORING as their status. Each storage volume will change to the AVAILABLE status automatically when it is finished being created. You can read and write to a storage volume and take a snapshot of it while it is RESTORING.</p>

Status	Description
RESTORING_PASS_THROUGH	<p>The volume is being restored from an existing snapshot and encountered an upload buffer issue. This status applies only for stored-volumes in the gateway-stored volume setup (see How AWS Storage Gateway Works (p. 3)).</p> <p>One reason that can cause the RESTORING_PASS_THROUGH status is if your gateway has run out of upload buffer space. Your applications can continue to read from and write data to your storage volumes while they are in RESTORING_PASS_THROUGH; however, no snapshots of the storage volume can occur in RESTORING_PASS_THROUGH. For information about what action to take when your storage volume is in RESTORING_PASS_THROUGH because upload buffer capacity is exceeded, see Troubleshooting Storage Volume Issues (p. 256).</p> <p>Infrequently, the RESTORING_PASS_THROUGH status can indicate that a disk allocated for an upload buffer has failed. For information about what action to take in this scenario, see Troubleshooting Storage Volume Issues (p. 256).</p>
UPLOAD_BUFFER_NOT_CONFIGURED	<p>The volume cannot be created or used because the gateway does not have an upload buffer configured. To add upload buffer capacity for a cached-volume gateway, see Adding and Removing Upload Buffer Capacity (Gateway-Cached) (p. 190). To add upload buffer capacity for a stored-volume gateway, see Configuring the Upload Buffer (Gateway-Stored) (p. 193).</p>

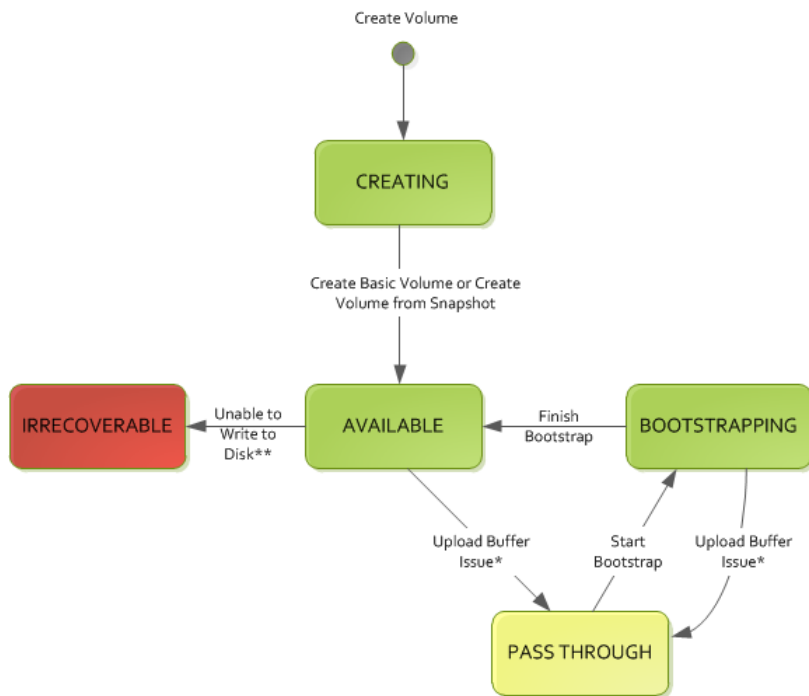
Cached Volume Status Transitions (Gateway-Cached)

The following state diagram describes the most common transitions between gateway-cached volume statuses. It is not necessary for you to understand the diagram in detail to use your gateway effectively. Rather, the diagram provides detailed information if you are interested in understanding more about how AWS Storage Gateway works.


The diagram shows neither the UPLOAD_BUFFER_NOT_CONFIGURED status nor the DELETING status. Volume states in the diagram are represented by green, yellow, and red boxes. The color of the boxes are interpreted as follows.

- **Green**—The gateway is operating normally. The volume status is AVAILABLE or will eventually become AVAILABLE.
- **Yellow**—Yellow (PASS THROUGH) indicates that there is a potential issue with the storage volume. If this status is because upload buffer space is filled, then in some cases more buffer space may become available again. At that point, the storage volume self-corrects and becomes AVAILABLE. In other cases, you may have to add more upload buffer space to your gateway to allow the storage volume status to become AVAILABLE. To troubleshoot when upload buffer capacity is exceeded, see [Troubleshooting Storage Volume Issues \(p. 256\)](#). To add upload buffer capacity, see [Adding and Removing Upload Buffer Capacity \(Gateway-Cached\) \(p. 190\)](#).
- **Red**—The storage volume has become IRRECOVERABLE. In this case, you should delete the volume (see [To remove a storage volume \(p. 183\)](#)).

In the diagram, a transition between two states is depicted with a labeled line. For example, the transition from the CREATING status to the AVAILABLE status is labeled as *Create Basic Volume or Create Volume from Snapshot* and represents creating a cached volume. For more information about creating storage volume, see [Adding a Storage Volume \(p. 182\)](#).



Key



* e.g. run out of upload buffer

** e.g. lost connectivity

Note

The volume status of PASS THROUGH is depicted as yellow in this diagram and does not match the color of this status icon in the **Status** field of the AWS Storage Gateway console.

Stored Volume Status Transitions (Gateway-Stored)

The following state diagram describes the most common transitions between gateway-stored volume statuses. It is not necessary for you to understand the diagram in detail to use your gateway effectively. Rather, the diagram provides detailed information if you are interested in understanding more about how AWS Storage Gateway works.

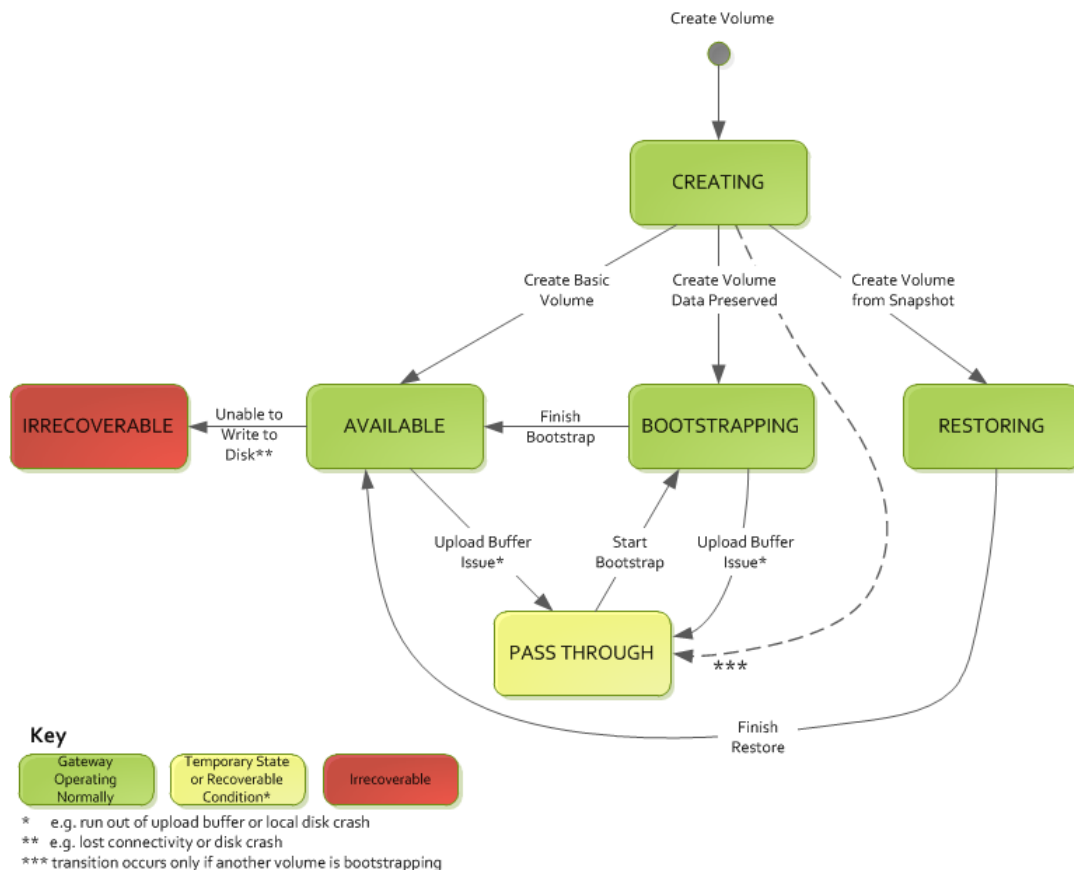
The diagram shows neither the UPLOAD BUFFER NOT CONFIGURED status nor the DELETING status. Volume states in the diagram are represented by green, yellow, and red boxes. The colors are interpreted as follows.

- **Green**—The gateway is operating normally. The volume status is AVAILABLE or will eventually become AVAILABLE.
- **Yellow**—When you are creating a storage volume and preserving data, then the path from CREATING to PASS THROUGH occurs if another volume is bootstrapping. In this case, the volume in PASS THROUGH will go to BOOTSTRAPPING and to AVAILABLE when the first volume is finished bootstrapping. Other than the specific scenario mentioned, yellow (PASS THROUGH) indicates that there is a potential issue with the storage volume, the most common being an upload buffer issue. If upload buffer capacity is exceeded, then in some cases more capacity may become available again. At that point, the storage volume self-corrects and becomes AVAILABLE. In other cases, you may have to add more upload buffer capacity to your gateway to allow the storage volume status to become AVAILABLE. To troubleshoot when upload buffer capacity is exceeded, see [Troubleshooting Storage](#)

[Volume Issues](#) (p. 256). To add upload buffer capacity, see [Configuring the Upload Buffer \(Gateway-Stored\)](#) (p. 193).

- **Red**—The storage volume has become IRRECOVERABLE. In this case, you should delete the volume (see [To remove the underlying local disk \(VMware ESXi\)](#) (p. 185)).

In the following diagram, a transition between two states is depicted with a labeled line. For example, the transition from the CREATING status to the AVAILABLE status is labeled as *Create Basic Volume* and represents creating a storage volume without preserving data or creating from a snapshot. For more information about creating storage volume, see [To create a storage volume using the console](#) (p. 159).



Note

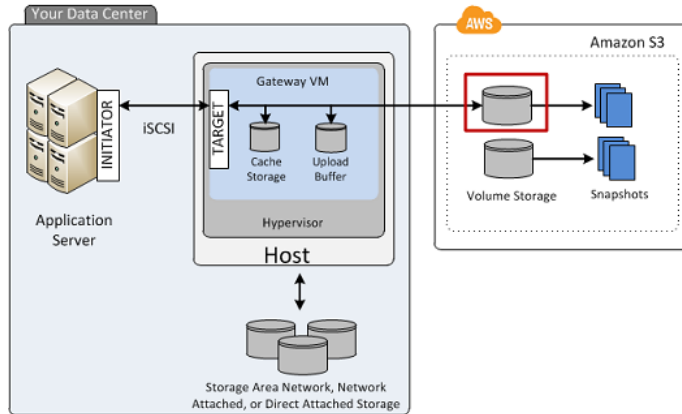
The volume status of PASS THROUGH is depicted as yellow in this diagram and does not match the color of this status icon in the **Status** field of the AWS Storage Gateway console.

Managing Storage Volumes (Gateway-Cached)

Topics

- [Adding a Storage Volume](#) (p. 182)
- [Removing a Storage Volume](#) (p. 183)

Gateway-cached volumes are storage volumes in Amazon S3 that are exposed as iSCSI targets on which you can store your application data. The following diagram highlights storage volumes in a gateway-cached setup (see [How AWS Storage Gateway Works](#) (p. 3)).



Your gateway exposes each volume as iSCSI target with a name you specify, prepended by `iqn.1997-05.com.amazon:`. For example, if you specify a target name of `myvolume`, then the iSCSI target you use to connect to the volume is `iqn.1997-05.com.amazon:myvolume`. For more information about how to configure your applications to mount volumes over iSCSI (see [Configuring Your Application Access to Storage Volumes](#) (p. 161)).

Note

Storage volumes can range from 1 GiB to 32 TiB in size and must be rounded to the nearest GiB. Each gateway can support up to 20 cached volumes and a total cache volume storage of 150 TiB.

Resizing a storage volume is not supported. To change the size of a volume, create a snapshot of the volume, and create a new cached volume from the snapshot. The new volume can be bigger than the volume from which the snapshot was created. For steps describing how to remove a storage volume, see [To remove a storage volume](#) (p. 183). For steps describing how to add a storage volume and preserve existing data, see [To create a storage volume using the console](#) (p. 157).

Important

Since a cached volume keeps your primary data in Amazon S3, you should avoid processes that read or write all data on the entire volume. For example, we strongly recommend against using virus scanning software that scans the entire cached volume. Such a scan, whether on-demand or scheduled, will cause all data stored in Amazon S3 to be downloaded locally for scanning, which results in high bandwidth usage and a dirty cache. Instead of doing a full disk scan, you can use real-time virus scanning—that is, scanning data as it is read from or written to the cached volume.

All gateway-cached volume data and snapshot data is stored in Amazon S3 encrypted at rest using Server Side Encryption (SSE). However, you cannot access this data using Amazon S3 APIs or with other tools such as the Amazon S3 console.

Adding a Storage Volume

You created a storage volume as part of your initial gateway setup, for example, see the [Getting Started with AWS Storage Gateway](#) (p. 7). As your application needs grow, you might need to add more storage volumes to your gateway. As you add more storage volumes, you must consider the size of your cache storage and upload buffer you allocated to the gateway. The gateway must have sufficient buffer and cache space for new volumes. For more information, see [Managing the Upload Buffer and Cache Storage \(Gateway-Cached\)](#) (p. 187).

You can add storage volumes using the AWS Storage Gateway API (see [CreateCachediSCSIVolume](#) (p. 318)) or the console. The following task demonstrates using the console and assumes that you already have a deployed and activated gateway.

To create a storage volume using the console

- Follow the steps in [Creating a Storage Volume \(Gateway-Cached\)](#) (p. 157) that you used to create your initial volume.

Removing a Storage Volume

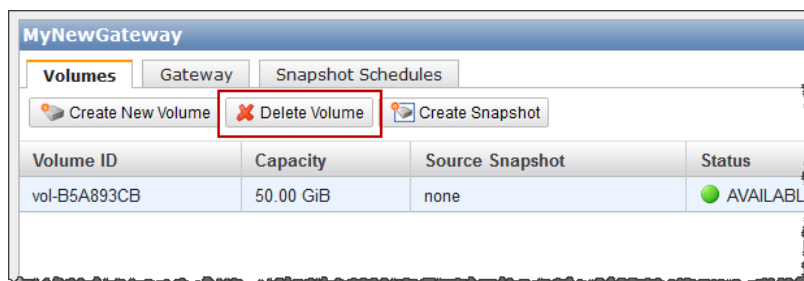
You might need to remove a storage volume as your application needs change, for example, if you migrate your application to use a larger storage volume. Before removing a storage volume, make sure that there are no applications currently writing to the volume. Also, make sure there are no snapshots in progress for the volume. If a snapshot schedule is defined for the volume, you can check it on the **Snapshot Schedules** tab of the console. For more information, see [Editing a Snapshot Schedule](#) (p. 208).

Before removing a storage volume, make sure that there are no applications currently writing to the volume. Also, make sure there are no snapshots in progress for the volume. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the console. For more information, see [Editing a Snapshot Schedule](#) (p. 208). Perform the following task to remove a storage volume with the gateway running.

You can remove storage volumes using the AWS Storage Gateway API (see [DeleteVolume](#) (p. 340)) or the console. The following task demonstrates using the console. Perform the task with the gateway running.

To remove a storage volume

1. In the AWS Storage Gateway console, in the **Volumes** tab, select the storage volume.



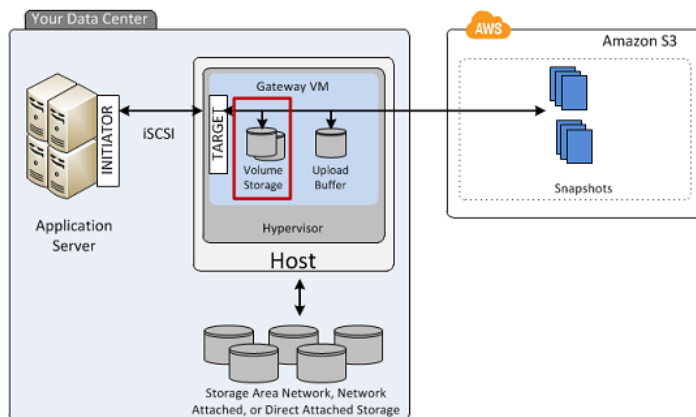
2. Click **Delete Volume**.

Managing Storage Volumes (Gateway-Stored)

Topics

- [Adding a Storage Volume](#) (p. 184)
- [Removing a Storage Volume](#) (p. 185)

Gateway-stored volumes are storage volumes that are exposed as iSCSI targets on which you can store your application data. The storage volumes are created on the local virtual disks that you added to your gateway VM (see [Provisioning Local Disks \(Gateway-Stored\)](#) (p. 102)). The following diagram highlights storage volumes in a gateway-stored setup (see [How AWS Storage Gateway Works](#) (p. 3)).



Your gateway exposes each volume as iSCSI target with a name you specify, prepended by `iqn.1997-05.com.amazon:`. For example, if you specify a target name of `myvolume`, then the iSCSI target you use to connect to the volume is `iqn.1997-05.com.amazon:myvolume`. For more information about how to configure your applications to mount volumes over iSCSI (see [Configuring Your Application Access to Storage Volumes \(p. 161\)](#)). Your gateway stores your application data locally on your storage volume's disk, while asynchronously uploading your data to AWS.

Note

Storage volumes can range from 1 GiB to 1 TiB in size. The size depends on the local disk from which you create the storage volume. The local disk should be rounded to the nearest GiB (see [Adding Local Disks for Volume Storage \(Gateway-Stored\) \(p. 102\)](#)). Each gateway can support up to 12 volumes and up to 12 TiB of local storage.

Resizing the underlying disk of a storage volume is not supported. To change the size of an underlying disk, delete the storage volume that is using the disk, resize the disk, and then create a new storage volume from the resized disk. When you recreate the storage volume, be sure to preserve the data on the disk. For steps describing how to remove a storage volume, see [To remove the underlying local disk \(VMware ESXi\) \(p. 185\)](#) or [To remove the underlying local disk \(Microsoft Hyper-V\) \(p. 187\)](#). For steps describing how to add a storage volume and preserve existing data, see [To create a storage volume using the console \(p. 159\)](#).

Adding a Storage Volume

You create a storage volume as part of your initial gateway setup, for example, see the [Getting Started with AWS Storage Gateway \(p. 7\)](#). As your application needs grow, you might need to add more storage volumes to your gateway. As you add more storage volumes, you must consider the size of your upload buffer you allocated to the gateway. The gateway must have sufficient buffer space. For more information, see [Managing the Upload Buffer \(Gateway-Stored\) \(p. 194\)](#).

You can add storage volumes using the AWS Storage Gateway API (see [CreateStoreiSCSIVolume \(p. 327\)](#)) or the console. The following task demonstrates using the console and assumes that you already have a deployed and activated gateway. Furthermore, it is assumed that there is at least one locally provisioned disk of the gateway that is not used and can be allocated as a storage volume. To provision a local disk for application storage, see [Provisioning Local Disks \(Gateway-Stored\) \(p. 102\)](#).

To add a storage volume

- Follow the steps in [Creating a Storage Volume \(Gateway-Stored\) \(p. 159\)](#).

Removing a Storage Volume

You might need to remove a storage volume as your application needs change, for example, if you migrate your application to use a larger storage volume and you want to reclaim the underlying, local disk space of the old storage volume. To reclaim the local disk space you need to remove the local disk from the VM.


Before removing a storage volume, make sure that there are no applications currently writing to the volume. Also, make sure there are no snapshots in progress for the volume. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the console. For more information, see [Editing a Snapshot Schedule \(p. 208\)](#).

You can remove storage volumes using the AWS Storage Gateway API (see [DeleteVolume \(p. 340\)](#)) or the console. The following task demonstrates using the console and either the vSphere client for a gateway deployed on the VMware ESXi platform or the Microsoft Hyper-V Manager for a gateway deployed on the Microsoft Hyper-V platform. Perform the following task to remove a storage volume with the gateway running.

To remove a storage volume in the AWS Storage Gateway console

1. In the AWS Storage Gateway console, in the **Volumes** tab, select the storage volume.
2. If you plan to remove the disk from the VM that backs the storage volume, then in the **Details** properties tab, note the value in the **Local Disk** field.

This value is the disk's **Virtual Device Node** value that you use in the hypervisor client to ensure that you remove the correct disk.



3. Click **Delete Volume**.
4. If you want to remove the underlying local disk do one of the following:

For a Gateway Hosted In...	Do This...
VMware ESXi	Follow the steps in To remove the underlying local disk (VMware ESXi) (p. 185) .
Microsoft Hyper-V	Follow the steps in To remove the underlying local disk (Microsoft Hyper-V) (p. 187) .

To remove the underlying local disk (VMware ESXi)

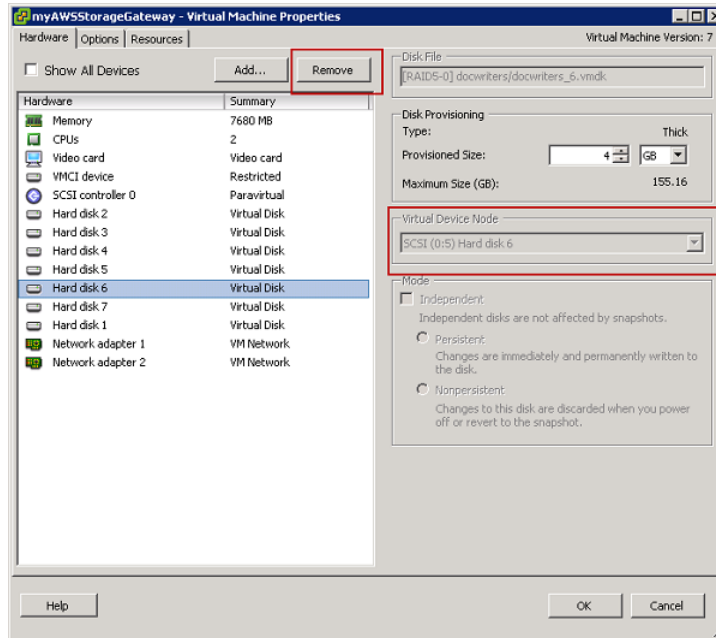
1. In the vSphere client, right-click the name of your gateway VM and click **Edit Settings...**

AWS Storage Gateway User Guide

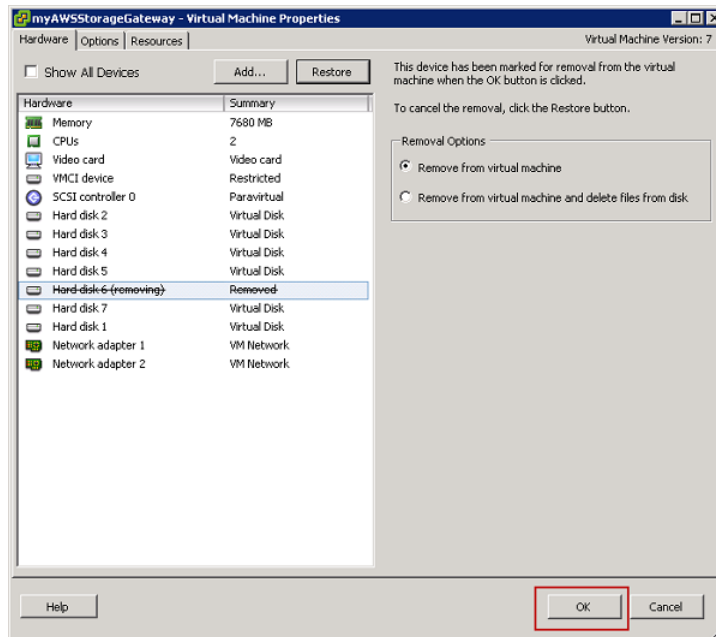
Managing Storage Volumes (Gateway-Stored)

2. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, select the disk to remove and click **Remove**.

Verify that the **Virtual Device Node** value in the **Virtual Machine Properties** dialog box has the same value that you noted from a previous step. This ensures that you remove the correct disk. The first SCSI controller displayed in the Microsoft Hyper-V Manager is controller 0.



3. Choose an option in the **Removal Options** panel, and click **OK** to complete the process of removing the disk.



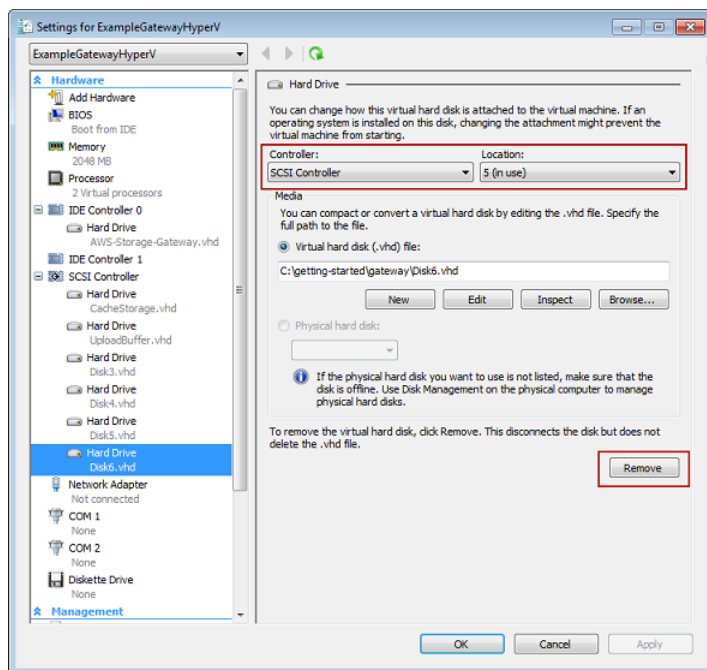
To remove the underlying local disk (Microsoft Hyper-V)

1. In the Microsoft Hyper-V Manager, right-click the name of your gateway VM and click **Settings....**
2. In the **Hardware** list of the **Settings** dialog box, select the disk to remove and click **Remove**.

The disks you add to a gateway are under the **SCSI Controller** entry in the **Hardware** list.

Verify that the **Controller** and **Location** value are the same value that you noted from a previous step. This ensures that you remove the correct disk.

The first SCSI controller displayed in the Microsoft Hyper-V Manager is controller 0.



3. Click **OK** to apply the change.

Managing the Upload Buffer and Cache Storage (Gateway-Cached)

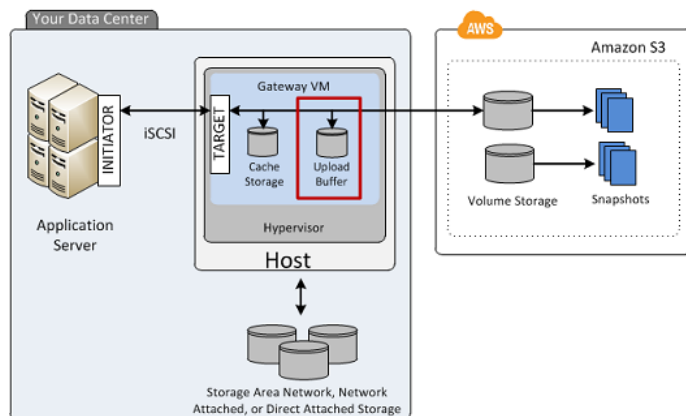
Topics

- [Managing the Upload Buffer \(Gateway-Cached\) \(p. 188\)](#)
- [Managing Cache Storage \(Gateway-Cached\) \(p. 189\)](#)
- [Adding and Removing Upload Buffer Capacity \(Gateway-Cached\) \(p. 190\)](#)
- [Adding Cache Storage \(Gateway-Cached\) \(p. 193\)](#)

In your initial deployment, you configure upload buffer and cache storage for your gateway. As you add storage volumes to provide more storage for your application data, you might need to add more upload buffer or cache storage capacity to the gateway. You might also need to remove a local disk allocated as upload buffer, for example, because you want to replace a local disk that has failed. This section reviews how to determine if you need to add more upload buffer or cache storage and how to do it.

Managing the Upload Buffer (Gateway-Cached)

Your gateway uses the upload buffer to temporarily buffer your volume data prior to uploading it to AWS. The following diagram highlights the upload buffer in the larger picture of the AWS Storage Gateway gateway-cached architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).



The amount of upload buffer space your gateway requires depends on several factors such as the rate of incoming data to the storage volumes, the rate of outgoing data to AWS, and your network bandwidth. If your applications continue to write data at a fast rate to your storage volumes, and network throughput is not sufficient for the gateway to upload data to AWS, then eventually your upload buffer will be filled with data waiting to be uploaded to AWS. Here is some guidance you can follow to avoid this situation:

- **Use the Sizing Formula**—As your application needs change, you should periodically review the recommended formula for sizing the upload buffer. For more information, see [Sizing the Upload Buffer \(Gateway-Cached\) \(p. 98\)](#).
- **Use Amazon CloudWatch Metrics**— You can proactively avoid the upload buffer from filling up by monitoring the percentage of upload buffer space your gateway is using in time. Amazon CloudWatch provides usage metrics such as the `UploadBufferPercentUsed` metric for monitoring your gateway's upload buffer (see [Monitoring the Upload Buffer \(p. 267\)](#)). You can set a threshold to trigger a notification to you when upload buffer usage exceeds a threshold. If your upload buffer is getting filled close to capacity, consider adding more buffer capacity to the gateway. For a full list of AWS Storage Gateway metrics, see [Understanding AWS Storage Gateway Metrics \(p. 272\)](#). For a full list of AWS Storage Gateway metrics, see [Understanding AWS Storage Gateway Metrics \(p. 272\)](#).

You can see the current percentage of upload buffer usage in the **Gateway** tab of the AWS Storage Gateway console.

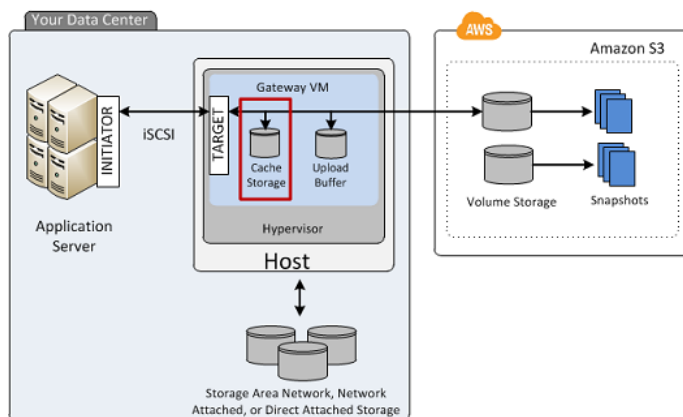
- **Monitor Volume Status**— A volume status can indicate an issue with the upload buffer. For example, if the upload buffer capacity is reached, the impacted storage volumes go into a **PASS THROUGH (p. 179)** mode, and while your applications can continue to operate, writing and reading data to and from your storage volume, snapshots are not taken during this time.
- **Optimize Your Environment**—If the speed of incoming writes is too high compared to the outgoing network bandwidth, then the gateway might never be able to catch up, no matter how much upload buffer capacity you provision. In this case, consider optimizing your gateway for better performance (see [Optimizing AWS Storage Gateway Performance \(p. 260\)](#)).

If you decide that you need to change your upload buffer capacity, take one of the following actions.

To...	Do This...
Add more upload buffer capacity to your gateway.	Follow the steps in Adding Upload Buffer Capacity (p. 190) .
Remove a disk allocated as upload buffer space.	Follow the steps in Removing Upload Buffer Capacity (p. 190) .

Managing Cache Storage (Gateway-Cached)

Your gateway uses cache storage to cache recently accessed application data. The following diagram highlights the cache storage in the larger picture of the AWS Storage Gateway gateway-cached architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).



The amount of cache storage your gateway requires depends on how much of your application data you want to provide low-latency access to. The cache storage must be at least the size of the upload buffer. This ensures the cache storage is large enough to be able to persistently hold all data that has not yet been uploaded to Amazon S3. When your cache storage has filled up with dirty data, application writes to your storage volume are blocked until more cache storage becomes available. Application reads from the storage volume, however, are still allowed. Here is some guidance you can follow to avoid this situation:

- **Use the Sizing Formula**—As your application needs change, you should periodically review the recommended formula for sizing cache storage. For more information, see [Sizing Cache Storage \(Gateway-Cached\) \(p. 94\)](#).
- **Use Amazon CloudWatch Metrics**— You can proactively avoid filling up cache storage with dirty data by monitoring how cache storage is being used—particularly, by reviewing cache misses. Amazon CloudWatch provides usage metrics such as the `CachePercentDirty` and `CacheHitPercent` metrics for monitoring how much of the gateway's cache storage has not been uploaded to Amazon S3. You can set a threshold to trigger a notification to you when cache percent dirty exceeds a threshold or cache hit percentage falls below a threshold, both potentially indicating that the cache storage is not adequate for the gateway. For a full list of AWS Storage Gateway metrics, see [Understanding AWS Storage Gateway Metrics \(p. 272\)](#).

If you decide that you need to increase your gateway's cache storage capacity, follow the steps in [Adding Cache Storage \(Gateway-Cached\) \(p. 193\)](#).

Adding and Removing Upload Buffer Capacity (Gateway-Cached)

After your initial gateway upload buffer configuration (see [Configuring Upload Buffer \(Gateway-Cached\)](#) (p. 150)), you can configure additional upload buffer capacity as your application needs change. To learn more about how to size your upload buffer based on your application needs, see [Sizing the Upload Buffer \(Gateway-Cached\)](#) (p. 98).

You can add more buffer capacity to your gateway without interrupting existing gateway functions. Note that when you add more upload buffer capacity, you do so with the gateway VM powered on; however, when you reduce the amount of upload buffer capacity, you must first power off the VM.

Adding Upload Buffer Capacity

As your application needs change and you add more storage volume capacity, you will might need to increase the gateway's upload buffer capacity. You can add more buffer capacity using the AWS Storage Gateway API (see [AddUploadBuffer](#) (p. 313)) or the AWS Storage Gateway console. The following procedure shows how to add more buffer capacity using the console. It assumes that your activated gateway has at least one local disk available on its VM that you can allocate as an upload buffer to the gateway.

To configure a local disk as upload buffer space for your gateway using the console

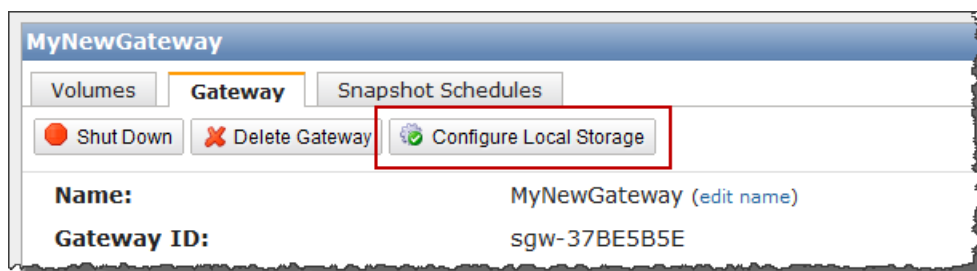
- Follow the steps in [To configure a local disk as an upload buffer for your gateway](#) (p. 151).

Removing Upload Buffer Capacity

As your application needs change and you change the storage volume configuration for a gateway, you might need to decrease the gateway's upload buffer capacity. Or, a local disk allocated as upload buffer space might fail and you need to remove that disk from your upload buffer and assign a new local disk. In both cases, you can remove more buffer capacity using the AWS Storage Gateway console. The following procedure assumes that your activated gateway has at least one local disk allocated as an upload buffer to the gateway. In the procedure you start in the AWS Storage Gateway console, leave the console and use the VMware vSphere client or the Microsoft Hyper-V Manager to remove the disk, and then return to the console.

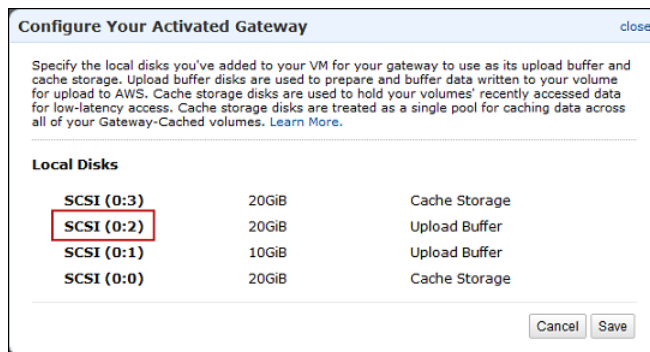
To use the console to find the disk ID of a disk allocated as an upload buffer

- In the AWS Storage Gateway console, in the **Gateway** tab, click **Configure Local Storage**.



- In the **Configure Your Activated Gateway** dialog box, note the value of the virtual device node for the local disk to be removed. You can find the node value in the **Local Disks** column. For example, in the following dialog box, the device node **SCSI (0:2)** is highlighted.

You use the disk's virtual device node in the vSphere client to ensure that you remove the correct disk.



- Shut down the gateway by following the steps in the [Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console](#) (p. 224) procedure.

Note

Before shutting down the gateway, ensure that it is not in use by an application that is writing data to it and that no snapshots are in progress. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the console. For more information, see [Editing a Snapshot Schedule](#) (p. 208).

- To remove the underlying local disk do one of the following and then go to the next step:

For a Gateway Hosted In...	Do This...
VMware ESXi	Follow the steps in To remove the underlying disk allocated as an upload buffer disk (VMware ESXi) (p. 191).
Microsoft Hyper-V	Follow the steps in To remove the underlying disk allocated as an upload buffer disk (Microsoft Hyper-V) (p. 192).

- In the AWS Storage Gateway console, turn on the gateway.

Important

After removing a disk used as an upload buffer, you must first turn the gateway back on before adding new disks to the VM.

- In the AWS Storage Gateway console, in the **Volumes** tab, check that all storage volumes have a status of [AVAILABLE](#) (p. 179).

After a gateway restart, a storage volume may go through the [PASS THROUGH](#) (p. 179) and [BOOTSTRAPPING](#) (p. 179) states as the gateway adjusts to the upload buffer disk that you removed. A storage volume that passes through these two states will eventually come to the [AVAILABLE](#) (p. 179) state. You can use a storage volume during the pass through and bootstrapping states; however, you cannot take snapshots.

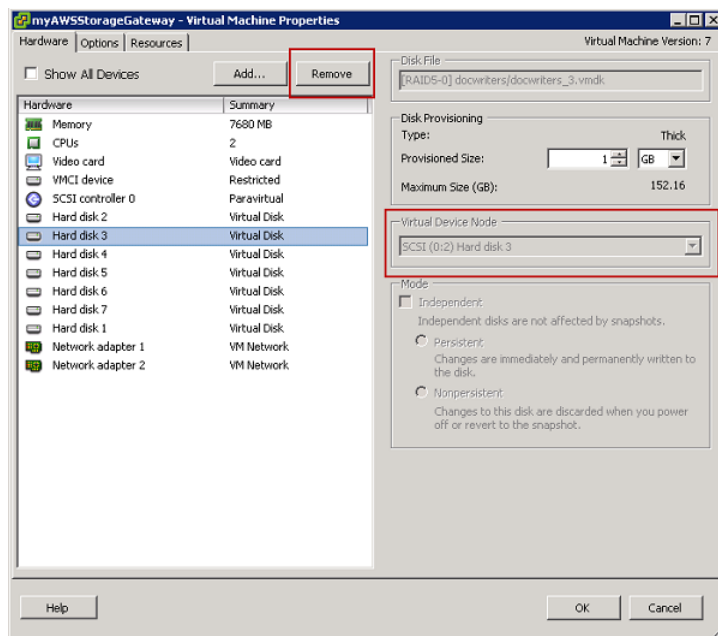
To remove the underlying disk allocated as an upload buffer disk (VMware ESXi)

- In the vSphere client, right-click the name of your gateway VM and click **Edit Settings...**
- In the **Hardware** tab of the **Virtual Machine Properties** dialog box, select the disk allocated as upload buffer space, and click **Remove**.

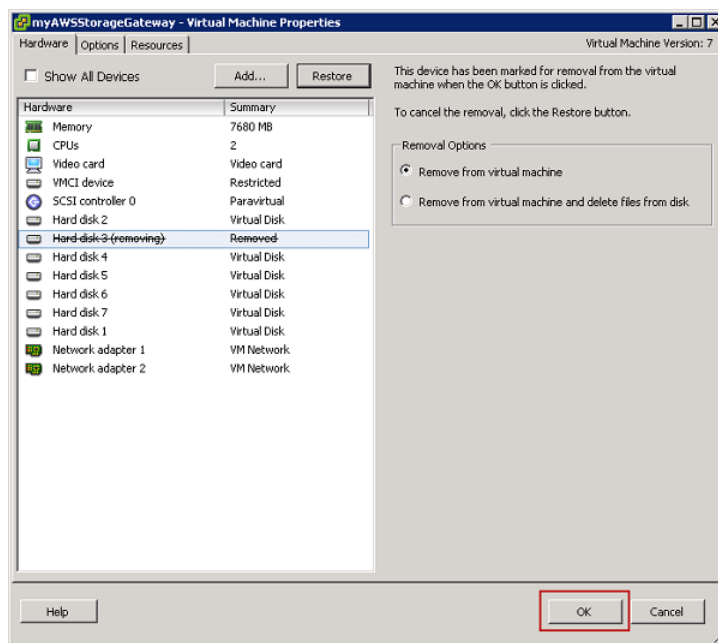
Verify that the **Virtual Device Node** value in the **Virtual Machine Properties** dialog box has the same value that you noted from a previous step. This ensures you remove the correct disk.

AWS Storage Gateway User Guide

Adding and Removing Upload Buffer Capacity



3. Choose an option in the **Removal Options** panel and click **OK** to complete the process of removing the disk.

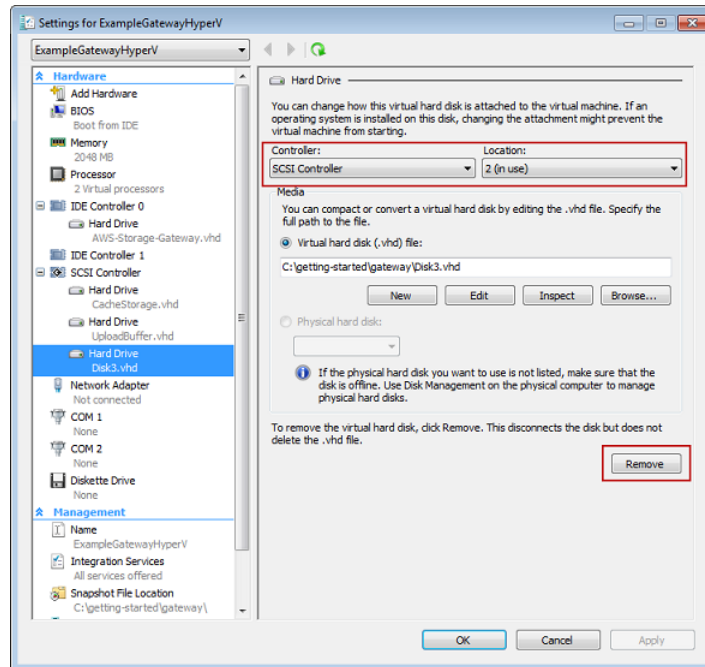


To remove the underlying disk allocated as an upload buffer disk (Microsoft Hyper-V)

1. In the Microsoft Hyper-V Manager, right-click the name of your gateway VM and click **Settings....**
2. In the **Hardware** list of the **Settings** dialog box, select the disk to remove and click **Remove**.

The disks you add to a gateway are under the **SCSI Controller** entry in the **Hardware** list. Verify that the **Controller** and **Location** value are the same value that you noted from a previous step. This ensures that you remove the correct disk.

The first SCSI controller displayed in the Microsoft Hyper-V Manager is controller 0.



3. Click **OK** to apply the change.

Adding Cache Storage (Gateway-Cached)

After your initial gateway cache storage configuration (see [Configuring Cache Storage \(Gateway-Cached\)](#) (p. 152)), you can configure additional cache storage to your gateway as your application needs change. To learn more about how to size your cache storage based on your application needs, see [Sizing Cache Storage \(Gateway-Cached\)](#) (p. 94).

You can add more cache storage to your gateway without interrupting existing gateway functions and with the gateway VM powered on.

Note

Removing a disk allocated as cache storage is currently not supported.

You can add more cache storage using the AWS Storage Gateway API (see [AddCache](#) (p. 310)) or the AWS Storage Gateway console. The following procedure assumes that your activated gateway has at least one local disk available on its VM that you can allocate as cache storage to the gateway.

To configure a local disk as cache storage for your gateway

- Follow the steps in [To configure a local disk as cache storage for your gateway](#) (p. 153).

Configuring the Upload Buffer (Gateway-Stored)

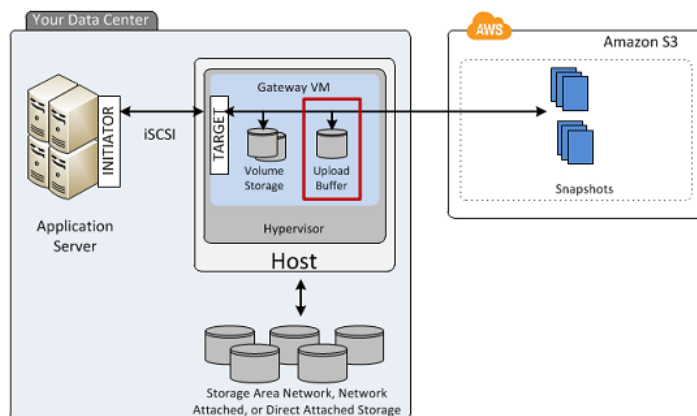
Topics

- [Managing the Upload Buffer \(Gateway-Stored\)](#) (p. 194)
- [Adding and Removing Upload Buffer Capacity \(Gateway-Stored\)](#) (p. 195)

As part of your initial deployment, you configure upload buffer for your gateway. As you add storage volumes to provide more storage for your application data, you might need to add more upload buffer capacity to the gateway. You might also need to remove a local disk allocated as upload buffer, for example, because you want to replace a local disk that has failed. This section reviews how to determine if you need to add more upload buffer capacity and how to do it.

Managing the Upload Buffer (Gateway-Stored)

Your gateway uses the upload buffer to temporarily buffer your volume data prior to uploading it to AWS. The following diagram highlights the upload buffer in the larger picture of the AWS Storage gateway-stored architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).



The amount of upload buffer space that is required by your gateway depends on several factors such as the rate of incoming data to the storage volumes, the rate of outgoing data to AWS, and your network bandwidth. If your applications continue to write data at a fast rate to your storage volumes, and network throughput is not sufficient for the gateway to upload data to AWS, then eventually your upload buffer will be filled with data waiting to be uploaded to AWS. Here is some guidance you can follow to avoid this situation:

- **Use the Sizing Formula**—As your application needs change, you should periodically review the recommended formula for sizing the upload buffer. For more information, see [Sizing the Upload Buffer \(Gateway-Stored\) \(p. 106\)](#).
- **Use Amazon CloudWatch Metrics**— You can proactively avoid the upload buffer from filling up by monitoring the percentage of upload buffer space your gateway is using in time. Amazon CloudWatch provides usage metrics such as the `UploadBufferPercentUsed` metric for monitoring your gateway's upload buffer usage (see [Monitoring the Upload Buffer \(p. 267\)](#)). You can set a threshold to trigger a notification to you when upload buffer usage exceeds a threshold. If your upload buffer is getting filled close to capacity, consider adding more buffer capacity to the gateway.

You can see the current value of the upload buffer percentage usage in the **Gateway** tab of the AWS Storage Gateway console.

- **Monitor Volume Status**— If the upload buffer capacity is reached, the impacted storage volumes go into a **PASS THROUGH** (p. 179) mode and while your applications can continue to operate, writing and reading data to and from your storage volume, snapshots are not taken during this time.
- **Optimize Your Environment**—If the speed of incoming writes is too high compared to the outgoing network bandwidth, then the gateway might never be able to catch up, no matter how much upload buffer capacity you provision. In this case, you should then consider optimizing your gateway for better performance (see [Optimizing AWS Storage Gateway Performance \(p. 260\)](#)).

To...	Do This...
Add more upload buffer capacity to your gateway.	Follow the steps in Adding Upload Buffer Capacity (Gateway-Stored) (p. 195) .
Remove a disk allocated as upload buffer.	Follow the steps in Removing Upload Buffer Capacity (Gateway-Stored) (p. 195) .

Adding and Removing Upload Buffer Capacity (Gateway-Stored)

After your initial gateway upload buffer configuration (see [Configuring Upload Buffer \(Gateway-Stored\) \(p. 154\)](#)), you can configure additional upload buffer capacity as your application needs change. To learn more about how to size your upload buffer based on your application needs, see [Sizing the Upload Buffer \(Gateway-Stored\) \(p. 106\)](#).

You can add more buffer capacity to your gateway without interrupting existing gateway functions. Note that when you add more upload buffer capacity, you do so with the gateway VM powered on; however, when you reduce the amount of upload buffer capacity, you must first power off the VM.

Adding Upload Buffer Capacity (Gateway-Stored)

As your application needs change and you add more storage volume capacity, you might need to increase the gateway's upload buffer capacity. You can add more buffer capacity using the AWS Storage Gateway API (see [AddUploadBuffer \(p. 313\)](#)) or the AWS Storage Gateway console. The following procedure shows how to add more buffer capacity using the console. It assumes that your activated gateway has at least one local disk available on its VM that you can allocate as an upload buffer to the gateway.

To configure an upload buffer for your gateway using the console

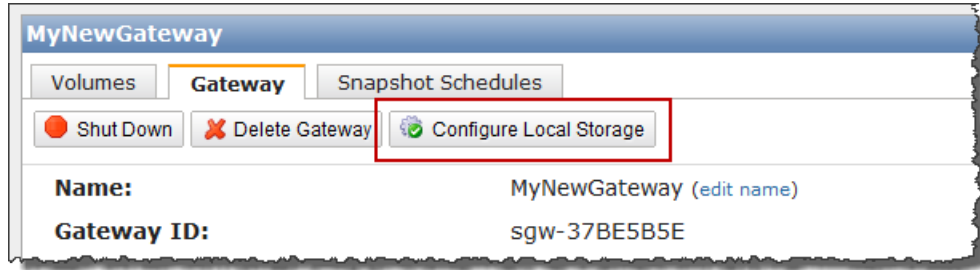
- Follow the steps in [To configure a local disk as an upload buffer for your gateway \(p. 155\)](#).

Removing Upload Buffer Capacity (Gateway-Stored)

As your application needs change and you change the storage volume configuration for a gateway, you might need to decrease the gateway's upload buffer capacity. Or, a local disk allocated to upload buffer might fail and you need to remove that disk from the upload buffer and assign a new local disk. In both cases, you can remove more buffer capacity using the AWS Storage Gateway console. The following procedure assumes that your activated gateway has at least one local disk allocated as an upload buffer to the gateway. In the procedure you start in the AWS Storage Gateway console, leave the console and use the VMware vSphere client or the Microsoft Hyper-V Manager to remove the disk, and then return to the console.

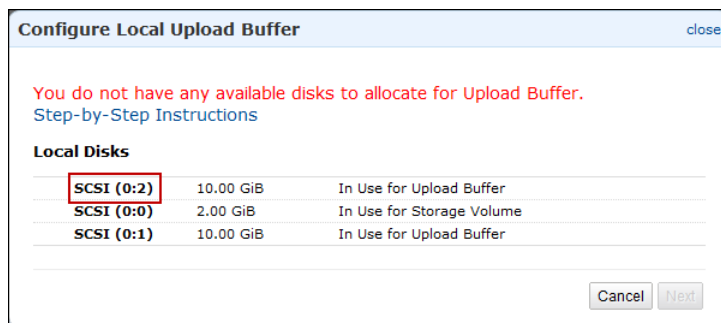
To use the console to find the disk ID of a disk allocated as an upload buffer

1. In the AWS Storage Gateway console, in the **Gateway** tab, click **Configure Local Storage**.



- In the **Configure Local Upload Buffer** dialog box, note the value of the virtual device node for the local disk to be removed. You can find the node value in the **Local Disks** column. For example, in the following dialog box, the device node **SCSI (0:2)** is highlighted.

You use the disk's virtual device node in the hypervisor client to ensure that you remove the correct disk.



- Shut down the gateway by following the steps in the [Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console](#) (p. 224) procedure.

Note

Before shutting down the gateway, ensure that it is not in use by an application that is writing data to it and that no snapshots are progress. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the console. For more information, see [Editing a Snapshot Schedule](#) (p. 208).

- To remove the underlying local disk, do one of the following and then go to the next step:

For a Gateway Hosted In...	Do This...
VMware ESXi	Follow the steps in To remove the underlying disk disk allocated as an upload buffer (VMware ESXi) (p. 197).
Microsoft Hyper-V	Follow the steps in To remove the underlying disk disk allocated as an upload buffer (Microsoft Hyper-V) (p. 198).

- In the AWS Storage Gateway console, turn on the gateway.

Important

After removing a disk used as an upload buffer, you must turn the gateway back on before adding new disks to the VM.

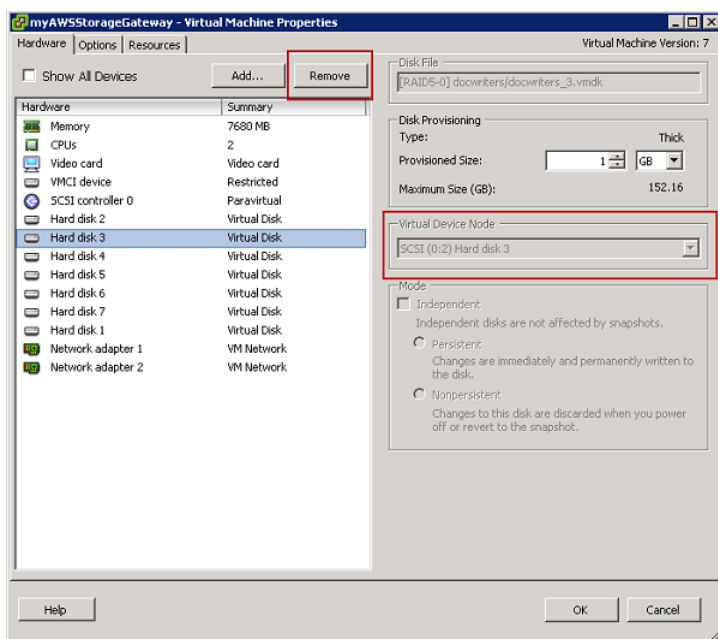
- In the AWS Storage Gateway console, in the **Volumes** tab, check that all storage volumes have a status of **AVAILABLE** (p. 179).

After a gateway restart, a storage volume may go through the [PASS THROUGH](#) (p. 179) and [BOOTSTRAPPING](#) (p. 179) states as the gateway adjusts to the upload buffer disk that you removed. A storage volume that passes through these two states will eventually come to the [AVAILABLE](#) (p. 179) state. You can use a storage volume during the pass through and bootstrapping states; however, you cannot take snapshots.

To remove the underlying disk allocated as an upload buffer (VMware ESXi)

1. In the vSphere client, right-click the name of your gateway VM and click **Edit Settings...**
2. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, select the disk allocated as an upload buffer, and click **Remove**.

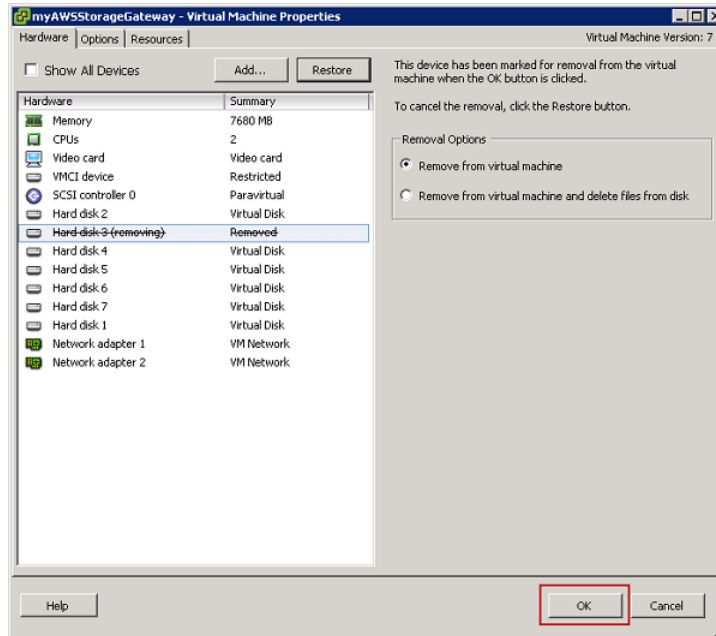
Verify that the **Virtual Device Node** value in the **Virtual Machine Properties** dialog box has the same value that you noted from a previous step. This ensures you remove the correct disk.



3. Choose an option in the **Removal Options** panel, and click **OK** to complete the process of removing the disk.

AWS Storage Gateway User Guide

Adding and Removing Upload Buffer Capacity

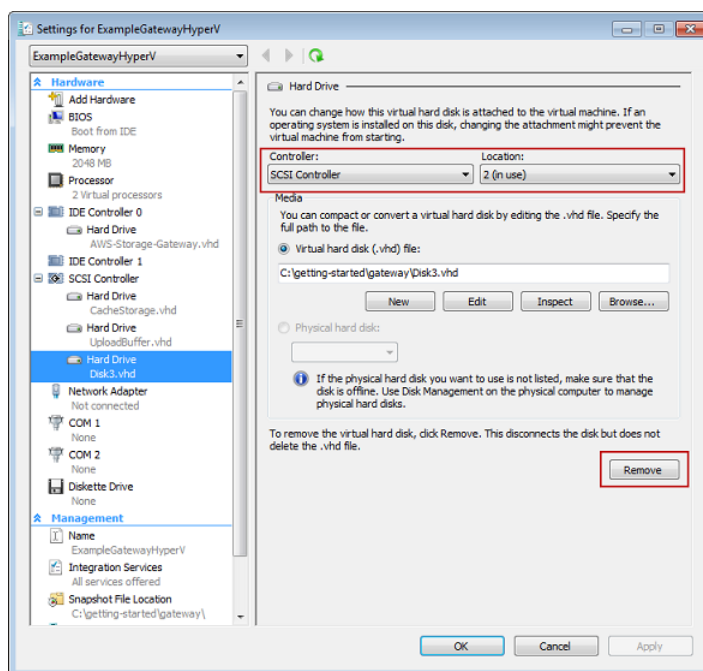


To remove the underlying disk disk allocated as an upload buffer (Microsoft Hyper-V)

1. In the Microsoft Hyper-V Manager, right-click the name of your gateway VM and click **Settings....**
2. In the **Hardware** list of the **Settings** dialog box, select the disk to remove and click **Remove**.

The disks you add to a gateway are under the **SCSI Controller** entry in the **Hardware** list. Verify that the **Controller** and **Location** value are the same value that you noted from a previous step. This ensures that you remove the correct disk.

The first SCSI controller displayed in the Microsoft Hyper-V Manager is controller 0.



3. Click **OK** to apply the change.

Working with Snapshots

Topics

- [Overview \(p. 199\)](#)
- [Finding a Snapshot \(p. 200\)](#)
- [Editing a Snapshot Schedule \(p. 208\)](#)
- [Creating an Ad-Hoc Snapshot \(p. 209\)](#)
- [Deleting a Snapshot \(p. 209\)](#)
- [Restoring a Snapshot \(p. 219\)](#)

Overview

AWS Storage Gateway provides the ability to back up point-in-time snapshots of your data to Amazon S3 for durable recovery that can be used later on-premises or in Amazon EC2. You can take snapshots on a scheduled or ad-hoc basis. In this section, we show you the most common tasks that you can perform with snapshots including creating a snapshot and restoring the snapshot to a volume which can then be mounted as an iSCSI device, and restoring a snapshot to an Amazon EBS volume, which can then be attached to an Amazon EC2 instance.

AWS Storage Gateway continually and asynchronously uploads data to AWS to keep your local data synchronized with a copy stored in AWS. A benefit of this is that when snapshots are initiated, some or all the data has already been uploaded and snapshots complete quickly. Furthermore, snapshots are incremental—that is, the gateway uploads only the blocks of your volume that have changed since the last snapshot. For example, if you have 100 GiB of data and only 5 GiB data changed since the last snapshot, then the gateway uploads only the 5 GiB of changed data. You can delete any snapshot. AWS Storage Gateway removes only the snapshot data that is not needed by other snapshots, enabling you to restore a volume from any of the active snapshots.

How snapshots can be effectively used in your AWS Storage Gateway setup depends on the type of gateway you set up—that is, a gateway-cached or gateway-stored architecture (see [How AWS Storage Gateway Works \(p. 3\)](#))

- For gateway-cached volumes, your volume data is already stored in Amazon S3, so snapshots can be used to preserve older versions of your data.
- For gateway-stored volumes, your volume data is stored on-premises, so snapshots provide durable, off-site backups in Amazon S3.

Summary of Snapshot Tasks

Since snapshots are key to using the AWS Storage Gateway service, you should understand at a high level what each snapshot operation does and why it is done. Each task is covered in detail in the linked section.

To work with the tasks, you should have one or more gateways that have been running for enough time so there are snapshots to work with. You can work with snapshots using the AWS Storage Gateway console, an AWS Software Development Kit (SDK), or the AWS Storage Gateway REST API (see [Operations in AWS Storage Gateway \(p. 305\)](#)). In this section, we primarily show how to work with the console to perform gateway tasks.

Snapshot Action	Common Scenarios
Finding	You might want to find a snapshot to see if it is complete, what time it was taken, what the size of the snapshot is, or the name of the volume the snapshot was taken from. For more information, see Finding a Snapshot (p. 200) .
Scheduling	When you first set up a stored volume, a default snapshot schedule of once per day is set. You can change the frequency and timing of the snapshot schedule to fit your application needs. For more information, see Editing a Snapshot Schedule (p. 208) .
Creating	Snapshots for stored volumes are automatically created on a schedule by default and you can change the schedule of the snapshots. However, you can take an instantaneous snapshot at any time for both stored and cached volumes. For more information, see Creating an Ad-Hoc Snapshot (p. 209) .
Restoring	You can restore the snapshot locally to a new AWS Storage Gateway volume, or you can use the snapshot to create an Amazon Elastic Block Store (EBS) volume and attach that to an Amazon EC2 instance. For more information, see Restoring a Snapshot to an AWS Storage Gateway Volume (p. 219) and Restoring a Snapshot to an Amazon EBS Volume (p. 222) .
Deleting	If you don't need a snapshot anymore, you can delete it. Since snapshots are incremental backups, the deletion process is such that if you delete a snapshot, only the data that is not needed in other snapshots is deleted. For more information, see Deleting a Snapshot (p. 209) .

Snapshot Consistency

Snapshots provide a point-in-time view of data that has been written to your AWS Storage Gateway volumes. However, snapshots only capture data that has been written to your storage volumes, which may exclude any data that has been buffered by your client application or OS. Your application and OS will eventually flush this buffered data to your storage volumes. If you need to guarantee that your application data is flushed to disk prior to taking a snapshot, you should consult your specific application's documentation to understand if and how your application buffers data and how to flush this data. If you need to guarantee that your OS and file system have flushed their buffered data to disk prior to taking a snapshot, you can do this by taking your storage volume offline before taking a snapshot. This forces your OS to flush its data to disk. After the snapshot is complete, you can bring the volume back online. In Windows, use Disk Management (`diskmgmt.msc`) to select the storage volume and take it online or offline. To script this process in Windows, you can use a command line tool such as [Diskpart.exe](#). In Linux, use the `mount` and `umount` commands.

Finding a Snapshot

You will need to find a snapshot associated with a volume when you want to restore the snapshot to a new volume—for example, in a disaster recovery scenario—or to restore a previous version of your application data. To restore a snapshot, you need to know its snapshot ID. There are several ways you can find a snapshot ID including using the AWS Storage Gateway console, the Amazon Elastic Compute Cloud (EC2) console, or programmatically using one of the AWS Software Development Kits (SDKs).

If you list snapshots using the AWS Storage Gateway console, the list includes all your snapshots generated from your gateway and snapshots that you might have generated from Amazon Elastic Block Store (EBS) volumes. If you list snapshots in the Amazon EC2 console, more snapshot properties are shown to help you find your snapshot as well as a search filtering capability. Both console experiences are described below. In some scenarios, you might need to search using several snapshot properties at once—for example, status, start date, and description. In this case, you can use a programmatic approach.

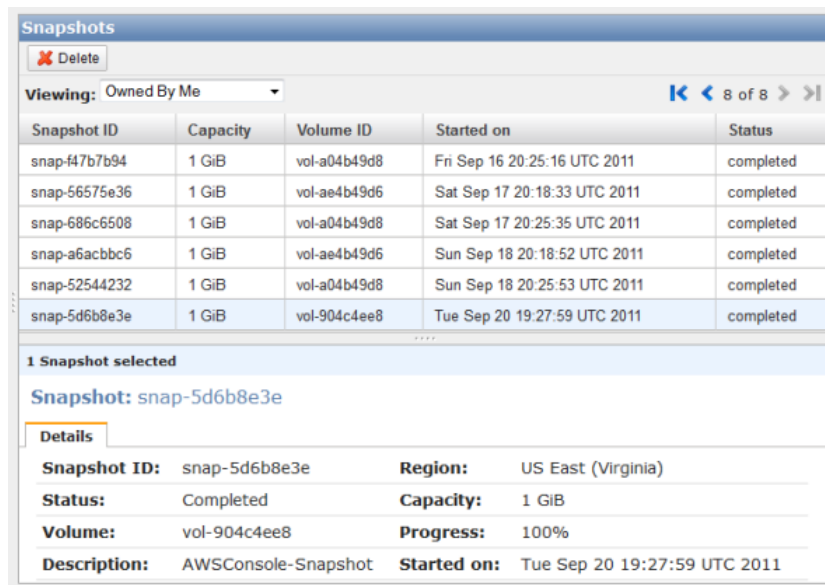
For examples, see [Finding Snapshots Using the AWS SDK for Java \(p. 202\)](#), [Finding Snapshots Using the AWS SDK for .NET \(p. 204\)](#), or [Finding Snapshots Using the AWS Tools for Windows PowerShell \(p. 206\)](#).

When you find your snapshot, you can view its details, including date and time the snapshot was started and the storage volume on your gateway that was the source for the snapshot.

To find a snapshot for a volume using the AWS Storage Gateway console

1. On the **AWS Storage Gateway** console, in the **Navigation** pane, click **Snapshots**.

The **Snapshots** window shows a list of your snapshots.



Snapshot ID	Capacity	Volume ID	Started on	Status
snap-f47b7b94	1 GiB	vol-a04b49d8	Fri Sep 16 20:25:16 UTC 2011	completed
snap-56575e36	1 GiB	vol-ae4b49d6	Sat Sep 17 20:18:33 UTC 2011	completed
snap-686c6508	1 GiB	vol-a04b49d8	Sat Sep 17 20:25:35 UTC 2011	completed
snap-a6acbbc6	1 GiB	vol-ae4b49d6	Sun Sep 18 20:18:52 UTC 2011	completed
snap-52544232	1 GiB	vol-a04b49d8	Sun Sep 18 20:25:53 UTC 2011	completed
snap-5d6b8e3e	1 GiB	vol-904c4ee8	Tue Sep 20 19:27:59 UTC 2011	completed

1 Snapshot selected

Snapshot: snap-5d6b8e3e

Details

Snapshot ID:	snap-5d6b8e3e	Region:	US East (Virginia)
Status:	Completed	Capacity:	1 GiB
Volume:	vol-904c4ee8	Progress:	100%
Description:	AWSConsole-Snapshot	Started on:	Tue Sep 20 19:27:59 UTC 2011

2. Find the snapshot that you are looking for in the list by looking for the volume ID in the **Volume ID** column.
3. Click the snapshot row to display the snapshot details.

To find a snapshot for a volume using the Elastic Block Store console

1. On the **Amazon EC2 Management** console, in the **Navigation** pane, click **Snapshots** under **ELASTIC BLOCK STORE**.

The **EBS Snapshots** window shows a list of your snapshots. This snapshot view offers more functionality for finding snapshots. In particular, this view shows descriptions that can be useful for filtering results—for instance, there is a pattern to snapshot descriptions that you can use to help find a snapshot. For example, in the console image below

- The row marked with 1 is a snapshot taken of an EBS volume. The name and description fields are specified in the snapshot creation. This snapshot is not from an AWS Storage Gateway operation.
- The row marked as 2 is an ad-hoc snapshot taken from the AWS Storage Gateway console. The description for ad-hoc snapshots contains "AWSConsole-Snapshot".
- The row marked as 3 is a snapshot of a volume taken from a snapshot schedule by AWS Storage Gateway. The description of scheduled snapshots gives the storage gateway ID, volume ID, and the word "Schedule" in the pattern "*gatewayID:volumeID:Schedule*".

Name	Snapshot ID	Capacity	Description	Status	Started
EBS Volume Snapshot	snap-8aa5c9fd	8 GiB	Description of EBS Snapshot	1 completed	2012-10-16 15:31 P
empty	snap-48a8c43f	1 GiB	AWSConsole-Snapshot	2 completed	2012-10-16 15:29 P
empty	snap-02a8c575	1 TiB	sgw-A7A346CE.vol-42404B6D.Schedule	3 completed	2012-10-16 13:58 P
empty	snap-288f1c5f	1 TiB	AWSConsole-Snapshot	completed	2012-10-15 16:07 P
empty	snap-db891aac	1 GiB	AWSConsole-Snapshot	completed	2012-10-15 16:06 P
empty	snap-de31a0a9	1 TiB	sgw-A7A346CE.vol-42404B6D.Schedule	completed	2012-10-15 13:58 P
empty	snap-fbe0718c	1 GiB	AWSConsole-Snapshot	completed	2012-10-15 12:57 P

- Find the snapshot that you are looking for in the list by typing some or all of the volume ID into the **Search** field.

In the following example, only results containing "vol-424" are shown. You can find the volume ID in the AWS Storage Gateway console in the **Volumes** tab.

Name	Snapshot ID	Capacity	Description	Status	Started	Progre
empty	snap-02a8c575	1 TiB	sgw-A7A346CE.vol-42404B6D.Schedule	completed	2012-10-16 13:58 PDT	availab
empty	snap-de31a0a9	1 TiB	sgw-A7A346CE.vol-42404B6D.Schedule	completed	2012-10-15 13:58 PDT	availab
empty	snap-a42baed3	1 TiB	sgw-A7A346CE.vol-42404B6D.Schedule	completed	2012-10-14 13:58 PDT	availab
empty	snap-b911b8ce	1 TiB	sgw-A7A346CE.vol-42404B6D.Schedule	completed	2012-10-13 13:58 PDT	availab
empty	snap-a75588d0	1 TiB	sgw-A7A346CE.vol-42404B6D.Schedule	completed	2012-10-12 13:58 PDT	availab
empty	snap-f717d680	1 TiB	sgw-A7A346CE.vol-42404B6D.Schedule	completed	2012-10-11 13:58 PDT	availab

- Click the snapshot row to display the snapshot details.

Finding Snapshots Using the AWS SDK for Java

You can use a programmatic approach to quickly find snapshots and filter the results returned using snapshot properties such as snapshot status, description, and the date the snapshot was initiated. The following example demonstrates how to find snapshots using the AWS SDK for Java using several snapshot properties. To use the example code, you should be familiar with running a Java console application. For more information, see [Getting Started](#) in the *AWS SDK for Java Developer Guide*.

Example : Finding Snapshots Using the AWS SDK for Java

The following Java code example finds snapshots for a specified volume of a gateway using several properties of the snapshot to filter the results returned. It uses the AWS SDK for Java and API for Amazon Elastic Compute Cloud (EC2). The Amazon EC2 API includes operations for working with snapshots.

You need to update the code and provide the service endpoint, a full or partial volume ID, a snapshot status, and the number of days to indicate a cutoff date for the snapshots returned. For a list of AWS service endpoints you can use with Amazon EC2, see [Regions and Endpoints](#) in the *Amazon Web Services Glossary*.

```
import java.io.IOException;
import java.util.Calendar;
import java.util.Date;
import java.util.GregorianCalendar;
import java.util.List;

import com.amazonaws.AmazonClientException;
import com.amazonaws.auth.PropertiesCredentials;
import com.amazonaws.services.ec2.AmazonEC2Client;
import com.amazonaws.services.ec2.model.DescribeSnapshotsRequest;
import com.amazonaws.services.ec2.model.DescribeSnapshotsResult;
import com.amazonaws.services.ec2.model.Filter;
import com.amazonaws.services.ec2.model.Snapshot;

public class FindingSnapshotsExample {

    static AmazonEC2Client ec2Client;
    // A full volume id or partial fragment with "*".
    static String volumeID = "vol-424*";
    // Snapshot status to filter on: "completed", "pending", "error".
    static String status = "completed";
    // The number of days before which to not return snapshot results.
    static int daysBack = 10;
    // Service end point. Should be same region as volume/gateway.
    public static String serviceURLEC2 = "https://ec2.us-east-1.amazonaws.com";

    public static void main(String[] args) throws IOException {

        ec2Client = new AmazonEC2Client(new PropertiesCredentials(
            FindingSnapshotsExample.class.getResourceAsStream("AwsCreden
tials.properties")));
        ec2Client.setEndpoint(serviceURLEC2);

        FindingSnapshotsForAVolume();

    }

    private static void FindingSnapshotsForAVolume() {

        try {
            Filter[] filters = new Filter[2];
            filters[0] = new Filter().withName("volume-id").withValues(volumeID);

            filters[1] = new Filter().withName("status").withValues(status);
```

```
DescribeSnapshotsRequest describeSnapshotsRequest =
    new DescribeSnapshotsRequest().withFilters(filters);
DescribeSnapshotsResult describeSnapshotResult =
    ec2Client.describeSnapshots(describeSnapshotsRequest);

List<Snapshot> snapshots = describeSnapshotResult.getSnapshots();
System.out.println("volume-id = " + volumeID);
for (Snapshot s : snapshots) {
    if (CompareDates(daysBack, s.getStartTime())) {
        StringBuilder sb = new StringBuilder();
        sb.append(s.getSnapshotId() + ", " + s.getStartTime() + ",
" + s.getDescription());
        System.out.println(sb.toString());
    }
}
} catch (AmazonClientException ace) {
    System.err.println(ace.getMessage());
}
}

public static boolean CompareDates(int daysBack, Date snapshotDate) {
    Date today = new Date();
    Calendar cal = new GregorianCalendar();
    cal.setTime(today);
    cal.add(Calendar.DAY_OF_MONTH, -daysBack);
    Date cutoffDate = cal.getTime();
    return (snapshotDate.compareTo(cutoffDate) > 0) ? true : false;
}
}
```

Finding Snapshots Using the AWS SDK for .NET

You can use a programmatic approach to quickly find snapshots and filter the results returned using snapshot properties such as snapshot status, description, and the date the snapshot was initiated. The following example demonstrates how to find snapshots using the AWS SDK for .NET using several snapshot properties. To use the example code, you should be familiar with running a .NET console application. For more information, see [Getting Started](#) in the *AWS SDK for .NET Developer Guide*.

Example : Finding Snapshots Using the AWS SDK for .NET

The following C# code example finds snapshots for a specified volume of a gateway using several properties of the snapshot to filter the results returned. It uses the AWS SDK for .NET and API for Amazon Elastic Compute Cloud (EC2). The Amazon EC2 API includes operations for working with snapshots.

You need to update the code and provide the service endpoint, a full or partial volume ID, a snapshot status, and the number of days to indicate a cutoff date for the snapshots returned. For a list of AWS service endpoints you can use with Amazon EC2, see [Regions and Endpoints](#) in the *Amazon Web Services Glossary*.

```
using System;
using System.Text;
using System.Collections.Generic;
using Amazon.EC2;
using Amazon.EC2.Model;

namespace AWSStorageGateway
{
    class FindingSnapshotsExample
    {
        static AmazonEC2Config ec2Config;
        static AmazonEC2Client ec2Client;
        // A full volume id or partial fragment with "*".
        static String volumeID = "vol-424*";
        // Snapshot status to filter on: "completed", "pending", "error".
        static String status = "completed";
        // The number of days before which to not return snapshot results.
        static int daysBack = 4;
        // Service endpoint. Should be same region as volume/gateway.
        static String serviceURLEC2 = "https://ec2.us-east-1.amazonaws.com";

        public static void Main(string[] args)
        {
            //Create a ec2 client
            ec2Config = new AmazonEC2Config();
            ec2Config.ServiceURL = serviceURLEC2;
            ec2Client = new AmazonEC2Client(ec2Config);

            FindingSnapshotsForAVolume();

            Console.WriteLine("\nTo continue, press Enter.");
            Console.Read();
        }

        private static void FindingSnapshotsForAVolume()
        {
            try
            {
                Filter[] filters = new Filter[2];
                filters[0] = new Filter().WithName("volume-id").With
Value(volumeID);
                filters[1] = new Filter().WithName("status").WithValue(status);

                DescribeSnapshotsRequest describeSnapshotsRequest =
                    new DescribeSnapshotsRequest().WithFilter(filters);
            }
        }
    }
}
```

```
DescribeSnapshotsResponse describeSnapshotsResponse =
    ec2Client.DescribeSnapshots(describeSnapshotsRequest);

List<Snapshot> snapshots = describeSnapshotsResponse.Describe
SnapshotsResult.Snapshot;
Console.WriteLine("volume-id = " + volumeID);
foreach (Snapshot s in snapshots)
{
    if (CompareDates(daysBack, s.StartTime))
    {
        StringBuilder sb = new StringBuilder();
        sb.Append(s.SnapshotId + ", " + s.StartTime + ", " +
s.Description);
        Console.WriteLine(sb.ToString());
    }
}
catch (AmazonEC2Exception ex)
{
    Console.WriteLine(ex.Message);
}
}

public static Boolean CompareDates(int daysBack, String d)
{
    DateTime snapshotDate = DateTime.Parse(d);
    DateTime cutoffDate = DateTime.Now.Add(new TimeSpan(-daysBack, 0,
0, 0));
    return (DateTime.Compare(snapshotDate, cutoffDate) < 0) ? true :
false;
}
}
```

Finding Snapshots Using the AWS Tools for Windows PowerShell

You can use a programmatic approach to quickly find snapshots and filter the results returned using snapshot properties such as snapshot status, description, and the date the snapshot was initiated. The following example demonstrates how to find snapshots using the AWS Tools for Windows PowerShell using several snapshot properties. To use the example code, you should be familiar with running a PowerShell script. For more information, see [Getting Started](#) in the *AWS Tools for Windows PowerShell User Guide*.

Example : Finding Snapshots Using the AWS Tools for Windows PowerShell

The following PowerShell script example finds snapshots for a specified volume of a gateway using several properties of the snapshot to filter the results returned. It uses AWS Tools for Windows PowerShell cmdlets for Amazon Elastic Compute Cloud (Amazon EC2). The Amazon EC2 cmdlets include operations for working with snapshots.

You need to update the script and provide a full or partial volume ID, a snapshot status, and the number of days to indicate a cutoff date for the snapshots returned.

```
<#
.DESCRIPTION
    Finds snapshots for a given volume and criteria about the snapshot.

.NOTES
    PREREQUISITES:
    1) AWS Tools for PowerShell from http://aws.amazon.com/powershell/
    2) Credentials and region stored in session using Initialize-AWSDefault.
    For more info, see http://docs.aws.amazon.com/powershell/latest/userguide/specifying-your-aws-credentials.html

.EXAMPLE
    powershell.exe .\SG_FindSnapshots.ps1
#>

# Criteria to use to filter the results returned.
$volumeID = "vol-424*"
$status = "completed"
$daysBack = 4

# Define filters.
$filter1 = New-Object Amazon.EC2.Model.Filter
$filter1.Name = "volume-id"
$filter1.Value.Add($volumeID)

$filter2 = New-Object Amazon.EC2.Model.Filter
$filter2.Name = "status"
$filter2.Value.Add($status)

$snapshots = get-EC2Snapshot -Filter $filter1, $filter2
$count = 0

foreach ($s in $snapshots)
{
    $d = ([DateTime]::Now).AddDays(-$daysBack)
    if ([DateTime]::Compare($d, $s.StartTime) -gt 0)
    {
        # Meets criteria.
        $count +=1
        $sb = $s.SnapshotId + ", " + $s.StartTime + ", " + $s.Description
        Write-Output($sb)
    }
}
Write-Output ("Found " + $count + " snapshots that matched the criteria.")
```

Editing a Snapshot Schedule

For gateway-stored volumes, AWS Storage Gateway creates a default snapshot schedule of once a day. This ensures that your gateway can keep up with the rate of incoming writes on your local storage volumes. You can change the default snapshot schedule, for example, by specifying both the time the snapshot occurs each day, as well as the frequency (every 1, 2, 4, 8, 12, or 24 hours).

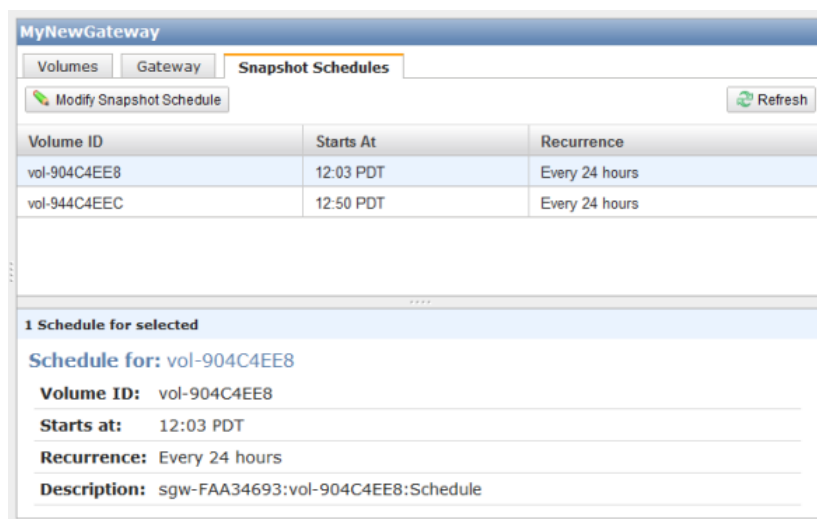
For gateway-cached volumes, AWS Storage Gateway does not create a default snapshot schedule; however, you can set up a snapshot schedule at any time if you need to. For gateway-cached volumes, since your data is stored in Amazon S3, the need for snapshots (and a snapshot schedule) for disaster recovery purposes is not needed.

In the following steps, we show you how to edit the snapshot schedule of a volume.

To edit snapshot schedule

1. In the **AWS Storage Gateway** console, select the gateway that contains the volume snapshot schedule that you want to edit.
2. Click the **Snapshot Schedules** tab.

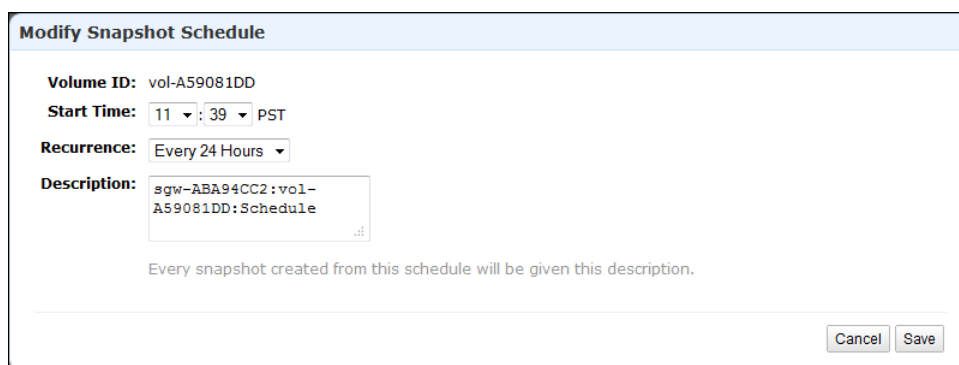
The tab shows a list of your storage volumes on the selected gateway.



3. Select a volume.

The AWS Storage Gateway console shows the snapshot schedule details for this volume.

4. Click **Modify Snapshot Schedule**.



5. In the **Modify Snapshot Schedule** dialog box, update the schedule fields as needed. For example, you can increase the default snapshot frequency of once a day or change the time.
6. Click **Save** to save the snapshot schedule updates.


Creating an Ad-Hoc Snapshot

In addition to scheduled snapshots, AWS Storage Gateway allows you to take ad-hoc snapshots, enabling you to back up your storage volume immediately without waiting for the next scheduled snapshot.

To take an ad-hoc snapshot of your storage volume

1. In the **AWS Storage Gateway** console, select the gateway that contains the storage volume of which you want to take a snapshot.
2. Click the **Volumes** tab.
3. Select a volume from the list and click **Create Snapshot**.

AWS Storage Gateway starts the snapshot process immediately.



4. Verify the snapshot at the console. For more information, see [Finding a Snapshot \(p. 200\)](#).

Deleting a Snapshot

Topics

- [Deleting Snapshots Using the AWS SDK for Java \(p. 211\)](#)
- [Deleting Snapshots Using the AWS SDK for .NET \(p. 214\)](#)
- [Deleting Snapshots Using the AWS Tools for Windows PowerShell \(p. 217\)](#)

You might want to delete a snapshot, for example, if you have taken many snapshots of a storage volume over a period of time and you don't need older snapshots. Since snapshots are incremental backups, the deletion process is such that if you delete a snapshot, only the data that is not needed in other snapshots is deleted.

In the AWS Storage Gateway console, you can delete a snapshots one at a time. To delete many snapshots, you use one of the AWS SDKs that supports AWS Storage Gateway operations. For examples, see [Deleting Snapshots Using the AWS SDK for Java \(p. 211\)](#), [Deleting Snapshots Using the AWS SDK for .NET \(p. 214\)](#), or [Deleting Snapshots Using the AWS Tools for Windows PowerShell \(p. 217\)](#).

To delete a snapshot using the console

1. In the AWS Storage Gateway console, click **Snapshots** in the **Navigation** pane.

A list of snapshots appears in the main pane.

The screenshot shows the 'Snapshots' page in the AWS Storage Gateway console. At the top, there is a 'Delete' button with a red 'X' icon. Below it, a 'Viewing:' dropdown menu is set to 'Owned By Me'. A table lists eight snapshots with columns for Snapshot ID, Capacity, Volume ID, Started on, and Status. The last row, 'snap-5d6b8e3e', is highlighted in blue. Below the table, a section titled '1 Snapshot selected' shows the details for 'Snapshot: snap-5d6b8e3e'. The details are as follows:

Details			
Snapshot ID:	snap-5d6b8e3e	Region:	US East (Virginia)
Status:	Completed	Capacity:	1 GiB
Volume:	vol-904c4ee8	Progress:	100%
Description:	AWSConsole-Snapshot	Started on:	Tue Sep 20 19:27:59 UTC 2011

2. Select the snapshot that you want to delete and click **Delete**.

This screenshot is identical to the previous one, but the 'Delete' button at the top left of the console is highlighted with a red rectangular box, indicating the next step in the process.

3. Clicking **OK** to confirm that you want to delete the snapshot.

Deleting Snapshots Using the AWS SDK for Java

To delete many snapshots associated with a volume, you can use a programmatic approach. The example below demonstrates how to delete snapshots using the AWS SDK for Java. To use the example code, you should be familiar with running a Java console application. For more information, see [Getting Started](#) in the *AWS SDK for Java Developer Guide*. If you need to just delete a few snapshots, use the console as described in [To delete a snapshot using the console](#) (p. 210).

Example : Deleting Snapshots Using the AWS SDK for Java

The following Java code example lists the snapshots for each volume of a gateway and whether the snapshot start time is before or after a specified date. It uses the AWS SDK for Java API for AWS Storage Gateway and Amazon Elastic Compute Cloud (EC2). The Amazon EC2 API includes operations for working with snapshots.

You need to update the code and provide the service endpoint, your gateway Amazon Resource Name (ARN), and the number of days before which snapshots are to be deleted, and you need to specify the boolean `viewOnly` indicating whether to view what would be deleted or actually perform the snapshot deletions. You should run the code first with just the view option (`viewOnly` set to `true`) to see what the code would delete. For a list of AWS service endpoints you can use with AWS Storage Gateway, see [Regions and Endpoints](#) in the *Amazon Web Services Glossary*.

```
import java.io.IOException;
import java.util.ArrayList;
import java.util.Calendar;
import java.util.Collection;
import java.util.Date;
import java.util.GregorianCalendar;
import java.util.List;

import com.amazonaws.auth.PropertiesCredentials;
import com.amazonaws.services.ec2.AmazonEC2Client;
import com.amazonaws.services.ec2.model.DeleteSnapshotRequest;
import com.amazonaws.services.ec2.model.DescribeSnapshotsRequest;
import com.amazonaws.services.ec2.model.DescribeSnapshotsResult;
import com.amazonaws.services.ec2.model.Filter;
import com.amazonaws.services.ec2.model.Snapshot;
import com.amazonaws.services.storagegateway.AWSStorageGatewayClient;
import com.amazonaws.services.storagegateway.model.ListVolumesRequest;
import com.amazonaws.services.storagegateway.model.ListVolumesResult;
import com.amazonaws.services.storagegateway.model.VolumeInfo;

public class ListDeleteVolumeSnapshotsExample {

    public static AWSStorageGatewayClient sgClient;
    public static AmazonEC2Client ec2Client;
    static String serviceURLSG = "https://storagegateway.us-east-1.amazonaws.com";
    static String serviceURLEC2 = "https://ec2.us-east-1.amazonaws.com";

    // The gatewayARN
    public static String gatewayARN = "*** provide gateway ARN ***";

    // The number of days back you want to save snapshots. Snapshots before
    // this cutoff are deleted
    // if viewOnly = false.
    public static int daysBack = 10;

    // true = show what will be deleted; false = actually delete snapshots that
    // meet the daysBack criteria
    public static boolean viewOnly = true;

    public static void main(String[] args) throws IOException {

        // Create a storage gateway and amazon ec2 client
        sgClient = new AWSStorageGatewayClient(new PropertiesCredentials(
```

```
        ListDeleteVolumeSnapshotsExample.class.getResourceAsStream("AwsCredentials.properties"));
        sgClient.setEndpoint(serviceURLSG);

        ec2Client = new AmazonEC2Client(new PropertiesCredentials(
            ListDeleteVolumeSnapshotsExample.class.getResourceAsStream("AwsCredentials.properties")));
        ec2Client.setEndpoint(serviceURLEC2);

        List<VolumeInfo> volumes = ListVolumesForGateway();
        DeleteSnapshotsForVolumes(volumes, daysBack);
    }
    public static List<VolumeInfo> ListVolumesForGateway()
    {
        List<VolumeInfo> volumes = new ArrayList<VolumeInfo>();

        String marker = null;
        do {
            ListVolumesRequest request = new ListVolumesRequest().withGatewayARN(gatewayARN);
            ListVolumesResult result = sgClient.listVolumes(request);
            marker = result.getMarker();

            for (VolumeInfo vi : result.getVolumeInfos())
            {
                volumes.add(vi);
                System.out.println(OutputVolumeInfo(vi));
            }
        } while (marker != null);

        return volumes;
    }
    private static void DeleteSnapshotsForVolumes(List<VolumeInfo> volumes,
        int daysBack2) {

        // Find snapshots and delete for each volume
        for (VolumeInfo vi : volumes) {

            String volumeARN = vi.getVolumeARN();
            String volumeId = volumeARN.substring(volumeARN.lastIndexOf("/") + 1).toLowerCase();
            Collection<Filter> filters = new ArrayList<Filter>();
            Filter filter = new Filter().withName("volume-id").withValues(volumeId);
            filters.add(filter);

            DescribeSnapshotsRequest describeSnapshotsRequest =
                new DescribeSnapshotsRequest().withFilters(filters);
            DescribeSnapshotsResult describeSnapshotsResult =
                ec2Client.describeSnapshots(describeSnapshotsRequest);

            List<Snapshot> snapshots = describeSnapshotsResult.getSnapshots();

            System.out.println("volume-id = " + volumeId);
            for (Snapshot s : snapshots){
                StringBuilder sb = new StringBuilder();
                boolean meetsCriteria = !CompareDates(daysBack, s.getStart
```

```
Time());
        sb.append(s.getSnapshotId() + ", " + s.getStartTime().to
String());
        sb.append(", meets criteria for delete? " + meetsCriteria);
        sb.append(", deleted? ");
        if (!viewOnly & meetsCriteria) {
            sb.append("yes");
            DeleteSnapshotRequest deleteSnapshotRequest =
                new DeleteSnapshotRequest().withSnapshotId(s.getSnap
shotId());
                ec2Client.deleteSnapshot(deleteSnapshotRequest);
            }
        else {
            sb.append("no");
        }
        System.out.println(sb.toString());
    }
}

private static String OutputVolumeInfo(VolumeInfo vi) {

    String volumeInfo = String.format(
        "Volume Info:\n" +
        "  ARN: %s\n" +
        "  Type: %s\n",
        vi.getVolumeARN(),
        vi.getVolumeType());
    return volumeInfo;
}

// Returns the date in two formats as a list
public static boolean CompareDates(int daysBack, Date snapshotDate) {
    Date today = new Date();
    Calendar cal = new GregorianCalendar();
    cal.setTime(today);
    cal.add(Calendar.DAY_OF_MONTH, -daysBack);
    Date cutoffDate = cal.getTime();
    return (snapshotDate.compareTo(cutoffDate) > 0) ? true : false;
}
}
```

Deleting Snapshots Using the AWS SDK for .NET

To delete many snapshots associated with a volume, you can use a programmatic approach. The example below demonstrates how to delete snapshots using the AWS SDK for .NET. To use the example code, you should be familiar with running a .NET console application. For more information, see [Getting Started](#) in the *AWS SDK for .NET Developer Guide*. If you need to just delete a few snapshots, use the console as described in [To delete a snapshot using the console \(p. 210\)](#).

Example : Deleting Snapshots Using the AWS SDK for .NET

The following C# code example lists the snapshots for each volume of a gateway and whether the snapshot start time is before or after a specified date. It uses the AWS SDK for .NET API for AWS Storage Gateway and Amazon Elastic Compute Cloud (EC2). The Amazon EC2 API includes operations for working with snapshots.

You need to update the code and provide the service endpoint, your gateway Amazon Resource Name (ARN), and the number of days before which snapshots are to be deleted, and you need to specify the boolean `viewOnly` indicating whether to view what would be deleted or actually perform the snapshot deletions. You should run the code first with just the view option (`viewOnly` set to `true`) to see what the code would delete. For a list of AWS service endpoints you can use with AWS Storage Gateway, see [Regions and Endpoints](#) in the *Amazon Web Services Glossary*.

```
using System;
using System.Collections.Generic;
using System.Text;
using Amazon.StorageGateway.Model;
using Amazon.StorageGateway;
using Amazon.EC2;
using Amazon.EC2.Model;

namespace AWSStorageGateway
{
    class ListDeleteVolumeSnapshotsExample
    {
        static AmazonStorageGatewayClient sgClient;
        static AmazonStorageGatewayConfig sgConfig;
        static AmazonEC2Config ec2Config;
        static AmazonEC2Client ec2Client;
        static String serviceURLSG = "https://storagegateway.us-east-1.amazonaws.com";
        static String serviceURLEC2 = "https://ec2.us-east-1.amazonaws.com";

        // The gatewayARN
        public static String gatewayARN = "*** provide gateway ARN ***";

        // The number of days back you want to save snapshots. Snapshots before
        // this cutoff are deleted
        // if viewOnly = false.
        public static int daysBack = 10;

        // true = show what will be deleted; false = actually delete snapshots
        // that meet the daysBack criteria
        public static Boolean viewOnly = true;

        public static void Main(string[] args)
        {
            // Create a storage gateway client
            sgConfig = new AmazonStorageGatewayConfig();
            sgConfig.ServiceURL = serviceURLSG;
            sgClient = new AmazonStorageGatewayClient(sgConfig);

            //Create a ec2 client
            ec2Config = new AmazonEC2Config();
            ec2Config.ServiceURL = serviceURLEC2;
            ec2Client = new AmazonEC2Client(ec2Config);
        }
    }
}
```

```
List<VolumeInfo> volumes = ListVolumesForGateway();
DeleteSnapshotsForVolumes(volumes, daysBack);

Console.WriteLine("\nTo continue, press Enter.");
Console.Read();
}
public static List<VolumeInfo> ListVolumesForGateway()
{
    ListVolumesRequest request = new ListVolumesRequest().WithGatewayARN(gatewayARN);
    ListVolumesResponse response = sgClient.ListVolumes(request);
    ListVolumesResult volumes = response.ListVolumesResult;

    foreach (VolumeInfo vi in volumes.VolumeInfos)
    {
        Console.WriteLine(OutputVolumeInfo(vi));
    }

    return volumes.VolumeInfos;
}

public static void DeleteSnapshotsForVolumes(List<VolumeInfo> volumes,
int cutoffDays)
{
    foreach (VolumeInfo vi in volumes)
    {
        String volumeARN = vi.VolumeARN;
        String volumeId = volumeARN.Substring(volumeARN.LastIndexOf("/")
+ 1).ToLower();
        Filter[] filters = new Filter[1];
        filters[0] = new Filter().WithName("volume-id").WithValue(volumeId);

        DescribeSnapshotsRequest describeSnapshotsRequest =
            new DescribeSnapshotsRequest().WithFilter(filters);
        DescribeSnapshotsResponse describeSnapshotsResponse =
            ec2Client.DescribeSnapshots(describeSnapshotsRequest);

        List<Snapshot> snapshots = describeSnapshotsResponse.DescribeSnapshotsResult.Snapshot;
        Console.WriteLine("volume-id = " + volumeId);
        foreach (Snapshot s in snapshots)
        {
            StringBuilder sb = new StringBuilder();
            Boolean meetsCriteria = CompareDates(daysBack, s.StartTime);

            sb.Append(s.SnapshotId + ", " + s.StartTime);
            sb.Append(", meets criteria for delete? " + meetsCriteria);

            sb.Append(", deleted? ");
            if (!viewOnly & meetsCriteria)
            {
                sb.Append("yes");
                DeleteSnapshotRequest deleteSnapshotRequest =
                    new DeleteSnapshotRequest().WithSnapshotId(s.SnapshotId);
                ec2Client.DeleteSnapshot(deleteSnapshotRequest);
            }
        }
    }
}
```

```
        }
        else
        {
            sb.Append("no");
        }
        Console.WriteLine(sb.ToString());
    }
}

private static String OutputVolumeInfo(VolumeInfo vi)
{
    String volumeInfo = String.Format(
        "Volume Info:\n" +
        "  ARN: {0}\n" +
        "  Type: {1}\n",
        vi.VolumeARN,
        vi.VolumeType);
    return volumeInfo;
}

public static Boolean CompareDates(int daysBack, String d)
{
    DateTime snapshotDate = DateTime.Parse(d);
    DateTime cutoffDate = DateTime.Now.Add(new TimeSpan(-daysBack, 0,
0, 0));
    return (DateTime.Compare(snapshotDate, cutoffDate) < 0) ? true :
false;
}
}
```

Deleting Snapshots Using the AWS Tools for Windows PowerShell

To delete many snapshots associated with a volume, you can use a programmatic approach. The example below demonstrates how to delete snapshots using the AWS Tools for Windows PowerShell. To use the example script, you should be familiar with running a PowerShell script. For more information, see [Getting Started](#) in the *AWS Tools for Windows PowerShell*. If you need to delete just a few snapshots, use the console as described in [To delete a snapshot using the console](#) (p. 210).

Example : Deleting Snapshots Using the AWS Tools for Windows PowerShell

The following PowerShell script example lists the snapshots for each volume of a gateway and whether the snapshot start time is before or after a specified date. It uses the AWS Tools for Windows PowerShell cmdlets for AWS Storage Gateway and Amazon Elastic Compute Cloud (Amazon EC2). The Amazon EC2 API includes operations for working with snapshots.

You need to update the script and provide your gateway Amazon Resource Name (ARN), and the number of days before which snapshots are to be deleted, and you need to specify the boolean `viewOnly` indicating whether to view what would be deleted or actually perform the snapshot deletions. You should run the code first with just the view option (`viewOnly` set to `true`) to see what the code would delete.

```
<#
.DESCRIPTION
    Delete Snapshots of a specified volume that match given criteria.

.NOTES
    PREREQUISITES:
    1) AWS Tools for PowerShell from http://aws.amazon.com/powershell/
    2) Credentials and region stored in session using Initialize-AWSDefault.
    For more info see, http://docs.aws.amazon.com/powershell/latest/userguide/specifying-your-aws-credentials.html

.EXAMPLE
    powershell.exe .\SG_DeleteSnapshots.ps1
#>

# Criteria to use to filter the results returned.
$daysBack = 18
$gatewayARN = "*** provide gateway ARN ***"
$viewOnly = $true;

#ListVolumes
$volumesResult = Get-SGVolume -GatewayARN $gatewayARN
$volumes = $volumesResult.VolumeInfos
Write-Output("`nVolume List")
foreach ($vi in $volumes)
{
    Write-Output("`nVolume Info:")
    Write-Output("ARN: " + $vi.VolumeARN)
    Write-Output("Type: " + $vi.VolumeType)
}

Write-Output("`nWhich snapshots meet the criteria?")
foreach ($vi in $volumes)
{
    $volumeARN = $vi.VolumeARN
    $volumeId = $volumeARN.Substring($volumeARN.LastIndexOf("/")+1).ToLower()

    $filter = New-Object Amazon.EC2.Model.Filter
    $filter.Name = "volume-id"
    $filter.Value.Add($volumeId)

    $snapshots = get-EC2Snapshot -Filter $filter
    Write-Output("`nFor volume-id = " + $volumeId)
    foreach ($s in $snapshots)
    {
        $d = ([DateTime]::Now).AddDays(-$daysBack)
        $meetsCriteria = $false
    }
}
}
```



```
    if ([DateTime]::Compare($d, $s.StartTime) -gt 0)
    {
        $meetsCriteria = $true
    }

    $sb = $s.SnapshotId + ", " + $s.StartTime + ", meets criteria for delete?"
" + $meetsCriteria
    if (!$ViewOnly -AND $meetsCriteria)
    {
        $resp = Remove-EC2Snapshot -SnapshotId $s.SnapshotId
        #Can get RequestId from response for troubleshooting.
        $sb = $sb + ", deleted? yes"
    }
    else {
        $sb = $sb + ", deleted? no"
    }
    Write-Output($sb)
}
}
```

Restoring a Snapshot

Topics

- [Restoring a Snapshot to an AWS Storage Gateway Volume \(p. 219\)](#)
- [Restoring a Snapshot to an Amazon EBS Volume \(p. 222\)](#)

You can restore a snapshot of a volume to a new AWS Storage Gateway volume, or you can use the snapshot to create an Amazon Elastic Block Store (EBS) volume and attach this volume to an Amazon EC2 instance. When you restore the snapshot to a new AWS Storage Gateway volume, you can mount the volume as an iSCSI device to your on-premises application server and access the contents of this snapshot similar to when you create a new volume.

The use cases for restoring snapshots depends on the type of gateway you set up (see [How AWS Storage Gateway Works \(p. 3\)](#)).

- For gateway-cached volumes, your volume data is already stored in Amazon S3, so snapshots are typically used to preserve older versions of your data. After initiating a snapshot restore to a gateway-cached volume, snapshot data is downloaded to the local cache only upon first access of the data.
- For gateway-stored volumes, your volume data is stored on-premises, so snapshots provide durable, off-site backups in Amazon S3. For example, if a local disk allocated as a storage volume crashes, you can provision a new local disk and restore a snapshot to it during the volume creation process (see [Creating a Storage Volume \(Gateway-Stored\) \(p. 159\)](#)).

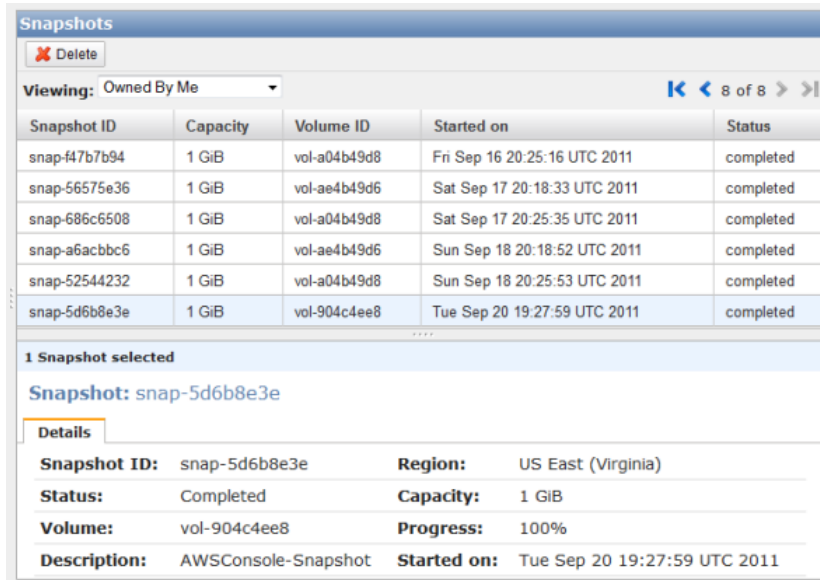
After you initiate a snapshot restore to a gateway-stored volume, snapshot data is downloaded in the background. This means that once you create a volume from a snapshot, there is no need to wait for all of the data to transfer from Amazon S3 to your volume before your application can start accessing the volume and all of its data. If your application accesses a piece of data that has yet to be loaded, the gateway immediately downloads the requested data from Amazon S3, and then continues loading the rest of the volume's data in the background.

Restoring a Snapshot to an AWS Storage Gateway Volume

The following procedure applies to both gateway-cached and gateway-stored volumes.

To create a storage volume from an existing snapshot

1. On the AWS Storage Gateway console, click **Snapshots**.



Snapshots

Delete

Viewing: Owned By Me 8 of 8

Snapshot ID	Capacity	Volume ID	Started on	Status
snap-f47b7b94	1 GiB	vol-a04b49d8	Fri Sep 16 20:25:16 UTC 2011	completed
snap-56575e36	1 GiB	vol-ae4b49d6	Sat Sep 17 20:18:33 UTC 2011	completed
snap-686c6508	1 GiB	vol-a04b49d8	Sat Sep 17 20:25:35 UTC 2011	completed
snap-a6acbbc6	1 GiB	vol-ae4b49d6	Sun Sep 18 20:18:52 UTC 2011	completed
snap-52544232	1 GiB	vol-a04b49d8	Sun Sep 18 20:25:53 UTC 2011	completed
snap-5d6b8e3e	1 GiB	vol-904c4ee8	Tue Sep 20 19:27:59 UTC 2011	completed

1 Snapshot selected

Snapshot: snap-5d6b8e3e

Details

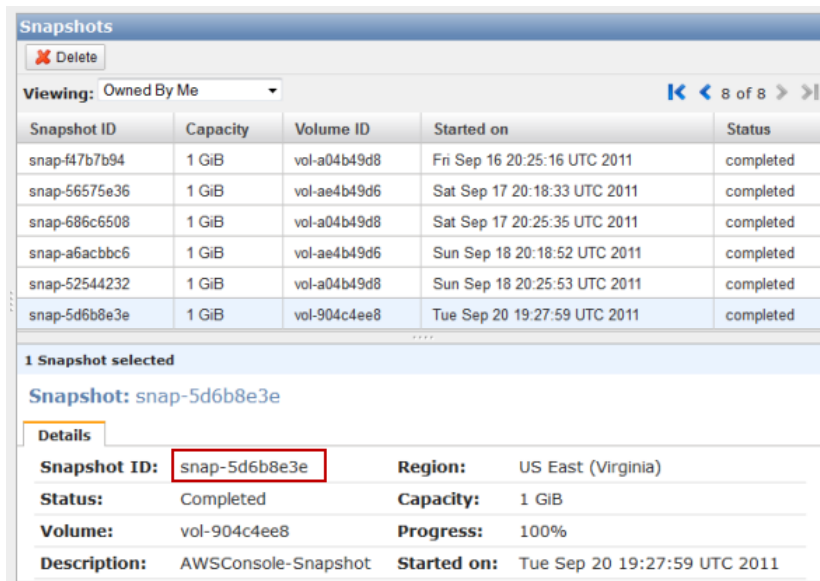
Snapshot ID: snap-5d6b8e3e **Region:** US East (Virginia)

Status: Completed **Capacity:** 1 GiB

Volume: vol-904c4ee8 **Progress:** 100%

Description: AWSConsole-Snapshot **Started on:** Tue Sep 20 19:27:59 UTC 2011

2. In the snapshot list, select the snapshot you want to create a storage volume from and note the Snapshot ID for use in a subsequent step.



Snapshots

Delete

Viewing: Owned By Me 8 of 8

Snapshot ID	Capacity	Volume ID	Started on	Status
snap-f47b7b94	1 GiB	vol-a04b49d8	Fri Sep 16 20:25:16 UTC 2011	completed
snap-56575e36	1 GiB	vol-ae4b49d6	Sat Sep 17 20:18:33 UTC 2011	completed
snap-686c6508	1 GiB	vol-a04b49d8	Sat Sep 17 20:25:35 UTC 2011	completed
snap-a6acbbc6	1 GiB	vol-ae4b49d6	Sun Sep 18 20:18:52 UTC 2011	completed
snap-52544232	1 GiB	vol-a04b49d8	Sun Sep 18 20:25:53 UTC 2011	completed
snap-5d6b8e3e	1 GiB	vol-904c4ee8	Tue Sep 20 19:27:59 UTC 2011	completed

1 Snapshot selected

Snapshot: snap-5d6b8e3e

Details


Snapshot ID: snap-5d6b8e3e **Region:** US East (Virginia)

Status: Completed **Capacity:** 1 GiB

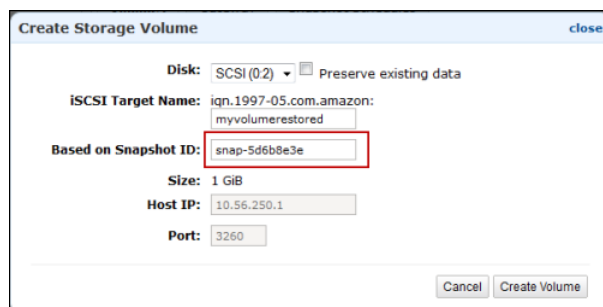
Volume: vol-904c4ee8 **Progress:** 100%

Description: AWSConsole-Snapshot **Started on:** Tue Sep 20 19:27:59 UTC 2011

3. In the **Navigation** pane, select the gateway to which you want to restore the snapshot.
4. Click **Create Volume**.



5. Depending on the type of gateway you configured, choose one of the following steps.
 - a. To create a new gateway-stored volume.
 - i. In the **Create Storage Volume** dialog box, paste the Snapshot ID you copied previously into the **Based on Snapshot ID** field.



- ii. Select a disk and a unique target name and click **Create Volume**.

The size of your storage volume must be greater than or equal to the size of the snapshot. To add a disk to your gateway VM that can be used as a stored-volume, see [Adding Local Disks for Volume Storage \(Gateway-Stored\)](#) (p. 102). You can now access the contents of this volume from your on-premises applications (see [Configuring Your Application Access to Storage Volumes](#) (p. 161)).

- b. To create a new gateway-cached volume.
 - i. In the **Configure Your Gateway** dialog box, paste the Snapshot ID you copied previously into the **Based on Snapshot ID** field.

The screenshot shows a dialog box titled "Configure Your Activated Gateway" with a "close" button in the top right corner. Below the title is a descriptive paragraph: "Create an iSCSI storage volume up to 32 TBs in size. This volume will be stored in Amazon S3, with only a cache of recently accessed data kept locally. Your client applications will connect to this volume over an iSCSI interface. [Learn More.](#)". Below this are several input fields: "Capacity" is set to "1" with a dropdown for "TBs" and "(Max: 32 TBs)"; "iSCSI Target Name" is "iqn.1997-05.com.amazon:myvolume"; "Based on Snapshot ID" is "snap-5d6b8e3e" and is highlighted with a red box; "Host IP" is "192.168.99.227"; and "Port" is "3260". At the bottom right are "Cancel" and "Create Volume" buttons.

- ii. Select a capacity for the disk and a unique target name and click **Create Volume**.

The size of your storage volume must be greater than or equal to the size of the snapshot. You can now access the contents of this volume from your on-premises applications (see [Configuring Your Application Access to Storage Volumes](#) (p. 161)).

Restoring a Snapshot to an Amazon EBS Volume

Your snapshots of your local storage volumes taken by AWS Storage Gateway are stored in Amazon S3 as Amazon EBS snapshots. For snapshots up to 1 TiB in size, you can restore snapshots of your local storage volumes to an Amazon EBS volume, and you can then attach the Amazon EBS volume to an Amazon EC2 instance. This allows you to easily migrate data from your on-premises applications to your applications running on Amazon EC2 in the event that you need to utilize Amazon EC2's compute capacity for disaster recovery or data processing. To see detailed pricing for Amazon EC2 and Amazon EBS, go to the [Amazon EC2 Pricing](#) page.

To restore a snapshot to an Amazon EBS volume

1. Create an Amazon EBS volume.
 - Follow the instructions in [Creating an Amazon EBS Volume](#) in the *Amazon Elastic Compute Cloud User Guide*.

The volume size that you specify must be greater than or equal to the size of the snapshot. Select the snapshot ID in the drop-down list of the **Create Volume** wizard in the **EBS Volumes** pane of the Amazon EC2 console. Alternatively, you can use the Amazon EC2 API to create your Amazon EBS volumes.
2. Attach the Amazon EBS volume to an Amazon EC2 instance. For more information, go to [Attaching the Volume to an Instance](#) in the *Amazon Elastic Compute Cloud User Guide*.

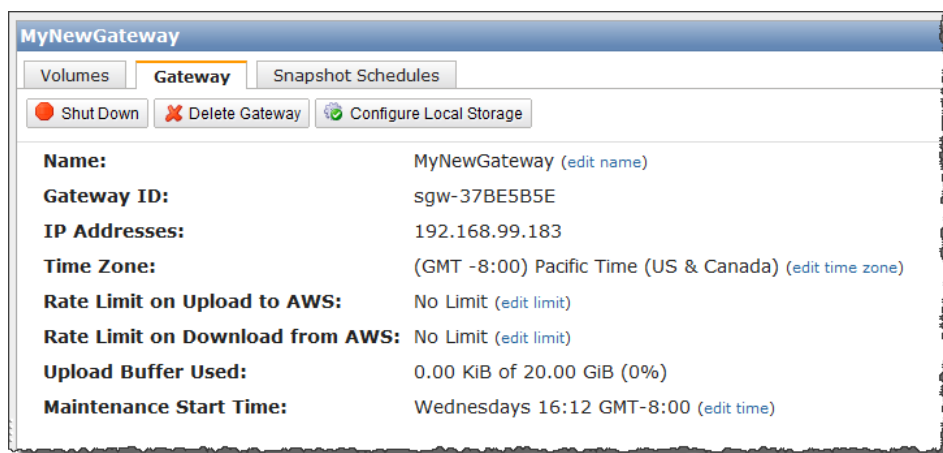
Performing Maintenance Tasks in AWS Storage Gateway

Topics

- [Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console](#) (p. 224)

- [Managing Gateway Updates Using the AWS Storage Gateway Console \(p. 226\)](#)
- [Updating Gateway Rate Limits \(p. 227\)](#)
- [Deleting a Gateway Using the AWS Storage Gateway Console \(p. 233\)](#)
- [Logging Into Your AWS Storage Gateway Local Console \(p. 234\)](#)
- [Routing AWS Storage Gateway Through a Proxy \(p. 238\)](#)
- [Configuring Your AWS Storage Gateway to Use a Static IP Address \(p. 239\)](#)
- [Testing Your AWS Storage Gateway Connection to the Internet \(p. 242\)](#)
- [Synchronizing Your Gateway VM Time \(p. 243\)](#)
- [Configuring AWS Storage Gateway for Multiple Network Adapters \(NICs\) \(p. 245\)](#)
- [Creating a Storage Volume in AWS Storage Gateway with Multiple Network Adapters \(p. 251\)](#)

You can perform many gateway maintenance-related tasks on the **Gateway** tab in the AWS Storage Gateway console. The following example shows the **Gateway** tab.



The following table summarizes the updatable fields on the **Gateway** tab. Click the **edit** link at the end of a field that can be edited to change the value.

Maintenance Item	Comments
Name	You can optionally change the name of your gateway. If you use Amazon CloudWatch to view your gateway metrics (see Using the Amazon CloudWatch Console (p. 261)), you might want to take note of the previous name and the new name to avoid confusion, or just use the gateway ID, which remains the same.
Gateway ID	AWS Storage Gateway assigns a unique identifier for each gateway. This value cannot be changed.
IP Addresses	Your storage applications can access a gateway's storage volumes using more than one IP address if the gateway is hosted on a server with more than one network interface card. In this scenario, all addresses that can be used to communicate with the gateway are listed in this field.
Time Zone	AWS Storage Gateway uses the time zone when displaying time-based information such as maintenance messages from AWS and snapshot scheduling.

Maintenance Item	Comments
Rate Limit on Upload to AWS	You can choose to limit the upload throughput from the gateway to AWS. Apply bandwidth throttling to your gateway to control the amount of network bandwidth used. Specify the rate limit as kilobits per second (kbps). The default is no rate limit on upload. For more information on updating this bandwidth, see Updating Gateway Rate Limits (p. 227) .
Rate Limit on Download from AWS	You can choose to limit the download throughput from AWS to your gateway. Apply bandwidth throttling to your gateway to control the amount of network bandwidth used. Specify the rate limit as kilobits per second (kbps). The default is no rate limit on download. For more information about updating this bandwidth, see Updating Gateway Rate Limits (p. 227) .
Upload Buffer Used	Displays the upload buffer used. For information about how to monitor the upload buffer and how it changes over time, see Monitoring the Upload Buffer (p. 267) .
Maintenance Start Time	Each gateway has a maintenance window of one time per week. During activation, a default time is assigned to your gateway. To change the time, click edit and specify a day of the week and time of the day in the time zone of the gateway.
Apply Update Now	If there is an update for your AWS Storage Gateway, a message appears in the console. Click Apply Update Now to apply the update immediately. If you do not apply the update, AWS Storage Gateway applies the update based on your Maintenance Start Time setting. For more information, see Managing Gateway Updates Using the AWS Storage Gateway Console (p. 226) .

Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console

This section discusses shutting down and turning on a gateway. You might need to shut down your gateway, for example, to apply a patch to your hypervisor host. While a gateway is shut down, your applications cannot access storage volumes and therefore cannot write any data to these storage volumes. The gateway also stops uploading any data to AWS.


Before shutting down your gateway, you must stop any applications that are writing to storage volumes by stopping your iSCSI Initiator connection. If a snapshot is in progress when the gateway is shut down, the snapshot will resume on gateway restart. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the console. For more information, see [Editing a Snapshot Schedule \(p. 208\)](#).

Note


Note that when you shut down a gateway using the AWS Storage Gateway console, you are stopping the gateway. However, the gateway VM remains on. If you need to power off the VM, use your VMware vSphere client or Microsoft Hyper-V Manager to connect to your host and then power off the specific VM. In most common scenarios in which you use the gateway after activation, you do not need to shut down the gateway VM.

To shut down a gateway


1. In the **Navigation** pane of the **AWS Storage Gateway** console, select the gateway.
2. Click **Shut Down**.



3. In the confirmation dialog box, click **Okay**.



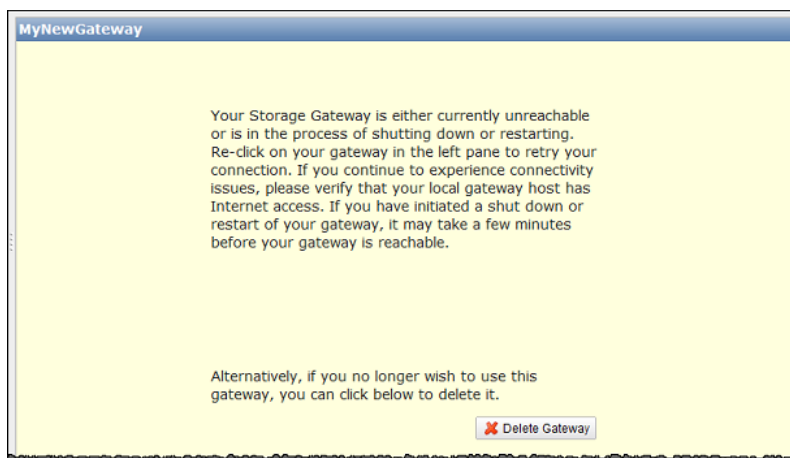
4. In the **Gateway Shut Down** dialog box, click **Close**.



5. While the gateway is shutting down, you may see a message that your gateway is in the process of shutting down.

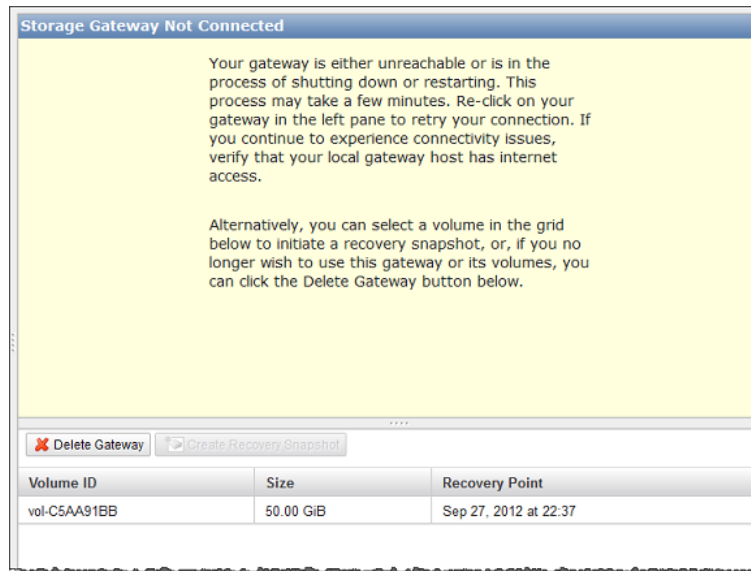
The options you have at this point depend on the type of gateway (cached-volume or stored-volume).

- a. For a gateway with stored-volumes, you have the option of deleting the gateway. Do not delete the gateway if you plan to restart the gateway and continue working with it.



- b. For a gateway with cached-volumes, you have the option of deleting the gateway or create a snapshot from a recovery point.

For more information on volume recovery points, see [Using Recovery Snapshots \(Gateway-Cached\)](#) (p. 258).



6. Select the gateway in the left navigation pane.

A **Restart** button is displayed.



To turn on a gateway

1. In the AWS Storage Gateway console, in the **Navigation** pane, select the gateway to restart.
2. Click **Restart**.

Managing Gateway Updates Using the AWS Storage Gateway Console

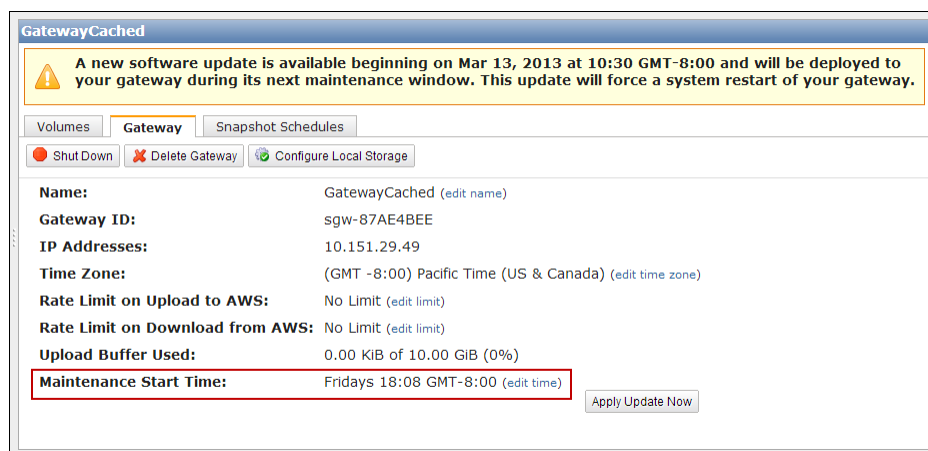
AWS Storage Gateway periodically deploys important updates and patches to your gateway that must be applied. Amazon will notify you via the AWS Storage Gateway console and via email in advance of any updates to your gateway. Software updates force a restart of your gateway which typically takes a few minutes to complete. You do not have to take any action, and in particular, you should not restart the VM manually. After the update, your gateway and its volumes will be in the same states as they were before the update. While the software update is being applied, application reads and writes from initiators to gateway storage volume targets are buffered and when the update is complete, the gateway processes

them. You can minimize the chance of any disruption to your applications by increasing your iSCSI Initiators' timeouts. For more information about increasing iSCSI Initiator timeouts for Windows and Linux, see [Customizing Your Windows iSCSI Settings \(p. 163\)](#) and [Customizing Your Linux iSCSI Settings \(p. 166\)](#), respectively.

You can choose to let AWS Storage Gateway apply updates according to the maintenance schedule for your gateway or you can apply the update yourself. When you deploy and activate your gateway, a default weekly maintenance schedule is set. You can modify this schedule at any time by clicking **edit** next to **Maintenance Start Time** in the **Gateway** tab. The following example shows the gateway maintenance tab with a maintenance message and the button in the UI for applying the update.

Important

A software update forces a system restart of your gateway. You can minimize the chance of any disruption to your applications by increasing your iSCSI Initiators' timeouts. For more information about increasing iSCSI Initiator timeouts for Windows and Linux, see [Customizing Your Windows iSCSI Settings \(p. 163\)](#) and [Customizing Your Linux iSCSI Settings \(p. 166\)](#), respectively.



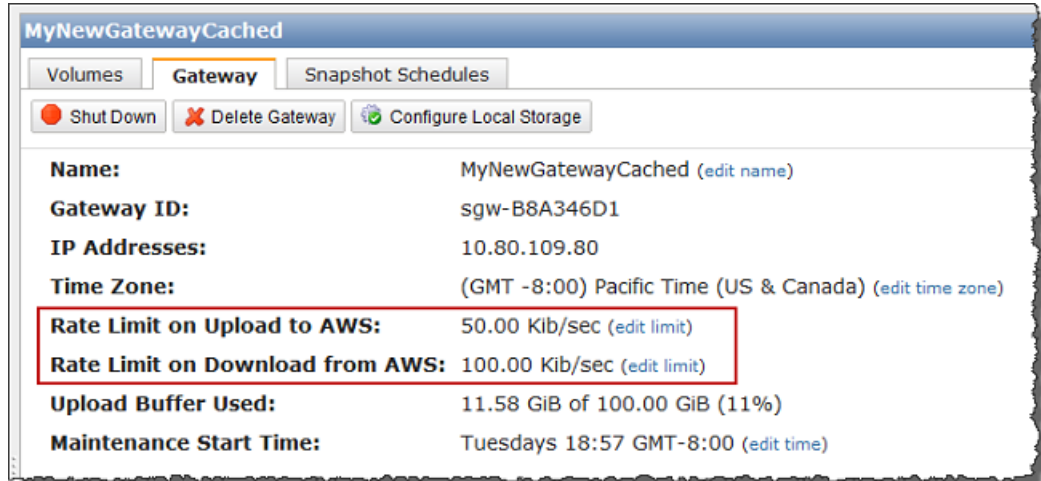
Updating Gateway Rate Limits

You can limit (or throttle) the upload throughput from the gateway to AWS or the download throughput from your AWS to your gateway. Using bandwidth throttling helps you to control the amount of network bandwidth used by your gateway. An activated gateway, by default, has no rate limits on upload or download.


You can specify rate limit using the AWS Management Console or programmatically using either the AWS API (see [UpdateBandwidthRateLimit \(p. 389\)](#)) or by using an AWS Software Development Kit (SDK). The ability to change throttling programmatically allows you to change limits automatically throughout the day, for example, by scheduling tasks to change the bandwidth. For examples of changing bandwidth rate limits programmatically, see [Updating Gateway Rate Limits Using the AWS SDK for Java \(p. 228\)](#), [Updating Gateway Rate Limits Using the AWS SDK for .NET \(p. 230\)](#) or [Updating Gateway Rate Limits Using the AWS Tools for Windows PowerShell \(p. 232\)](#).

To change a gateway's bandwidth throttling using the console

1. In the **Navigation** pane of the **AWS Storage Gateway** console, select the gateway you want to manage.
2. Click the **Gateway** tab in the right pane.
3. Click the **edit limit** text near the limit you want to change.



4. In the **Edit Rate Limits** dialog box, enter new limit values, and click **Save**.



Updating Gateway Rate Limits Using the AWS SDK for Java

Updating bandwidth rate limits programmatically provides a path for you to adjust limits automatically over a period of time, for example, by using scheduled tasks. The following example below demonstrates how to update a gateway's bandwidth rate limits using the AWS Software Development Kit (SDK) for Java. To use the example code, you should be familiar with running a Java console application. For more information, see [Getting Started](#) in the *AWS SDK for Java Developer Guide*.

Example : Updating Gateway Bandwidth Limits Using the AWS SDK for Java

The following Java code example updates a gateway's bandwidth rate limits. You need to update the code and provide the service endpoint, your gateway Amazon Resource Name (ARN), and the upload and download limits. For a list of AWS service endpoints you can use with AWS Storage Gateway, see [Regions and Endpoints](#) in the Amazon Web Services Glossary.

```
import java.io.IOException;

import com.amazonaws.AmazonClientException;
import com.amazonaws.auth.PropertiesCredentials;
import com.amazonaws.services.storagegateway.AWSStorageGatewayClient;
import com.amazonaws.services.storagegateway.model.UpdateBandwidthRateLimitRequest;
import com.amazonaws.services.storagegateway.model.UpdateBandwidthRateLimitResult;

public class UpdateBandwidthExample {

    public static AWSStorageGatewayClient sgClient;

    // The gatewayARN
    public static String gatewayARN = "*** provide gateway ARN ***";

    // The endpoint
    static String serviceURL = "https://storagegateway.us-east-1.amazonaws.com";

    // Rates
    static long uploadRate = 51200; // Bits per second, minimum 51200
    static long downloadRate = 102400; // Bits per second, minimum 102400

    public static void main(String[] args) throws IOException {

        // Create a storage gateway client
        sgClient = new AWSStorageGatewayClient(new PropertiesCredentials(
            ListDeleteVolumeSnapshotsExample.class.getResourceAsStream("AwsCredentials.properties")));
        sgClient.setEndpoint(serviceURL);

        UpdateBandwidth(gatewayARN, uploadRate, downloadRate);

    }

    private static void UpdateBandwidth(String gatewayARN2, long uploadRate2,
        long downloadRate2) {
        try
        {
            UpdateBandwidthRateLimitRequest updateBandwidthRateLimitRequest =
                new UpdateBandwidthRateLimitRequest()
                    .withGatewayARN(gatewayARN)
                    .withAverageDownloadRateLimitInBitsPerSec(downloadRate)
                    .withAverageUploadRateLimitInBitsPerSec(uploadRate);

            UpdateBandwidthRateLimitResult updateBandwidthRateLimitResult =
                sgClient.updateBandwidthRateLimit(updateBandwidthRateLimitRequest);
            String returnGatewayARN = updateBandwidthRateLimitResult.getGatewa
```

```
yARN();
    System.out.println("Updated the bandwidth rate limits of " + re
turnGatewayARN);
    System.out.println("Upload bandwidth limit = " + uploadRate + "
bits per second");
    System.out.println("Download bandwidth limit = " + downloadRate +
" bits per second");
    }
    catch (AmazonClientException ex)
    {
        System.err.println("Error updating gateway bandwidth.\n" + ex.to
String());
    }
}
```

Updating Gateway Rate Limits Using the AWS SDK for .NET

Updating bandwidth rate limits programmatically provides a path for you to adjust limits automatically over a period of time, for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth rate limits using the AWS Software Development Kit (SDK) for .NET. To use the example code, you should be familiar with running a .NET console application. For more information, see [Getting Started](#) in the *AWS SDK for .NET Developer Guide*.

Example : Updating Gateway Bandwidth Limits Using the AWS SDK for .NET

The following C# code example updates a gateway's bandwidth rate limits. You need to update the code and provide the service endpoint, your gateway Amazon Resource Name (ARN), and the upload and download limits. For a list of AWS service endpoints you can use with AWS Storage Gateway, see [Regions and Endpoints](#) in the Amazon Web Services General Reference.

```
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using Amazon.StorageGateway;
using Amazon.StorageGateway.Model;

namespace AWSStorageGateway
{
    class UpdateBandwidthExample
    {
        static AmazonStorageGatewayClient sgClient;
        static AmazonStorageGatewayConfig sgConfig;

        // The gatewayARN
        public static String gatewayARN = "*** provide gateway ARN ***";

        // The endpoint
        static String serviceURL = "https://storagegateway.us-east-1.amazon
aws.com";

        // Rates
        static long uploadRate = 51200; // Bits per second, minimum 51200
        static long downloadRate = 102400; // Bits per second, minimum 102400

        public static void Main(string[] args)
        {
            // Create a storage gateway client
            sgConfig = new AmazonStorageGatewayConfig();
            sgConfig.ServiceURL = serviceURL;
            sgClient = new AmazonStorageGatewayClient(sgConfig);

            UpdateBandwidth(gatewayARN, uploadRate, downloadRate);

            Console.WriteLine("\nTo continue, press Enter.");
            Console.Read();
        }

        public static void UpdateBandwidth(string gatewayARN, long uploadRate,
long downloadRate)
        {
            try
            {
                UpdateBandwidthRateLimitRequest updateBandwidthRateLimitRequest
=
                new UpdateBandwidthRateLimitRequest()
                .WithGatewayARN(gatewayARN)
                .WithAverageDownloadRateLimitInBitsPerSec(downloadRate)
                .WithAverageUploadRateLimitInBitsPerSec(uploadRate);
            }
        }
    }
}
```

```
UpdateBandwidthRateLimitResponse updateBandwidthRateLimitResponse
= sgClient.UpdateBandwidthRateLimit(updateBandwidthRateLimitRequest);
String returnGatewayARN = updateBandwidthRateLimitResponse.Up
dateBandwidthRateLimitResult.GatewayARN;
Console.WriteLine("Updated the bandwidth rate limits of " +
returnGatewayARN);
Console.WriteLine("Upload bandwidth limit = " + uploadRate + "
bits per second");
Console.WriteLine("Download bandwidth limit = " + downloadRate
+ " bits per second");
}
catch (AmazonStorageGatewayException ex)
{
    Console.WriteLine("Error updating gateway bandwidth.\n" +
ex.ToString());
}
}
```

Updating Gateway Rate Limits Using the AWS Tools for Windows PowerShell

Updating bandwidth rate limits programmatically provides a path for you to adjust limits automatically over a period of time, for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth rate limits using the AWS Tools for Windows PowerShell. To use the example code, you should be familiar with running a PowerShell script. For more information, see [Getting Started](#) in the *AWS Tools for Windows PowerShell User Guide*.

Example : Updating Gateway Bandwidth Limits Using the AWS Tools for Windows PowerShell

The following PowerShell script example updates a gateway's bandwidth rate limits. You need to update the script and provide your gateway Amazon Resource Name (ARN), and the upload and download limits.

```
<#
.DESCRIPTION
    Update Gateway bandwidth limits.

.NOTES
    PREREQUISITES:
    1) AWS Tools for PowerShell from http://aws.amazon.com/powershell/
    2) Credentials and region stored in session using Initialize-AWSDefault.
    For more info see, http://docs.aws.amazon.com/powershell/latest/userguide/specifying-your-aws-credentials.html

.EXAMPLE
    powershell.exe .\SG_UpdateBandwidth.ps1
#>

$UploadBandwidthRate = 51200
$DownloadBandwidthRate = 102400
$gatewayARN = "*** provide gateway ARN ***"

#Update Bandwidth Rate Limits
Update-SGBandwidthRateLimit -GatewayARN $gatewayARN `
    -AverageUploadRateLimitInBitsPerSec $UploadBandwidthRate `
    -AverageDownloadRateLimitInBitsPerSec $DownloadBandwidthRate

$limits = Get-SGBandwidthRateLimit -GatewayARN $gatewayARN

Write-Output("`nGateway: " + $gatewayARN);
Write-Output("`nNew Upload Rate: " + $limits.AverageUploadRateLimitInBitsPerSec)
Write-Output("`nNew Download Rate: " + $limits.AverageDownloadRateLimitInBitsPerSec)
```

Deleting a Gateway Using the AWS Storage Gateway Console

Deleting a gateway removes the gateway as an activated gateway that you can use to store application data. The deleted gateway no longer shows in the AWS Storage Gateway console and any existing iSCSI connections you have open to the gateway will be closed. Reusing the VM for a new gateway is not supported.

Important

You no longer pay software charges after the gateway is deleted; however, your existing Amazon EBS snapshots persist and you will continue to be billed for these snapshots. You can choose to remove all remaining Amazon EBS snapshots by canceling your Amazon EC2 subscription. If you prefer not to cancel your Amazon EC2 subscription, you can delete your snapshots using the Amazon EC2 console. For more information, see the [AWS Storage Gateway Detail Page](#).

If you accidentally delete your gateway, you can activate a new one (see [Activating AWS Storage Gateway \(p. 111\)](#)) and configure it to match the setup of the deleted gateway. Create a new VM for the


new gateway. For a gateway-stored setup, create the new gateway with the same upload buffer as that of the deleted gateway. For a gateway-cached setup, create the new gateway with the same upload buffer and cache storage as that of the deleted gateway. When you create storage volumes on the new gateway for the gateway-cached setup, create a new volume in Amazon S3 and restore from a snapshot. For the gateway-stored setup, you can use the same underlying disks as the deleted gateway and preserve the data on the disks. For more information, see [Managing Storage Volumes \(Gateway-Stored\)](#) (p. 183).

To delete a gateway

1. In the **Navigation** pane of the **AWS Storage Gateway** console, select the gateway you want to delete.
2. In the **Gateway** tab, click **Delete Gateway**.

Important

Be sure that there are no applications currently writing to the gateway's volumes. If you delete the gateway while it is in use, data loss may occur.



3. Confirm the deletion by clicking **OK**.

At this point, the deleted gateway is no longer an activated gateway. However, the gateway VM still exists in your virtualization environment. To remove the VM, use the VMware vSphere client or Microsoft Hyper-V Manager to connect to the host and remove the VM.

Logging Into Your AWS Storage Gateway Local Console

Some gateway maintenance tasks require that you log in to your gateway's local console. The local console is accessible through your hypervisor client software. The user is *sguser* and the password is *sgpassword*. These login credentials give you access to configuration menus, where you can configure gateway network settings.


In this topic, we show you how to access the local console of a gateway hosted in VMware ESXi (see [To access your gateway local console \(VMware ESXi\)](#) (p. 234)) or in Microsoft Hyper-V (see [To access into your gateway local console \(Microsoft Hyper-V\)](#) (p. 235)). After you access the console, you log into it (see [To log in to the gateway local console](#) (p. 237)).

To access your gateway local console (VMware ESXi)

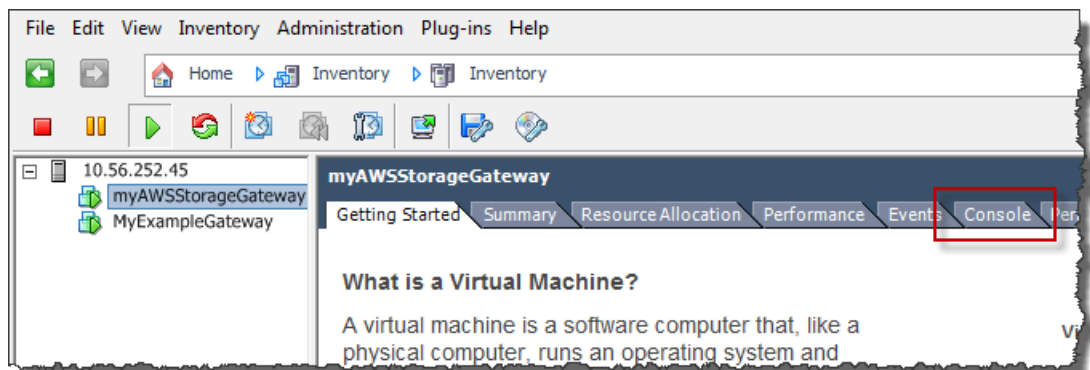
1. In the VMware vSphere client, select your gateway VM.
2. Ensure that the gateway is powered on.

Note

If your gateway VM is powered on, a green arrow icon appears with the VM icon as shown in the example below. If your gateway VM is not powered on, you can power it on by clicking the green **Power On** icon in the **Toolbar** menu.



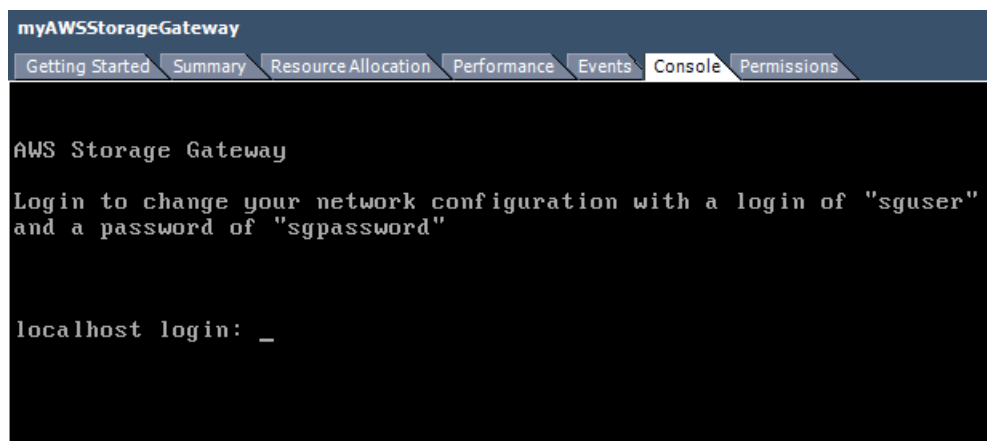
3. Click the **Console** tab.



4. After a few moments, the virtual machine is ready for you to log in.

Note

To release the cursor from the console window, press **Ctrl+Alt**.



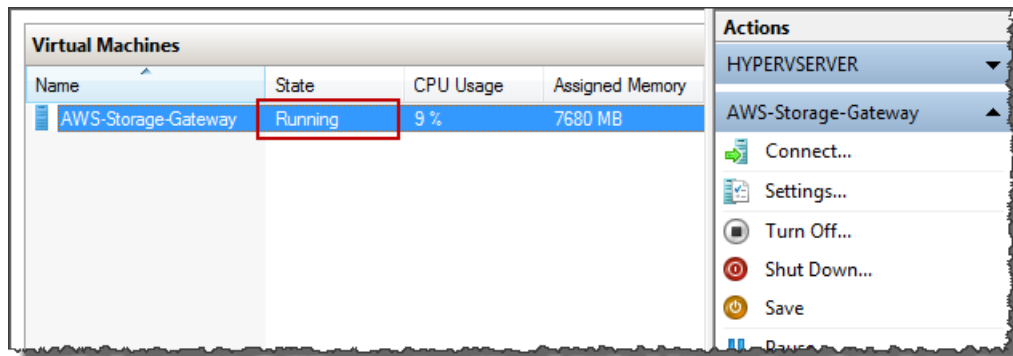
5. To log in, continue to the procedure [To log in to the gateway local console \(p. 237\)](#).

To access into your gateway local console (Microsoft Hyper-V)

1. In the Microsoft Hyper-V Manager, in the **Virtual Machines** list, select your gateway VM.
2. Ensure the gateway is powered on.


Note

If your gateway VM is powered on, `Running` is displayed as the **State** of the VM as shown in the example below. If your gateway VM is not powered on, you can power it on by clicking **Start** in the **Actions** pane.

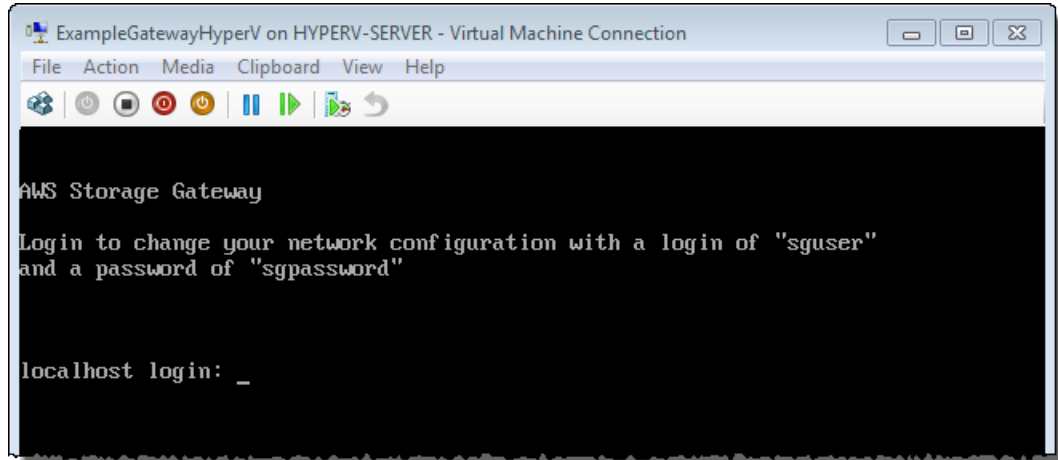


3. In the **Actions** pane, select **Connect...**

The **Virtual Machine Connection** window appears. If an authentication window appears, enter the user name and password provided to you by the hypervisor administrator.



4. After a few moments, the virtual machine is ready for you to log in.



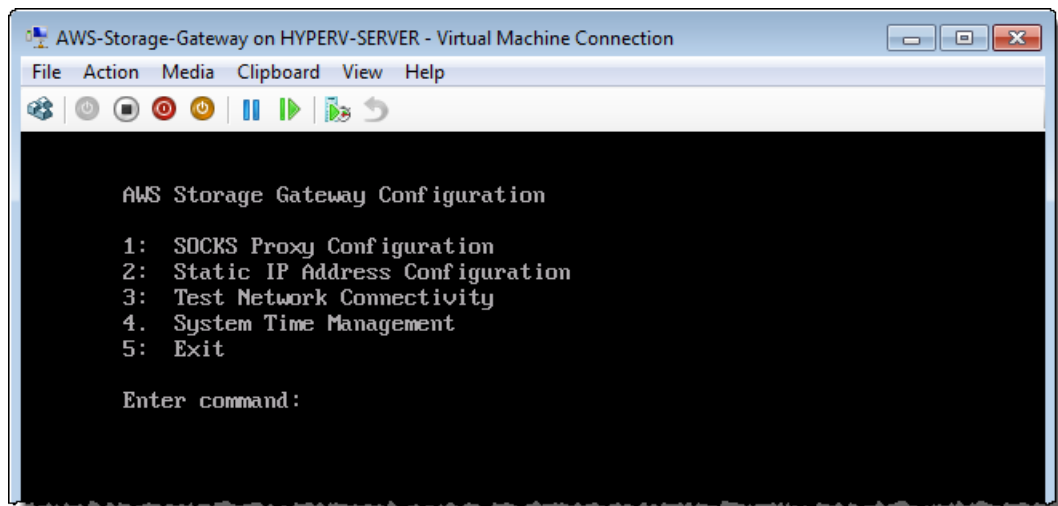
5. To log in, continue to the procedure [To log in to the gateway local console \(p. 237\)](#).

To log in to the gateway local console

1. In the login screen, log into the VM with the user name and password specified in the console window.
2. After you log in, you will see the **AWS Storage Gateway Configuration** main menu.

Note

For a gateway deployed on Microsoft Hyper-V, you will see an extra menu item for **System Time Management** that does not apply to gateways deployed on VMware ESXi.



To...	See...
Configure a SOCKS proxy for your gateway	Routing AWS Storage Gateway Through a Proxy (p. 238)
Configure static IP addresses for your gateway's interfaces	Configuring Your AWS Storage Gateway to Use a Static IP Address (p. 239)
Test network connectivity	Testing Your AWS Storage Gateway Connection to the Internet (p. 242)

To...	See...
Manage VM time (Microsoft Hyper-V only)	Synchronizing Your Gateway VM Time (p. 243)

Routing AWS Storage Gateway Through a Proxy

The AWS Storage Gateway supports the configuration of a SOCKS5 proxy between your gateway and AWS. If your gateway must use a proxy server to communicate to the Internet, then you need to configure a SOCKS proxy settings for your gateway. You do this by specifying an IP address and port number for the host running your proxy, and AWS Storage Gateway will route all HTTPS traffic through your proxy server.

To route your gateway Internet traffic through a local proxy server

1. Log into your gateway's local console. For instructions, see [Logging Into Your AWS Storage Gateway Local Console \(p. 234\)](#).
2. In the **AWS Storage Gateway Configuration** main menu, enter option 1.

Note

For a gateway deployed on Microsoft Hyper-V, you will see an extra menu item for **System Time Management** that does not apply to gateways deployed on VMware ESXi.

```
AWS Storage Gateway Configuration

1: SOCKS Proxy Configuration
2: Static IP Address Configuration
3: Test Network Connectivity
4: Exit

Enter command: _
```

3. Choose one of the following options in the **AWS Storage Gateway SOCKS Proxy Configuration** menu:

```
AWS Storage Gateway SOCKS Proxy Configuration

1: Configure SOCKS Proxy
2: View Current SOCKS Proxy Configuration
3: Remove SOCKS Proxy Configuration
4: Exit

Enter command: _
```

To...	Do this...
Configure a SOCKS proxy	Enter option 1. You will need to supply a host name and port to complete configuration.
View the current SOCKS proxy configuration	Enter option 2. If a SOCKS proxy is not configured, the message "SOCKS Proxy not configured" is displayed. If a SOCKS proxy is configured, the host name and port of the proxy are displayed.
Remove a SOCKS proxy configuration	Enter option 3. The message "SOCKS Proxy Configuration Removed" is displayed.
Exit this menu and return to the previous menu	Enter option 4.

Configuring Your AWS Storage Gateway to Use a Static IP Address

The default network configuration for the gateway is Dynamic Host Configuration Protocol (DHCP). With DHCP, your gateway is automatically assigned an IP address. In some cases, you may need to manually assign your gateway's IP as a static IP address. This topic explains how.

To configure your gateway to use static IP addresses

1. Log into your gateway's local console. For instructions, see [Logging Into Your AWS Storage Gateway Local Console](#) (p. 234).
2. In the **AWS Storage Gateway Configuration** main menu, select option 2.

Note

For a gateway deployed on Microsoft Hyper-V, you will see an extra menu item for **System Time Management** that does not apply to gateways deployed on VMware ESXi.

```
AWS Storage Gateway Configuration
1: SOCKS Proxy Configuration
2: Static IP Address Configuration
3: Test Network Connectivity
4: Exit

Enter command: _
```

3. Choose one of the following options in the **AWS Storage Gateway Static IP Address Configuration** menu:

```
AWS Storage Gateway Static IP Address Configuration

1: View Network Configuration
2: Configure Static IP
3: View DNS Configuration
4: Reset to DHCP
5: Set Default Route Adapter
6: View Routes
7: Exit

Enter command: _
```

To...	Do this...
View your gateway's network configuration	<p>Enter option 1.</p> <p>A list of adapter names is displayed and you are prompted to enter an adapter name, for example, "eth0". If the adapter you specify is in use, the following information about the adapter is displayed:</p> <ul style="list-style-type: none">• MAC address• IP address• Netmask• Gateway IP address• DHCP enabled status <p>You use the same adapter name when you configure a static IP address (option 2) or set your gateway's default route adapter (option 5).</p>

To...	Do this...
Configure a static IP address for your gateway	<p>Enter option 2.</p> <p>You are prompted to enter the following information to configure a static IP.</p> <ul style="list-style-type: none"> • Network adapter name • IP address • Netmask • Default gateway address • Primary DNS address • Secondary DNS address <p>Important If your gateway has already been activated, you must shut down and restart it from the AWS Storage Gateway console for the settings to take effect. For more information, see Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console (p. 224).</p> <p>If your gateway uses more than one network interface, you must set all enabled interfaces to use DHCP or static IP addresses. For example, if your gateway VM uses two interfaces configured as DHCP and you later set one interface to a static IP, the other interface is disabled. To enable the interface, you must set it to a static IP. If both interfaces are initially set to use static IP addresses and you then set the gateway to use DHCP, both interfaces will use DHCP.</p>
View your gateway's DNS configuration	<p>Enter option 3.</p> <p>The IP addresses of the primary and secondary DNS name servers are displayed.</p>
Reset your gateway's network configuration to DHCP	<p>Enter option 4.</p> <p>All network interfaces are set to use DHCP.</p> <p>Important If your gateway has already been activated, you must shut down and restart your gateway from the AWS Storage Gateway console for the settings to take effect. For more information, see Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console (p. 224).</p>

To...	Do this...
Set your gateway's default route adapter	Enter option 5 . The available adapters for your gateway are shown, and you are prompted to select one of the adapters, for example, "eth0".
View routing tables	Enter option 6 . The default route of your gateway is displayed.
Exit this menu and return to the previous menu	Enter option 7 .

Testing Your AWS Storage Gateway Connection to the Internet

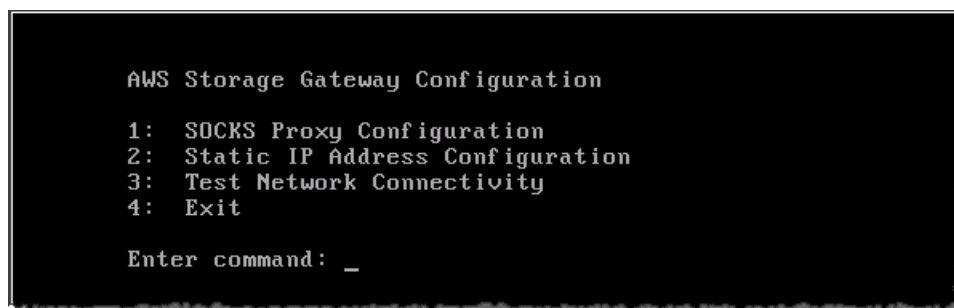
The AWS Storage Gateway configuration menus also let you test your gateway's connection to the Internet. This test can be useful when you are troubleshooting issues with your gateway. The network connectivity test does not take into account any SOCKS proxy you might have configured.

To test your gateway's connection to the Internet

1. Log into your gateway's local console. For instructions, see [Logging Into Your AWS Storage Gateway Local Console \(p. 234\)](#).
2. In the **AWS Storage Gateway Configuration** main menu, select option **3**.

Note

For a gateway deployed on Microsoft Hyper-V, you will see an extra menu item for **System Time Management** that does not apply to gateways deployed on VMware ESXi.



```
AWS Storage Gateway Configuration
1: SOCKS Proxy Configuration
2: Static IP Address Configuration
3: Test Network Connectivity
4: Exit
Enter command: _
```

The outcome from the testing network connectivity can be one of the following.

Connectivity is...	Message
Successful	AWS Storage Gateway has Internet connectivity
Not successful	AWS Storage Gateway does not have Internet connectivity

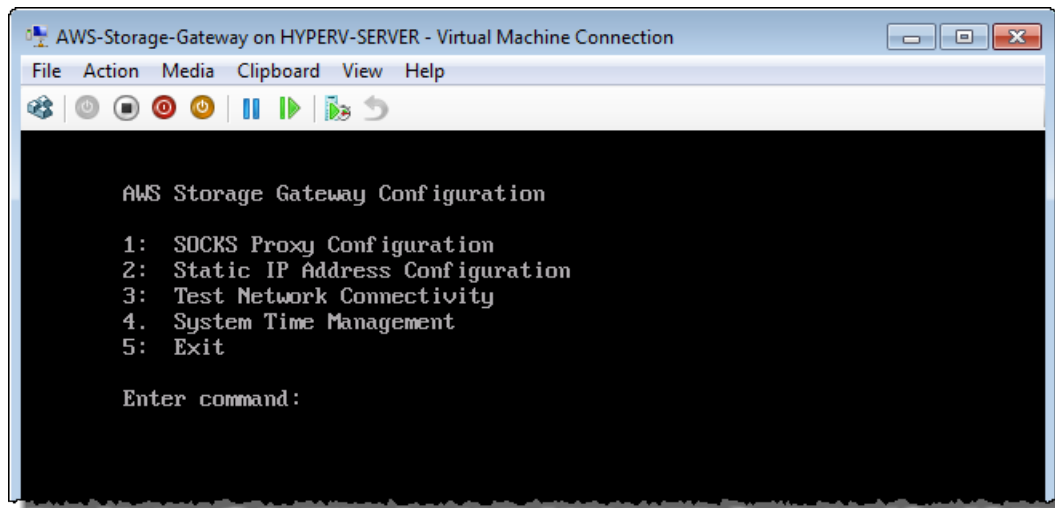
Synchronizing Your Gateway VM Time

After your gateway is deployed and running, there are some scenarios when the gateway VM's time may drift. For example, if there is a prolonged network outage and your hypervisor host and gateway do not get time updates, then the gateway VM's time will be off from the true time. When there is a time drift, there will be a discrepancy between actual times when operations like gateway updates or snapshots occur and the actual time that the operations occur.

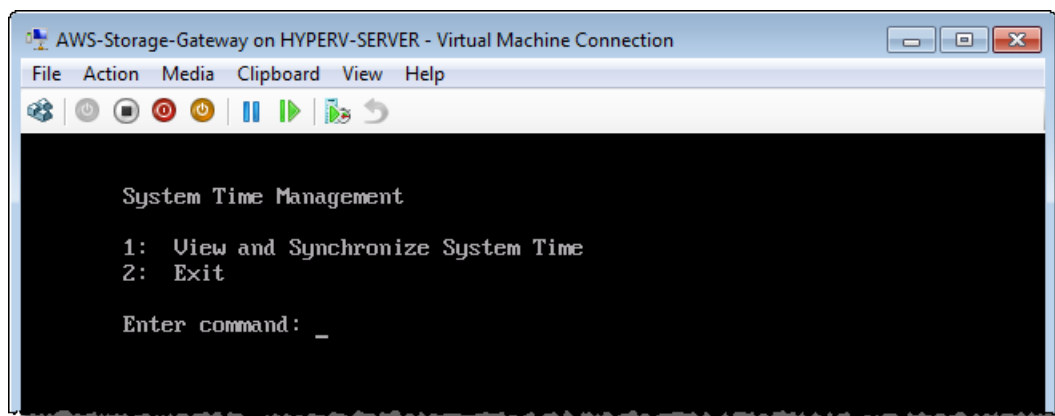
For a gateway deployed on VMware ESXi, setting the hypervisor host time and synchronizing the VM time to the host is sufficient to avoid clock drift. For more information, see (see [Synchronize VM Time with Host Time \(p. 15\)](#)). For a gateway deployed on Microsoft Hyper-V, you should periodically check your VM's time using the procedure described in this section.

To view and synchronize the time of a Hyper-V gateway VM to an NTP server

1. Log into your gateway's local console. For instructions, see [Logging Into Your AWS Storage Gateway Local Console \(p. 234\)](#).
2. In the **AWS Storage Gateway Configuration** main menu, select option **4** for **System Time Management**.



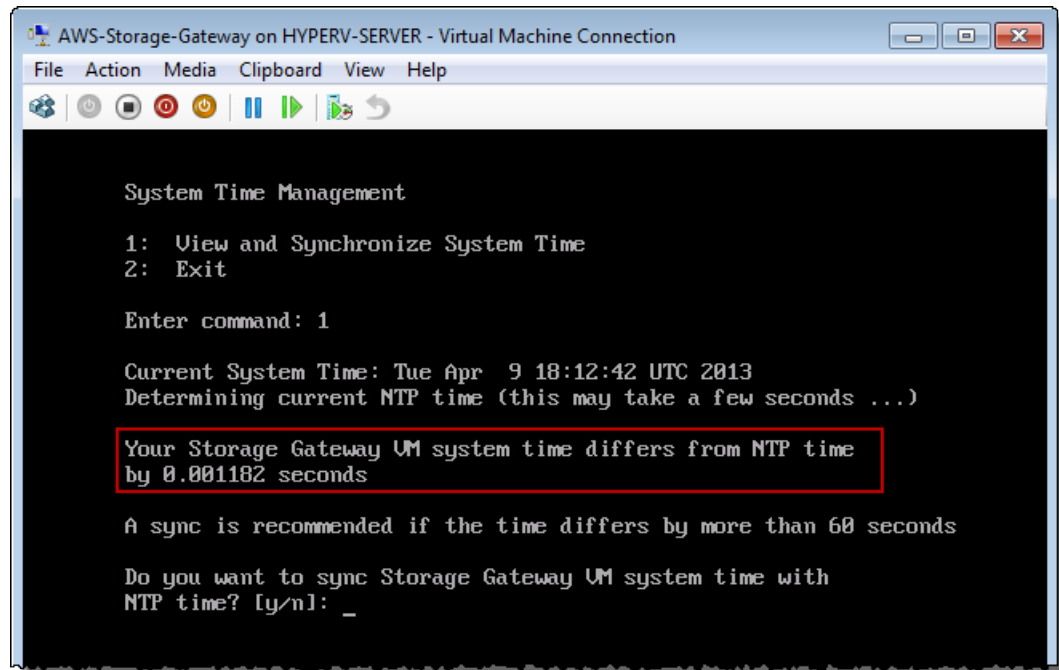
3. In the **System Time Management** menu, select option **1** for **View and Synchronize System Time**.



4. If the result indicates that you should sync your time, enter **y**; otherwise, enter **n**.

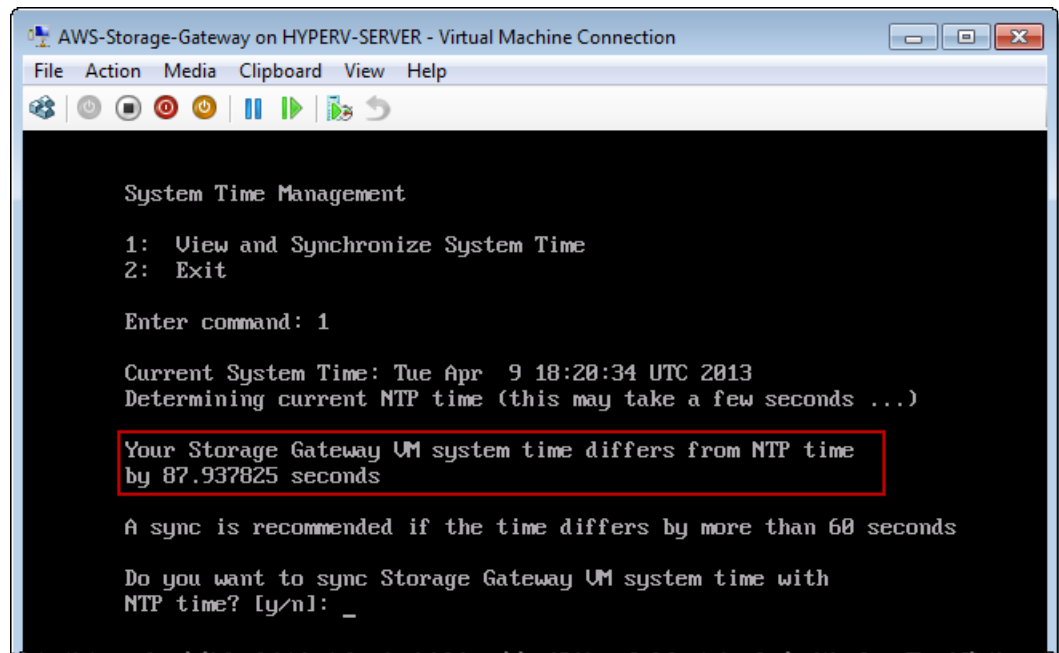
If you enter **y** to synchronize, the synchronization may take a few moments.

The following example shows a VM that does not require time synchronization.



```
AWS-Storage-Gateway on HYPERV-SERVER - Virtual Machine Connection
File Action Media Clipboard View Help
System Time Management
1: View and Synchronize System Time
2: Exit
Enter command: 1
Current System Time: Tue Apr 9 18:12:42 UTC 2013
Determining current NTP time (this may take a few seconds ...)
Your Storage Gateway VM system time differs from NTP time
by 0.001182 seconds
A sync is recommended if the time differs by more than 60 seconds
Do you want to sync Storage Gateway VM system time with
NTP time? [y/n]: _
```

The following example shows a VM that does require time synchronization.



```
AWS-Storage-Gateway on HYPERV-SERVER - Virtual Machine Connection
File Action Media Clipboard View Help
System Time Management
1: View and Synchronize System Time
2: Exit
Enter command: 1
Current System Time: Tue Apr 9 18:20:34 UTC 2013
Determining current NTP time (this may take a few seconds ...)
Your Storage Gateway VM system time differs from NTP time
by 87.937825 seconds
A sync is recommended if the time differs by more than 60 seconds
Do you want to sync Storage Gateway VM system time with
NTP time? [y/n]: _
```

Configuring AWS Storage Gateway for Multiple Network Adapters (NICs)

Gateways can be accessed by more than one IP address if you configure them to use multiple network adapters. Use cases when you would want to configure a gateway to use multiple network adapters include:

- **Maximizing Throughput** – You might want to maximize throughput to a gateway when network adapters are a bottleneck.
- **Application Separation** – You might need to separate your applications and how they write to a gateway's storage volumes. You might choose, for example, to have a critical storage application exclusively use one of the adapters defined for a gateway.
- **Network Constraints** – Your application environment may require that you keep your iSCSI targets and initiators that connect to them in an isolated network which is different from the network that the gateway uses to communicate to AWS.

In a typical multiple adapter use case, one adapter is configured as the route by which the gateway communicates with AWS (default gateway). Except for this one adapter, initiators must be in the same subnet as the adapter that contains the iSCSI targets to which they connect; otherwise, communication with intended targets may not be possible. If a target is configured on the same adapter that is used for communication with AWS, then iSCSI traffic for that target and AWS traffic will flow through the same adapter.

The following procedure assumes that your gateway VM already has one network adapter defined and you will add a second adapter. The first procedure shows how to add an adapter for VMware ESXi and the second procedure shows how for Microsoft Hyper-V.

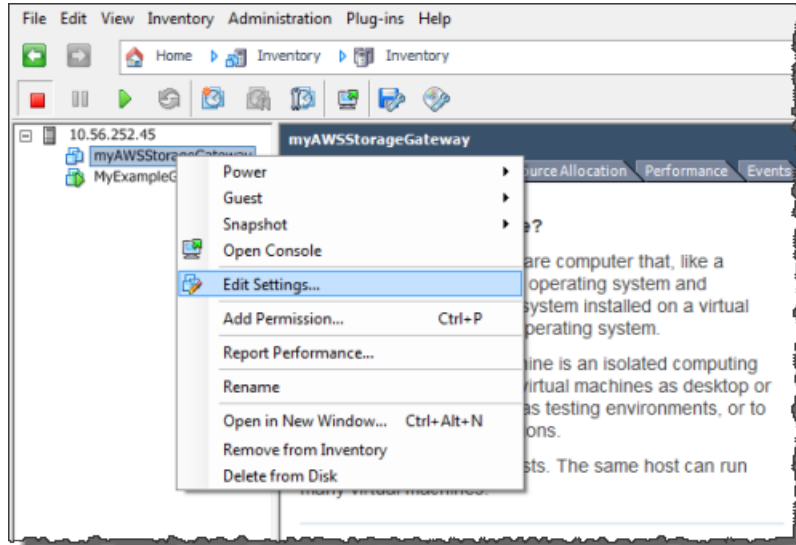
To configure your gateway to use an additional network adapter for VMware ESXi

1. In the AWS Storage Gateway console, power down the gateway.

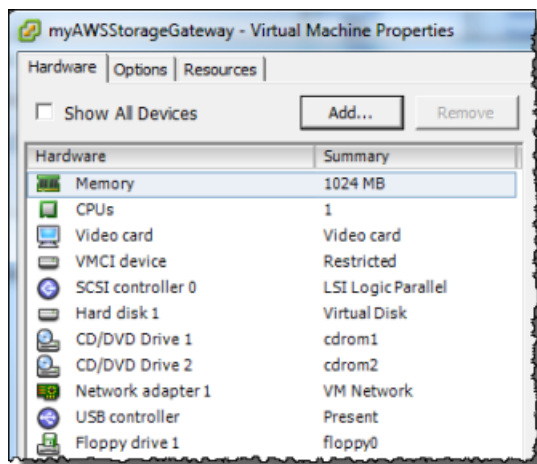
Follow the steps in [To shut down a gateway \(p. 224\)](#), return here, and then go to the next step.
2. In the VMware vSphere client, select your gateway VM.

The VM can remain powered on for this procedure.
3. In the client, right-click the name of your gateway VM and click **Edit Settings**.

AWS Storage Gateway User Guide Configuring Your Gateway for Multiple Network Adapters (NICs)

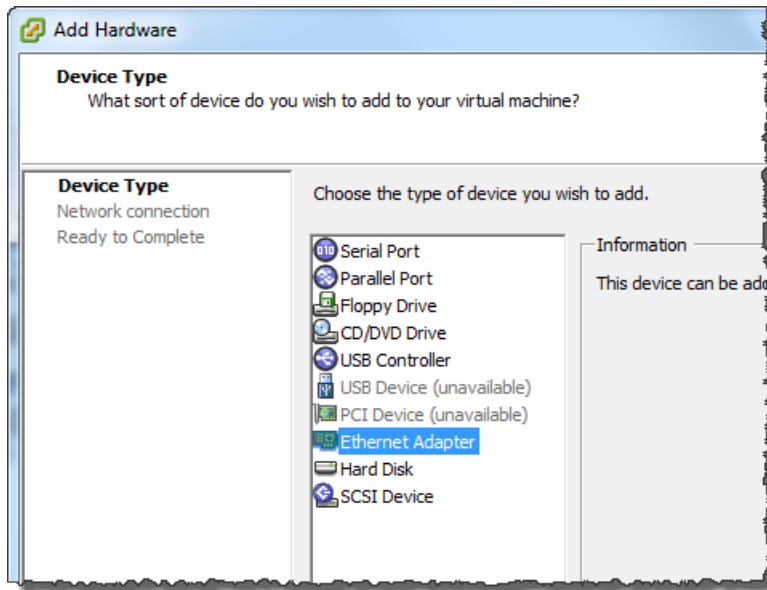


4. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, click **Add** to add a device.



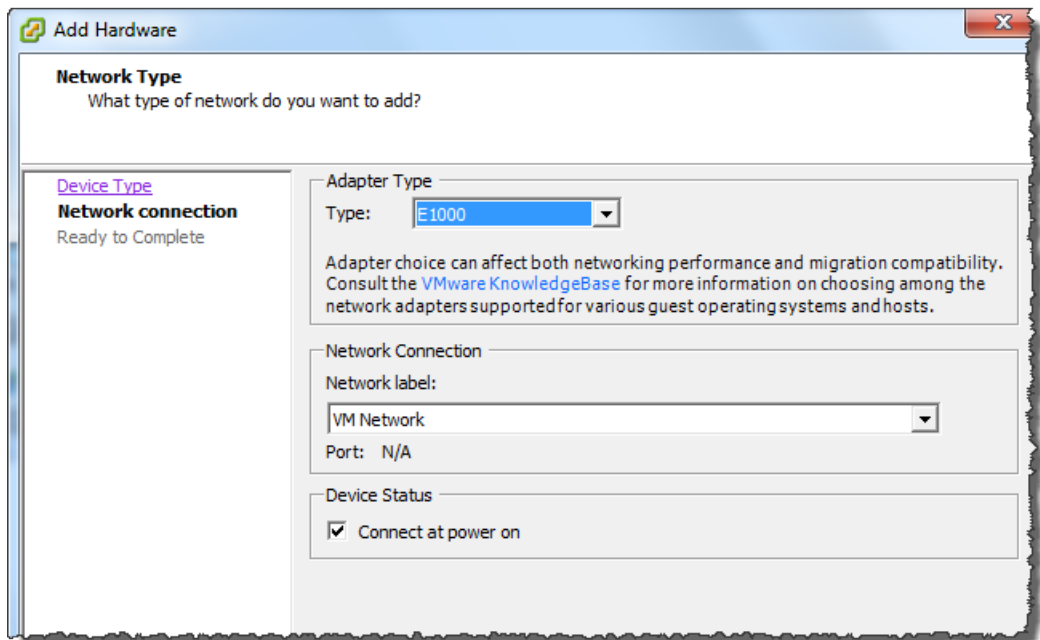
5. Follow the **Add Hardware** wizard to add a network adapter:
 - a. In the **Device Type** pane, click **Ethernet Adapter** to add an adapter, and click **Next**.

AWS Storage Gateway User Guide
Configuring Your Gateway for Multiple Network Adapters
(NICs)



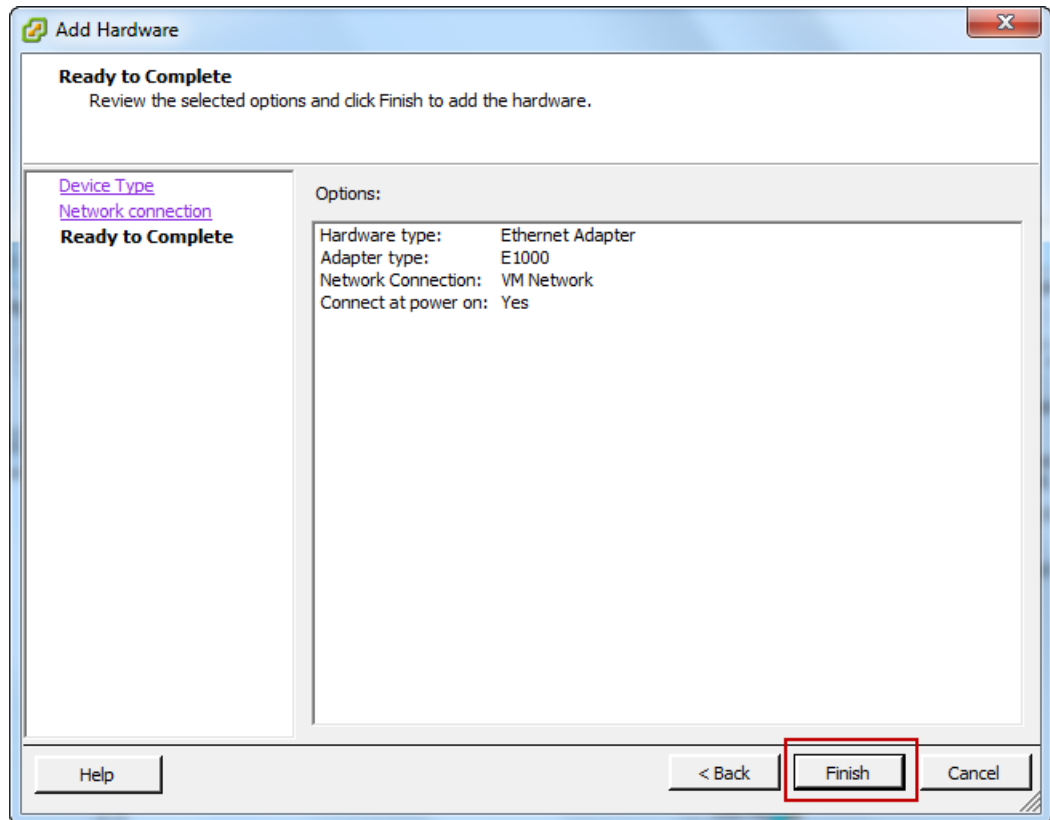
- b. In the **Network Type** pane, in the **Type** drop-down list, select an adapter type, ensure that **Connect at power on** is selected, and then click **Next**.

We recommend that you use the E1000 network adapter with AWS Storage Gateway. For more information on the adapter types that might appear in the adapter list, see *Network Adapter Types* in the [ESXi and vCenter Server Documentation](#).



- c. In the **Ready to Complete** pane, review the information and click **Finish**.

AWS Storage Gateway User Guide
Configuring Your Gateway for Multiple Network Adapters
(NICs)



6. Click the **Summary** tab of the VM, and click **View All** next to the **IP Address** field. A **Virtual Machine IP Addresses** window displays all the IP addresses you can use to access the gateway. Confirm that a second IP address is listed for the gateway.


Note

It might take several moments for the adapter changes to take effect and the VM summary information to refresh.

The following example is for illustration only. In practice, one of the IP addresses would be the address by which the gateway communicates to AWS and the other would be an address in a different subnet.


AWS Storage Gateway User Guide

Configuring Your Gateway for Multiple Network Adapters (NICs)



7. In the AWS Storage Gateway console, power on the gateway.
Follow the steps in [To turn on a gateway \(p. 226\)](#), return here, and then go to the next step.
8. In the AWS Storage Gateway console, in the **Navigation** pane, select the gateway to which you added the adapter and select the **Gateway** tab.

Confirm that the second IP address is listed in the **IP Addresses** field.



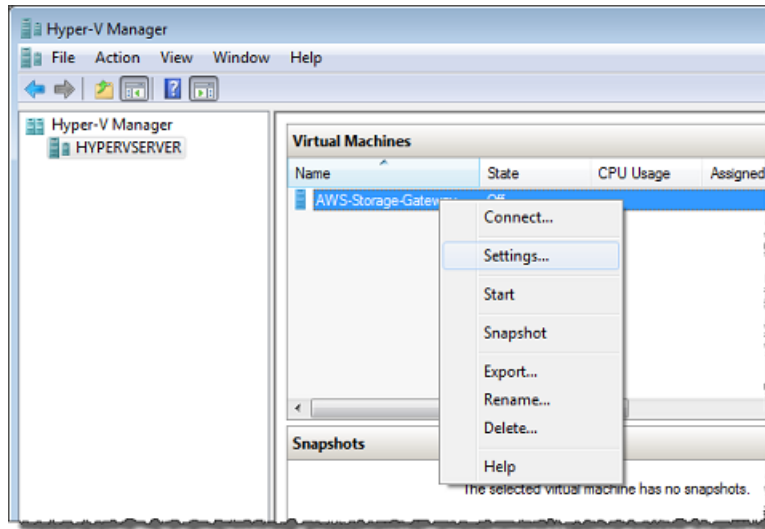
To configure your gateway to use an additional network adapter for Microsoft Hyper-V

1. In the AWS Storage Gateway console, power down the gateway.
Follow the steps in [To shut down a gateway \(p. 224\)](#), return here, and then go to the next step.
2. In the Microsoft Hyper-V Manager, select your gateway VM.
3. Power down the VM if it isn't already.
 - Right-click the gateway and select **Turn Off...**

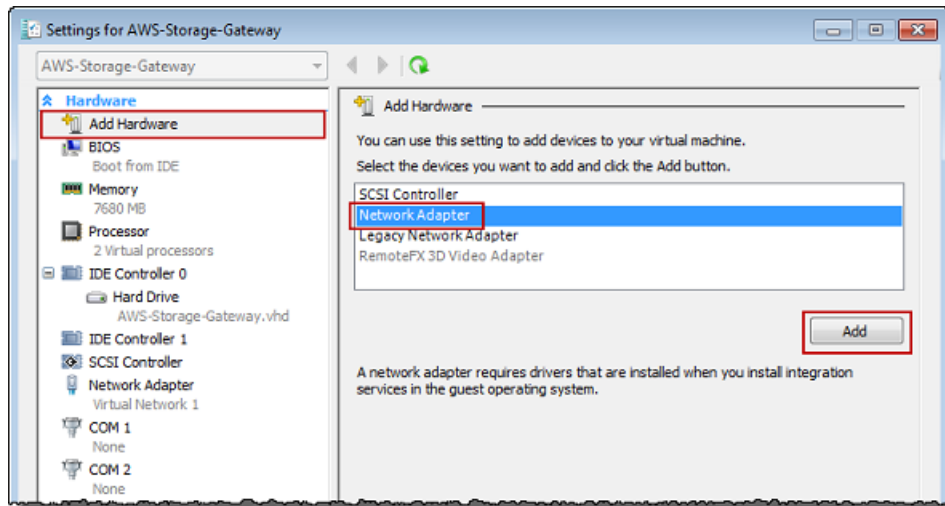
AWS Storage Gateway User Guide

Configuring Your Gateway for Multiple Network Adapters (NICs)

4. In the client, right-click the name of your gateway VM and click **Settings....**



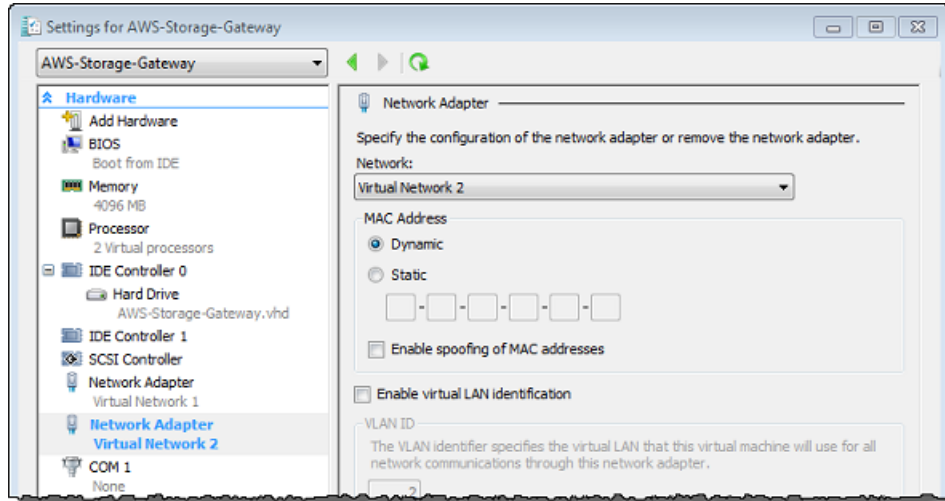
5. In the **Settings** dialog box for the VM, under the **Hardware** list, click **Add Hardware**.
6. In the **Add Hardware** pane, select **Network Adapter**, and click **Add** to add a device.



7. Configure the network adapter and click **Apply**.
- In the following example, **Virtual Network 2** is selected for the new adapter.

AWS Storage Gateway User Guide

Creating a Storage Volume on a Gateway with Multiple Network Adapters

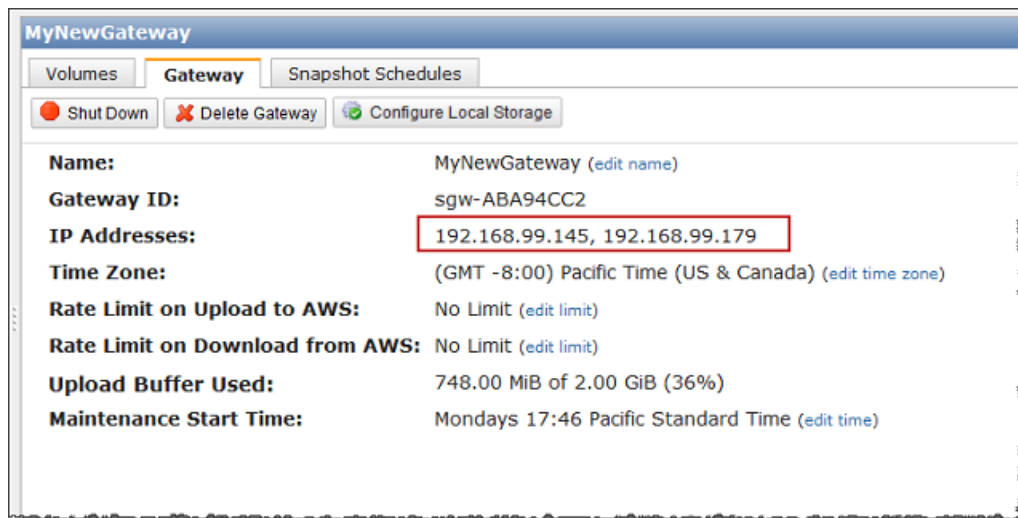


8. In the **Hardware** list of the **Settings** dialog box confirm that the second adapter was added and then click **OK**.
9. In the AWS Storage Gateway console, power on the gateway.

Follow the steps in [To turn on a gateway \(p. 226\)](#), return here, and then go to the next step.

10. In the AWS Storage Gateway console, in the **Navigation** pane, select the gateway to which you added the adapter and select the **Gateway** tab.

Confirm that the second IP address is listed in the **IP Addresses** field.



Creating a Storage Volume in AWS Storage Gateway with Multiple Network Adapters

If you have defined your gateway to use multiple network adapters, then when you create a storage volume for the gateway, you must choose which IP address your storage applications will use to access the storage volume. Each network adapter defined for a gateway will represent one IP address that you

can choose. For information about how to add a network adapter to your gateway, see [Configuring AWS Storage Gateway for Multiple Network Adapters \(NICs\)](#) (p. 245).

To create a storage volume using a specified network adapter.

1. In the AWS Storage Gateway console, in the **Navigation** pane, select the gateway you want to work with and select the **Volumes** tab.
2. Click **Create New Volume**.
3. Configure the storage volume as described in the procedure, [Managing Storage Volumes \(Gateway-Cached\)](#) (p. 181) or [Managing Storage Volumes \(Gateway-Stored\)](#) (p. 183).
4. Select an IP address to use to access the volume.

Note that the **Create Storage Volume** dialog box displays a drop-down list for **Host IP**, one IP address per adapter configured for the gateway VM. If the gateway VM is configured for only one network adapter, the drop-down list is disabled since there is only one IP address.

The screenshot shows the 'Create Storage Volume' dialog box. The 'Disk' dropdown is set to 'SCSI (0:4)' and the 'Preserve existing data' checkbox is checked. The 'iSCSI Target Name' text box contains 'iqn.1997-05.com.amazon:myvolume'. The 'Based on Snapshot ID' text box is empty. The 'Size' dropdown is set to '20 GiB'. The 'Host IP' dropdown is set to '192.168.99.145' and is highlighted with a red rectangular box. The 'Port' text box contains '3260'. At the bottom right, there are 'Cancel' and 'Create Volume' buttons.

5. Click **Create Volume**.

To create a connection to the storage volume, see [Configuring Your Application Access to Storage Volumes](#) (p. 161).

Troubleshooting in AWS Storage Gateway

Topics

- [Troubleshooting On-Premises Gateway Issues](#) (p. 253)
- [Troubleshooting Amazon EC2 Gateway Issues](#) (p. 255)
- [Troubleshooting Storage Volume Issues](#) (p. 256)
- [Using Recovery Snapshots \(Gateway-Cached\)](#) (p. 258)

This section discusses troubleshooting gateway-related and storage volume-related issues. The gateway troubleshooting issues are split into two sections: gateways that are on-premises and gateways that are deployed on Amazon EC2. The on-premises gateway troubleshooting issues cover gateways deployed on both the VMware ESXi or Microsoft Hyper-V platform. The troubleshooting issues for storage volumes apply to all gateway types.

Troubleshooting On-Premises Gateway Issues

The following table lists typical issues that you might encounter working with your on-premises gateways.

Issue	Action to Take
You cannot find the IP address of your gateway.	<p>Use the hypervisor client to connect to your host to find the gateway IP address.</p> <ul style="list-style-type: none">For VMware ESXi, the VM's IP address can be found in the vSphere client on the Summary tab (see Activating AWS Storage Gateway (p. 111)).For Microsoft Hyper-V, the VM's IP address can be found by logging into the local console (see Activating AWS Storage Gateway (p. 131)). <p>If you are still having trouble finding the gateway IP address:</p> <ul style="list-style-type: none">Check that the VM is powered on. Only when the VM is powered on does an IP address get assigned to your gateway.Wait for the VM to finish powering on. If you just powered on your VM, then it may take several minutes for the gateway to finish its boot sequence.
Your gateway's activation fails when you click the Proceed to Activation button in the AWS Storage Gateway console.	<ul style="list-style-type: none">Check that the gateway VM can be accessed by pinging the VM from your client.Check that your VM has network connectivity to the Internet; otherwise, you'll need to configure a SOCKS proxy. For more information, see Routing AWS Storage Gateway Through a Proxy (p. 238).Check that the host has the correct time and is configured to synchronize its time automatically to a Network Time Protocol (NTP) server and that the gateway VM has the correct time. For information about synchronizing the time of hypervisor hosts and VMs, see Synchronizing Your Gateway VM Time (p. 243).After performing these steps, you can retry the gateway deployment using the AWS Storage Gateway console and the Setup and Activate Gateway wizard.Check that your VM has at least 7.5 GB of RAM. Gateway allocation fails if there is less than 7.5 GB of RAM. For more information, see Requirements (p. 6).
You need to remove a disk allocated as upload buffer space because you want to reduce the amount of upload buffer space for a gateway or you need to replace a disk used as an upload buffer that has failed.	<p>For instructions about removing a disk allocated as upload buffer space, see Removing Upload Buffer Capacity (p. 190) or Removing Upload Buffer Capacity (Gateway-Stored) (p. 195).</p>

Issue	Action to Take
<p>You need to improve bandwidth between your gateway and AWS.</p>	<p>You can improve the bandwidth from your gateway to AWS by setting up your Internet connection to AWS on a NIC separate from that of the connection between your applications and the gateway VM. This is useful if you have a high-bandwidth connection to AWS and you want to avoid bandwidth contention, especially during a snapshot restore. For high-throughput workload needs, you can use AWS Direct Connect to establish a dedicated network connection between your on-premises gateway and AWS. To measure the bandwidth of the connection from your gateway to AWS, use the <code>CloudBytesDownloaded</code> and <code>CloudBytesUploaded</code> metrics of the gateway (see Measuring Performance Between Your Gateway and AWS (p. 264)). Improving your Internet connectivity helps to ensure that your upload buffer does not fill up.</p>
<p>Throughput to or from your gateway drops to zero.</p>	<ul style="list-style-type: none"> • In the AWS Storage Gateway console, on the Gateway tab, verify that the IP addresses for your gateway VM are the same that you see using your hypervisor client software (i.e., VMware Vsphere client or Microsoft Hyper-V Manager). If you find a mismatch, restart your gateway from the AWS Storage Gateway console as shown in Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console (p. 224). After the restart, the IP Addresses field in the Gateway tab of the AWS Storage Gateway console should match the IP addresses for your gateway that you determine from the hypervisor client. • For VMware ESXi, the VM's IP address can be found in the vSphere client on the Summary tab (see Activating AWS Storage Gateway (p. 111)). • For Microsoft Hyper-V, the VM's IP address can be found by logging into the local console (see Activating AWS Storage Gateway (p. 131)). • Check your gateway's connectivity to AWS as described in Testing Your AWS Storage Gateway Connection to the Internet (p. 242). • Check your gateway's network adapter configuration and ensure that all the interfaces you intended to be enabled for the gateway are enabled. To view the network adapter configuration for your gateway, follow the instructions for Configuring Your AWS Storage Gateway to Use a Static IP Address (p. 239) and select the option for viewing your gateway's network configuration. <p>You can view the throughput to and from your gateway from the Amazon CloudWatch console. For more information about measuring throughput to and from your gateway to AWS, see Measuring Performance Between Your Gateway and AWS (p. 264).</p>
<p>You are having trouble importing (deploying) AWS Storage Gateway on Microsoft Hyper-V.</p>	<p>See Troubleshooting Your Microsoft Hyper-V Setup (p. 429) in the Appendix. The tips there cover some of the common issues you will run into when deploying a gateway on Microsoft Hyper-V.</p>

Troubleshooting Amazon EC2 Gateway Issues

The following table lists typical issues that you might encounter working with your gateway deployed on Amazon Elastic Compute Cloud (Amazon EC2). For more information about the difference between an on-premises gateway and a gateway deployed in Amazon EC2, see [Deploying and Activating AWS Storage Gateway on Amazon EC2 \(p. 137\)](#).

Issue	Action to Take
<p>Your Amazon EC2 gateway activation fails when you click the Proceed to Activation button in the AWS Storage Gateway console.</p>	<p>If activation has not occurred in a few moments, check the following in the Amazon EC2 console:</p> <ul style="list-style-type: none"> • Port 80 is enabled in the security group you associated to the instance. For more information about adding a security group rule, go to Adding a Security Group Rule in the <i>Amazon EC2 User Guide</i>. • The gateway instance is marked as running. In the Amazon EC2 console, the State of the instance should be "running". <p>After correcting the problem, try activating the gateway again by going to the AWS Storage Gateway console, clicking Deploy a new Gateway on Amazon EC2, and re-entering the IP address of the instance.</p>
<p>You can't find your Amazon EC2 gateway instance in the list of instances.</p>	<p>If, for example, you did not give your instance a resource tag and you have many instances running so that it is hard to tell which instance you deployed the gateway in, you can take the following actions to find the gateway instance:</p> <ul style="list-style-type: none"> • Check the name of the Amazon Machine Image (AMI) name in the Description tab of the instance. An instance based on the AWS Storage Gateway AMI should start with the text "aws-storage-gateway-ami". • If you have several instances based off the AWS Storage Gateway AMI, check the instance launch time to find the correct instance.
<p>You created an Amazon EBS volume but can't attach it to your Amazon EC2 gateway instance.</p>	<p>Check that the Amazon EBS volume in question is in the same Availability Zone as the gateway instance. If there is a discrepancy in Availability Zones, create a new Amazon EBS volume in the same Availability Zone as your instance.</p>
<p>You can't attach an initiator to a storage volume target of your Amazon EC2 gateway.</p>	<p>Check that the security group you launched the instance with includes a rule allowing the port you are using for iSCSI. The port is usually set as 3260. For more information on connecting to storage volumes, see Configuring Your Application Access to Storage Volumes (p. 161).</p>
<p>You activated your Amazon EC2 gateway, but when you go to add storage volumes, you receive an error message indicating you have no disks available.</p>	<p>For a newly activated gateway, no volume storage is defined. Before you can define volume storage, you must allocate local disks to the gateway to use as upload buffer space and cache storage. For a gateway deployed to Amazon EC2, the local disks are Amazon EBS volumes attached to the instance. This error message likely occurs because no Amazon EBS volumes are defined for the instance. Check block devices defined for the instance that is running the gateway. If there are only two block devices (the default devices that come with the AMI), then you should add storage (see Adding and Removing Amazon EBS Volumes from Your Instance (p. 147)). After attaching two or more Amazon EBS volumes, try creating volume storage on the gateway.</p>

Issue	Action to Take
You need to remove a disk allocated as upload buffer space because you want to reduce the amount of upload buffer space.	Follow the steps in To remove an Amazon EBS volume from your Amazon EC2–hosted gateway (p. 148).
Throughput to or from your Amazon EC2 gateway drops to zero.	<ul style="list-style-type: none"> • Verify the gateway instance is running. If the instance is starting, for example, due to a reboot, wait for the instance to restart. • Verify that the gateway IP has not changed. If the instance was stopped and then restarted, the IP address of the instance may change. In this case, you need to reactivate a new gateway. <p>You can view the throughput to and from your gateway from the Amazon CloudWatch console. For more information about measuring throughput to and from your gateway to AWS, see Measuring Performance Between Your Gateway and AWS (p. 264).</p>

Troubleshooting Storage Volume Issues

The following table lists the most typical issues you might encounter when working with storage volumes.

Issue	Action to Take
The AWS Storage Gateway console indicates that your volume has a status of UPLOAD BUFFER NOT CONFIGURED (p. 179).	Add upload buffer capacity to your gateway. You cannot use a gateway to store your application data if the upload buffer for the gateway is not configured. For more information, see Resources: To configure a local disk as upload buffer space for your gateway using the console (p. 190) or To configure an upload buffer for your gateway using the console (p. 195).
The AWS Storage Gateway console indicates that your volume has a status of IRRECOVERABLE (p. 179).	<p>The storage volume is no longer usable. You can try to delete the volume in the AWS Storage Gateway console. If there is data on the volume, then you can recover the data when you create a new storage volume based on the local disk of the VM that was initially used to create the storage volume. When you create the new storage volume, select Preserve existing data. For more information, see Managing Storage Volumes in AWS Storage Gateway (p. 176). Delete pending snapshots of the volume before deleting the storage volume. For more information, see Deleting a Snapshot (p. 209).</p> <p>If deleting the volume in the AWS Storage Gateway console does not work, then the disk allocated for the storage volume may have been improperly removed from the VM and cannot be removed from the appliance.</p>

Issue	Action to Take
<p>The AWS Storage Gateway console indicates that your volume has a status of PASS THROUGH (p. 179).</p>	<p>A volume can be in PASS THROUGH (p. 179) for several reasons. Some of the reasons are a cause for action and some are not.</p> <p>An example of where you should take action when your storage volume is in PASS THROUGH is when your gateway has run out of upload buffer space. To verify if your upload buffer was exceeded in the past, you can view the <code>UploadBufferPercentUsed</code> metric in the Amazon CloudWatch console (see Monitoring the Upload Buffer (p. 267)). If your gateway is in PASS THROUGH because it has run out of upload buffer space, you should allocate more upload buffer space to your gateway. Adding more buffer space will cause your storage volume to transition from PASS THROUGH to BOOTSTRAPPING (p. 179) to AVAILABLE (p. 179) automatically. During BOOTSTRAPPING, the gateway reads data off the storage volume's disk, uploads this data to Amazon S3, and catches up as needed. Once, the gateway has caught up saving the storage volume data to Amazon S3, the volume status becomes AVAILABLE and snapshots can be started again. Note that when your storage volume is in PASS THROUGH or BOOTSTRAPPING, you can continue to read and write data from the storage volume disk. For more information about adding more upload buffer space, see Managing the Upload Buffer (Gateway-Stored) (p. 194).</p> <p>To take action before the upload buffer is exceeded, you can set a threshold alarm on a gateway's upload buffer. For more information, see To set an upper threshold alarm for a gateway's upload buffer (p. 268).</p> <p>Another example of not needing to take action when a storage volume is in PASS THROUGH is when the storage volume is waiting to be bootstrapped because another volume is currently being bootstrapped. The gateway bootstraps volumes one at a time.</p> <p>Infrequently, the PASS THROUGH status can indicate that a disk allocated for an upload buffer has failed. In this is the case, you should remove the disk. For more information, see Removing Upload Buffer Capacity (p. 190) or Removing Upload Buffer Capacity (Gateway-Stored) (p. 195).</p>
<p>Your storage volume's iSCSI target does not show up in the Disk Management Console (Windows).</p>	<p>Check that you have configured the upload buffer for the gateway. For more information, see To configure a local disk as upload buffer space for your gateway using the console (p. 190) or To configure an upload buffer for your gateway using the console (p. 195).</p>
<p>You want to change the iSCSI target name of your storage volume.</p>	<p>The target name is not configurable without deleting the volume and adding it again with a new target name. You can preserve the data on the volume. For information about creating a storage volume, see Managing Storage Volumes (Gateway-Cached) (p. 181) or Managing Storage Volumes (Gateway-Stored) (p. 183).</p>
<p>Your scheduled snapshot of a storage volume did not occur.</p>	<p>Check if your volume is in PASS THROUGH (p. 179), or if the gateway's upload buffer was filled just prior to the time the snapshot was scheduled to be taken. You can check the <code>UploadBufferPercentUsed</code> metric for the gateway in the Amazon CloudWatch console and create an alarm for it. For more information, see Monitoring the Upload Buffer (p. 267) and To set an upper threshold alarm for a gateway's upload buffer (p. 268).</p>

Issue	Action to Take
<p>You need to remove a storage volume because it isn't needed, or you need to replace a storage volume disk that has failed.</p>	<p>You should remove the volume first using the AWS Storage Gateway console (see To remove a storage volume (p. 183)) and then using the hypervisor client to remove the backing storage.</p> <ul style="list-style-type: none"> • For VMware ESXi, remove the backing storage as described in To remove the underlying local disk (VMware ESXi) (p. 185). • For Microsoft Hyper-V, remove the backing storage as describe in To remove the underlying local disk (Microsoft Hyper-V) (p. 187).
<p>Throughput from your application to a storage volume has dropped to zero.</p>	<ul style="list-style-type: none"> • Check that your storage volume's Host IP address matches one of the addresses that appears in the vSphere client on the Summary tab. You can find the Host IP field for a storage volume in the AWS Storage Gateway console in the iSCSI Target Info tab for the storage volume. A discrepancy in the IP address can occur, for example, when you assign a new static IP address to your gateway. If there is a discrepancy, restart your gateway from the AWS Storage Gateway console as shown in Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console (p. 224). After the restart, the Host IP address in the iSCSI Target Info tab for a storage volume should match an IP address shown in the vSphere client on the Summary tab for the gateway. • Check to see if IPAddressNotFound appears in the Host IP field for the storage volume. This can occur, for example, when you create a storage volume associated with an IP address of a network adapter of a gateway that is configured with two or more network adapters. When you remove or disable the network adapter that the storage volume is associated with, the IPAddressNotFound message is displayed. To address this issue, delete the storage volume and then re-create it preserving its existing data. For more information, see Managing Storage Volumes in AWS Storage Gateway (p. 176). • Check that the iSCSI initiator your application uses is correctly mapped to the iSCSI target for the storage volume. For more information about connecting to storage volumes, see Configuring Your Application Access to Storage Volumes (p. 161). <p>You can view the throughput for storage volumes and create alarms from the Amazon CloudWatch console. For more information about measuring throughput from your application to a storage volume, see Measuring Performance Between Your Application and Gateway (p. 262).</p>

Using Recovery Snapshots (Gateway-Cached)

AWS Storage Gateway provides recovery points for each volume in a gateway-cached volume architecture. A volume recovery point is a point in time at which all data of the volume is consistent and from which you can create a snapshot. You can use the snapshot to create a new volume in the event that your gateway becomes unreachable or one gateway-cached volume becomes irrecoverable.

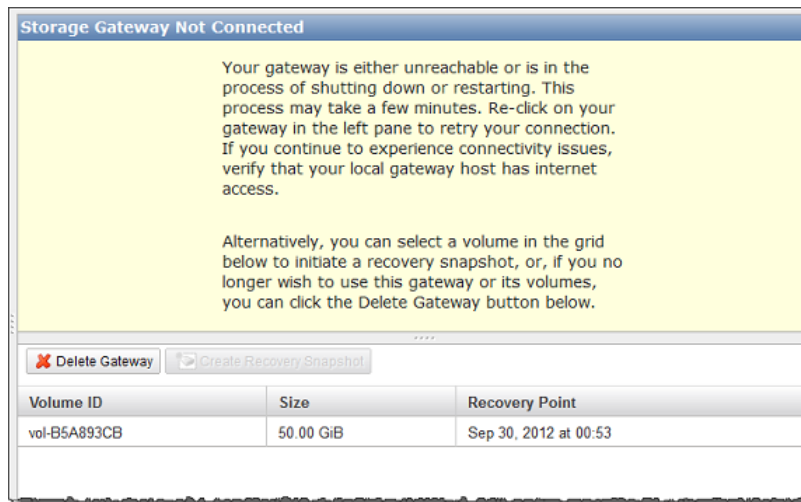
Volume recovery points are maintained automatically for each gateway-cached volume. You can also take snapshots on an ad-hoc basis or set up a snapshot schedule for the volume. For more information about snapshots, see [Working with Snapshots \(p. 199\)](#).

When the gateway becomes unreachable (such as when you shut it down), you have the option of creating a snapshot from a volume recovery point.

To create and use a recovery snapshot of a volume from an unreachable gateway

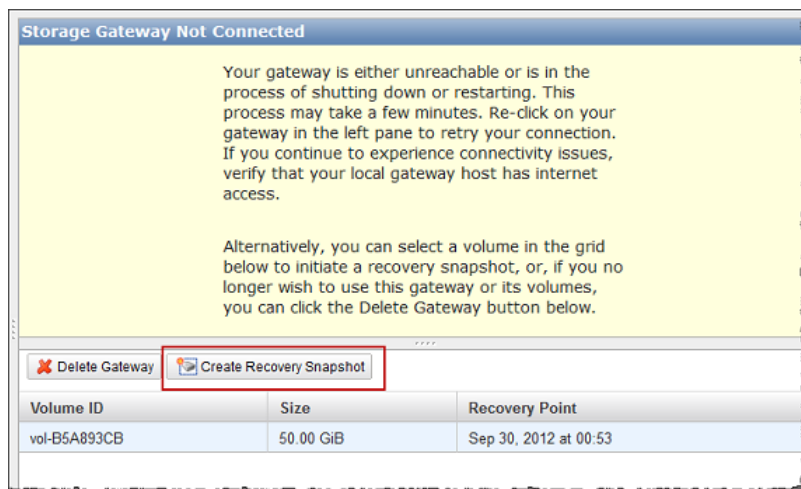
1. In the **AWS Storage Gateway** console navigation pane, select the unreachable gateway.

A list of volumes of the gateway is displayed.




2. Select the volume from which to create a recovery snapshot.
3. Click the **Create Recovery Snapshot** button.

AWS Storage Gateway initiates the snapshot process.



4. In the **Recovery Point Created** dialog box, click **Close**.



5. Find the snapshot using the steps in the procedure [Finding a Snapshot \(p. 200\)](#).
6. Restore the snapshot using one of the procedures in [Restoring a Snapshot \(p. 219\)](#).

Optimizing AWS Storage Gateway Performance

This section provides information about how to optimize the performance of your gateway. The guidance is based on adding resources to your gateway and adding resources to your application server.

Add Resources to Your Gateway

- **Use Higher Performance Disks**—You can add high performance disks such as Serial Attached SCSI (SAS) disks and Solid-State Drives (SSDs), or you can attach virtual disks to your VM directly from a SAN instead of through VMware's VMFS layer or Microsoft Hyper-V's NTFS. Improved disk performance generally results in better throughput and input/output operations per second (IOPS). To measure throughput, use the `ReadBytes` and `WriteBytes` metrics with the `Samples` Amazon CloudWatch statistic. For example, the `Samples` statistic of the `ReadBytes` metric over a sample period of five minutes divided by 300 seconds, gives you the input/output operations per second (IOPS). As a general rule, when you review these metrics for a gateway, look for low throughput and low IOPS trends to indicate disk-related bottlenecks. For more information about gateway metrics, see [Measuring Performance Between Your Gateway and AWS \(p. 264\)](#).
- **Add CPU Resources to Your Gateway Host**—The minimum requirement for a gateway host server is four virtual processors. You should confirm that the four virtual processors that are assigned to the gateway VM are backed by four cores and that you are not oversubscribing the CPUs of the host server. When you add additional CPUs to your gateway host server, you increase the processing capability of the gateway to deal with, in parallel, both storing data from your application to your local storage and uploading this data to Amazon S3. Additional CPUs also ensure that your gateway gets enough CPU resources when the host is shared with other VMs. This has the general effect of improving throughput.
- **Change the Storage Volumes Configuration**—If you find that adding more storage volumes to a gateway reduces the throughput to the gateway, then you can consider adding the storage volume to a separate gateway. In particular, if the storage volume is used for a high-throughput application, then you should consider creating a separate gateway for the high-throughput application. However, as a general rule, you should not use one gateway for all of your high-throughput applications and another gateway for all of your low-throughput applications. To measure your storage volume throughput, use the `ReadBytes` and `WriteBytes` metrics (see [Measuring Performance Between Your Application and Gateway \(p. 262\)](#)).
- **Back Gateway Virtual Disks Using Separate Physical Disks**—When you provision disks in a gateway-cached volume setup, we strongly recommend that you do not provision local disks for upload buffer and cache storage that use the same underlying physical storage disk. Similarly, for a gateway-stored volume setup, we strongly recommend that you do not provision local disks for upload buffer and application storage that use the same underlying physical storage disk. For example, for VMware ESXi, the underlying physical storage resources are represented as a datastore in VMware ESXi. When you deploy the gateway VM, you choose a datastore on which to store the VM files. When you provision a virtual disk (e.g., to use as upload buffer), you have the option to store the virtual disk in the same datastore as the VM or a different datastore. If you have more than one datastore, then it is strongly recommended that you choose one datastore for each type of local storage you are creating. A datastore that is backed by only one underlying physical disk, or that is backed by a less-performant RAID configuration such as RAID 1, may lead to poor performance in some situations when, for example, used to back both the cache storage and upload buffer in a gateway-cached volume setup.

Add Resources to Your Application Environment

- **Increase the Bandwidth Between Your Application Server and Your Gateway**—Ensure that the network bandwidth between your application and the gateway can sustain your application needs. You can use the `ReadBytes` and `WriteBytes` metrics of the gateway (see [Measuring Performance Between Your Gateway and AWS](#) (p. 264)) to measure the total data throughput. Compare the measured throughput with the desired throughput (specific to your application). If the measured throughput is less than the desired throughput, then increasing the bandwidth between your application and gateway can improve performance if the network is the bottleneck. Similarly, you can increase the bandwidth between your VM and your local disks (if they're not direct-attached).
- **Add CPU Resources to Application Environment**—If your application can make use of additional CPU resources, then adding more CPUs may allow your application to scale its IO load.

Monitoring Your AWS Storage Gateway

Topics

- [Using the Amazon CloudWatch Console](#) (p. 261)
- [Measuring Performance Between Your Application and Gateway](#) (p. 262)
- [Measuring Performance Between Your Gateway and AWS](#) (p. 264)
- [Monitoring the Upload Buffer](#) (p. 267)
- [Monitoring Cache Storage](#) (p. 271)
- [Understanding AWS Storage Gateway Metrics](#) (p. 272)

In this section, we discuss how to monitor your gateway, including its volumes associated with the gateway (both gateway-cached and gateway-stored) and the upload buffer. You use the AWS Management Console to view metrics for your gateway. For example, you can view the number of bytes used in read and write operations, the time spent in read and write operations, and the time to retrieve data from the AWS cloud. With metrics, you can track the health of your gateway and set up alarms to notify you when one or more metrics are outside a defined threshold.

AWS Storage Gateway provides Amazon CloudWatch metrics at no additional charge. AWS Storage Gateway metrics are recorded for a period of two weeks, allowing you access to historical information and providing you with a better perspective of how your gateway and volumes are performing. For detailed information about Amazon CloudWatch, go to the [Amazon CloudWatch Developer Guide](#).

Using the Amazon CloudWatch Console

You can get monitoring data for your gateway using either the AWS Management Console or the Amazon CloudWatch API. The console displays a series of graphs based on the raw data from the Amazon CloudWatch API. The Amazon CloudWatch API can be also be used through one of the [Amazon AWS Software Development Kits \(SDKs\)](#) or the [Amazon CloudWatch API](#) tools. Depending on your needs, you might prefer to use either the graphs displayed in the console or retrieved from the API.

Regardless of which method you choose to use to work with metrics, you must specify the following information.

- First, you specify the metric dimension to work with. A dimension is a name-value pair that helps you to uniquely identify a metric. The dimensions for AWS Storage Gateway are `GatewayId`, `GatewayName`, and `VolumeId`. In the Amazon CloudWatch console, the `Gateway Metrics` and `Volume Metrics` views are provided to easily select gateway and volume-specific dimensions. For more information about dimensions, see [Dimensions](#) in the *Amazon CloudWatch Developer Guide*.
- Second, you specify the metric name, such as `ReadBytes`.

Tip

If the name of your gateway was changed for the time range that you are interested in viewing metrics, then you should use the `GatewayId` to specify the metrics for your analysis.

The following table summarizes the types of AWS Storage Gateway metric data that are available to you.

Amazon CloudWatch Namespace	Dimension	Description
AWS/StorageGateway	GatewayId, GatewayName	<p>These dimensions filter for metric data that describes aspects of the gateway. You can identify a gateway to work with either the <code>GatewayId</code> or the <code>GatewayName</code>.</p> <p>Throughput and latency data of a gateway is based on all the volumes in the gateway.</p> <p>Data is available automatically in 5-minute periods at no charge.</p>
	VolumeId	<p>This dimension filters for metric data that is specific to a storage volume. Identify a storage volume to work with by <code>VolumeId</code>.</p> <p>Data is available automatically in 5-minute periods at no charge.</p>

Working with gateway and volume metrics is similar to working with other service metrics. Many of the common tasks are outlined in the Amazon CloudWatch documentation and are listed below for your convenience:

- [Listing Available Metrics](#)
- [Getting Statistics for a Metric](#)
- [Creating CloudWatch Alarms](#)

Measuring Performance Between Your Application and Gateway

Data throughput, data latency, and operations per second are three measures that you can use to understand how your application storage using the AWS Storage Gateway is performing. These three values can be measured using the AWS Storage Gateway metrics that are provided for you when you use the correct aggregation statistic. A statistic is an aggregation of a metric over a specified period of time. When you view the values of a metric in Amazon CloudWatch, use the `Average` statistic for data latency (milliseconds), use the `Sum` statistic for data throughput (bytes per second), and use the `Samples` statistic for operations per second (IOPS). For more information, see [Statistics](#) in the *Amazon CloudWatch Developer Guide*.

The following table summarizes the metrics and corresponding statistic to use to measure the throughput, latency, and IOPS between your applications and gateways.

AWS Storage Gateway User Guide
Measuring Performance Between Your Application and Gateway

Item of Interest	How to Measure
Throughput	Use the <code>ReadBytes</code> and <code>WriteBytes</code> metrics with the <code>Sum</code> Amazon CloudWatch statistic. For example, the <code>Sum</code> of the <code>ReadBytes</code> over a sample period of five minutes divided by 300 seconds, gives you the throughput as bytes/second rate.
Latency	Use the <code>ReadTime</code> and <code>WriteTime</code> metrics with the <code>Average</code> Amazon CloudWatch statistic. For example, the <code>Average</code> of the <code>ReadTime</code> gives you the latency per operation over the sample period of time.
IOPS	Use the <code>ReadBytes</code> and <code>WriteBytes</code> metrics with the <code>Samples</code> Amazon CloudWatch statistic. For example, the <code>Samples</code> of the <code>ReadBytes</code> over a sample period of five minutes divided by 300 seconds, gives you input/output operations per second (IOPS).


For the average latency graphs and average size graphs, the average is calculated over the total number of operations (read or write, whichever is applicable to the graph) that completed during the period.

The following tasks assume that you are starting in the Amazon CloudWatch console.

To measure the data throughput from an application to a storage volume

1	Select the StorageGateway: Volume Metrics dimension and find the storage volume that you want to work with.
2	Select the <code>ReadBytes</code> and <code>WriteBytes</code> metrics.
3	Select a Time Range .
4	Select the <code>Sum</code> statistic.
5	Select a Period of 5 minutes or greater.
6	In the resulting time-ordered sets of data points (one for <code>ReadBytes</code> and one for <code>WriteBytes</code>), divide each data point by the period (in seconds) to get the throughput at the sample point. The total throughput is the sum of the throughputs.


The following example shows the `ReadBytes` and `WriteBytes` metrics for a storage volume with the `Sum` statistic. In the example, the cursor over a data point displays information about the data point including its value and the number of bytes. Divide the bytes value by the **Period** (5 minutes) to get the data throughput at that sample point. For the point highlighted, the read throughput is 2,384,199,680 bytes divided by 300 seconds, which is 7.6 MB/s.



To measure the data input/output operations per second from an application to a storage volume

1	Select the StorageGateway: Volume Metrics dimension and find the storage volume that you want to work with.
2	Select the <code>ReadBytes</code> and <code>WriteBytes</code> metrics.
3	Select a Time Range .
4	Select the <code>Samples</code> statistic.
5	Select a Period of 5 minutes or greater.
6	In the resulting time-ordered sets of data points (one for <code>ReadBytes</code> and one for <code>WriteBytes</code>), divide each data point by the period (in seconds) to get the input/output operations per second.

The following example shows the `ReadBytes` and `WriteBytes` metrics for a storage volume with the `Samples` statistic. In the example, the cursor over a data point displays information about the data point, including its value and the number of samples. Divide the samples value by the **Period** (5 minutes) to get the operations per second at that sample point. For the point highlighted, the number of write operations is 24,373 bytes divided by 300 seconds, which is 81 write operations per second.



Measuring Performance Between Your Gateway and AWS

Data throughput, data latency, and operations per second are three measures that you can use to understand how your application storage using the AWS Storage Gateway is performing. These three values can be measured using the AWS Storage Gateway metrics provided for you when you use the correct aggregation statistic. The following table summarizes the metrics and corresponding statistic to use to measure the throughput, latency, and IOPS between your gateway and AWS.

Item of Interest	How to Measure
Throughput	Use the <code>ReadBytes</code> and <code>WriteBytes</code> metrics with the <code>Sum</code> Amazon CloudWatch statistic. For example, the <code>Sum</code> of the <code>ReadBytes</code> over a sample period of five minutes divided by 300 seconds, gives you the throughput as bytes/second rate.
Latency	Use the <code>ReadTime</code> and <code>WriteTime</code> metrics with the <code>Average</code> Amazon CloudWatch statistic. For example, the <code>Average</code> of the <code>ReadTime</code> gives you the latency per operation over the sample period of time.

AWS Storage Gateway User Guide
Measuring Performance Between Your Gateway and AWS


Item of Interest	How to Measure
IOPS	Use the <code>ReadBytes</code> and <code>WriteBytes</code> metrics with the <code>Samples</code> Amazon CloudWatch statistic. For example, the <code>Samples</code> of the <code>ReadBytes</code> over a sample period of five minutes divided by by 300 seconds, gives you the input/output operations per second (IOPS).
Throughput to AWS	Use the <code>CloudBytesDownloaded</code> and <code>CloudBytesUploaded</code> metrics with the <code>Sum</code> Amazon CloudWatch statistic. For example, the <code>Sum</code> of the <code>CloudBytesDownloaded</code> over a sample period of five minutes divided by 300 seconds, gives you the throughput from AWS to the gateway as bytes/per second.
Latency of data to AWS	Use the <code>CloudDownloadLatency</code> metric with the <code>Average</code> statistic. For example, the <code>Average</code> statistic of the <code>CloudDownloadLatency</code> metric gives you the latency per operation.

The following tasks assume that you are starting in the Amazon CloudWatch console.

To measure the upload data throughput from a gateway to AWS

1	Select the StorageGateway: Gateway Metrics dimension and find the gateway that you want to work with.
2	Select the <code>CloudBytesUploaded</code> metric.
3	Select a Time Range .
4	Select the <code>Sum</code> statistic.
5	Select a Period of 5 minutes or greater.
6	In the resulting time-ordered set of data points, divide each data point by the period (in seconds) to get the throughput at that sample period.

The following example shows the `CloudBytesUploaded` metric for a gateway volume with the `Sum` statistic. In the example, the cursor over a data point displays information about the data point, including its value and bytes uploaded. Divide this value by the **Period** (5 minutes) to get the throughput at that sample point. For the point highlighted, the throughput from the gateway to AWS is 555,544,576 bytes divided by 300 seconds, which is 1.7 MB/s.



To measure the latency per operation of a gateway

1	Select the StorageGateway: Gateway Metrics dimension and find the gateway that you want to work with.
2	Select the <code>ReadTime</code> and <code>WriteTime</code> metrics.
3	Select a Time Range .
4	Select the <code>Average</code> statistic.
5	Select a Period of 5 minutes to match the default reporting time.
6	In the resulting time-ordered set of points (one for <code>ReadTime</code> and one for <code>WriteTime</code>), add data points at the same time sample to get to the total latency in milliseconds.

To measure the data latency from a gateway to AWS

1	Select the StorageGateway: GatewayMetrics dimension and find the gateway that you want to work with.
2	Select the <code>CloudDownloadLatency</code> metric.
3	Select a Time Range .
4	Select the <code>Average</code> statistic.
5	Select a Period of 5 minutes to match the default reporting time.
6	The resulting time-ordered set of data points contains the latency in milliseconds.

To set an upper threshold alarm for a gateway's throughput to AWS

1	Start the Create Alarm Wizard .
2	Select the StorageGateway: Gateway Metrics dimension and find the gateway that you want to work with.
3	Select the <code>CloudBytesUploaded</code> metric.
4	Define the alarm by defining the alarm state when the <code>CloudBytesUploaded</code> metric is greater than or equal to a specified value for a specified time. For example, you can define an alarm state when the <code>CloudBytesUploaded</code> metric is greater than 10 MB for 60 minutes.
5	Configure the actions to take for the alarm state.
6	Create the alarm.

To set an upper threshold alarm for reading data from AWS

1	Start the Create Alarm Wizard .
2	Select the StorageGateway: Gateway Metrics dimension and find the gateway that you want to work with.
3	Select the <code>CloudDownloadLatency</code> metric.

4	Define the alarm by defining the alarm state when the <code>CloudDownloadLatency</code> metric is greater than or equal to a specified value for a specified time. For example, you can define an alarm state when the <code>CloudDownloadLatency</code> is greater than 60,000 milliseconds for greater than 2 hours.
5	Configure the actions to take for the alarm state.
6	Create the alarm.

Monitoring the Upload Buffer

The following section discusses how to monitor a gateway's upload buffer and how to create an alarm so that you get a notification when the buffer exceeds a specified threshold. This enables you to proactively add buffer storage to a gateway before it fills completely and your storage application stops backing up to AWS.

Monitoring upload buffer applies to both the gateway-cached and gateway-stored architectures. For more information, see [How AWS Storage Gateway Works \(p. 3\)](#)

Note

`WorkingStoragePercentUsed`, `WorkingStorageUsed`, and `WorkingStorageFree` metrics represent the upload buffer for only the gateway-stored volume setup prior to the release of the cached-volume feature in AWS Storage Gateway. Now you should use the equivalent upload buffer metrics: `UploadBufferPercentUsed`, `UploadBufferUsed`, and `UploadBufferFree` which apply to both gateway architectures.

Item of Interest	How to Measure
Upload buffer usage	Use the <code>UploadBufferPercentUsed</code> , <code>UploadBufferUsed</code> , <code>UploadBufferFree</code> metrics with the <code>Average</code> statistic. For example, use the <code>UploadBufferUsed</code> with the <code>Average</code> statistic to analyze the storage usage over a time period.

The following tasks assume that you are starting in the Amazon CloudWatch console.


To measure upload buffer percent used

1	Select the StorageGateway: Gateway Metrics dimension and find the gateway that you want to work with.
2	Select the <code>UploadBufferPercentUsed</code> metric.
3	Select a Time Range .
4	Select the <code>Average</code> statistic.
5	Select a Period of 5 minutes to match the default reporting time.
6	The resulting time-ordered set of data points that contains the percent used of upload buffer.

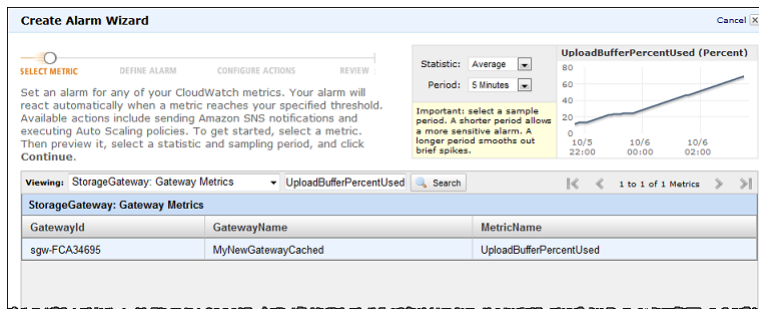
The upload buffer task below shows you how to create an alarm using the Amazon CloudWatch console and the **Create Alarm Wizard**. To learn more about alarms and thresholds, see [Creating CloudWatch Alarms](#).

To set an upper threshold alarm for a gateway's upload buffer

1. Start the **Create Alarm Wizard**.
 - a. In the Amazon CloudWatch console, click the **Alarms** link in the **Navigation** pane.
 - b. In the **Your CloudWatch Alarms** pane, click **Create Alarm**.



2. Specify a metric for your alarm.
 - a. In the **SELECT METRIC** page of the **Create Alarm Wizard**, select the **AWS/StorageGateway:GatewayId,GatewayName** dimension and find the gateway that you want to work with.
 - b. Select the `UploadBufferPercentUsed` metric. Use the **Average** statistic and a period of 5 minutes.



- c. Click **Continue**.
3. Define the alarm name, description, and threshold.
 - a. In the **DEFINE ALARM** page of the **Create Alarm Wizard**, identify your alarm by giving it a name and description in the **Name** and **Description** fields, respectively.
 - b. Define the alarm threshold.

In the example below, the alarm state is defined for `UploadBufferPercentUsed` greater than or equal to 50 percent for 5 minutes.

AWS Storage Gateway User Guide

Monitoring the Upload Buffer

Create Alarm Wizard Cancel X

SELECT METRIC **DEFINE ALARM** CONFIGURE ACTIONS REVIEW

Provide the details and threshold for your alarm. Use the graph below to help set the appropriate threshold.

Identify Your Alarm
Assign your alarm a name and description.

Name: Alarm Name
Description: Alarm Description

Define Alarm Threshold
Alarms have three states: ALARM, OK, and INSUFFICIENT DATA. The state of your alarm changes according to a threshold you specify. First, define the criterion for entering the ALARM state. Later, you can specify an action to be taken when your alarm enters any of the three states.

This alarm will enter the ALARM state when UploadBufferPercentUsed is \geq 50 for 5 minutes.

Metric: UploadBufferPercentUsed
Period: 5 Minutes
Statistic: Average

UploadBufferPercentUsed (Percent)

Back Continue

c. Click **Continue**.

4. Configure an email action for the alarm.

- In the **CONFIGURE ACTIONS** page of the **Create Alarm Wizard**, select **ALARM** from the **Alarm State** drop-down list.
- Select **Select or create email topic...** from the **Topic** drop-down list.

Define an email topic means you set up an Amazon SNS topic. For more information about Amazon SNS, see [Set Up Amazon SNS](#).

- In the **Topic** field, enter a descriptive name for the topic.
- Click **ADD ACTION**.

Edit Alarm Wizard Cancel X

SELECT METRIC DEFINE ALARM **CONFIGURE ACTIONS** REVIEW

Define what actions are taken when your alarm changes.

You can define multiple actions for a single alarm. For example, you may want to scale out your fleet and send an email to your pager when this alarm enters the ALARM state, and then send another all-clear email when it returns to the OK state.

Define Your Actions
Actions define what steps you want to automate when the alarm state changes. For example, you can send a message using email via the Simple Notification Service (SNS). You can also execute an Auto Scaling Policy, if you have one configured ([learn about policies](#)).

When Alarm state is	Take action	Action details	
ALARM	Send Notification	Topic: my-alarm-topic Email(s): user@example.com	ADD ACTION

Back Continue

e. Click **Continue**.

5. Review the alarm settings and create the alarm.

- In the **REVIEW** page of the **Create Alarm Wizard**, review the alarm definition, metric, and associated actions from this step.

AWS Storage Gateway User Guide Monitoring the Upload Buffer

Edit Alarm Wizard Cancel X

SELECT METRIC DEFINE ALARM CONFIGURE ACTIONS **REVIEW**

Please review the alarm information below. If you would like to proceed with this configuration, click **Save Alarm**. If you want to make any changes to this alarm, click **Back** or select a step on the right to edit.

Alarm Definition Edit Definition

Name: Alarm Name
Description: Alarm Description
In ALARM state when: the value is >= 50 for 5 minutes

Metric Edit Metric

Namespace: AWS/StorageGateway
MetricName: UploadBufferPercentUsed
GatewayName: MyNewGatewayCached
GatewayId: sgw-FCA34695
Period / Statistic: 5 Minutes / Average

Alarm Actions Edit Actions

Actions:

When alarm state is *ALARM*
Action Type: Send Notification to New Topic
Action: Notify topic: my-alarm-topic(user@example.com)

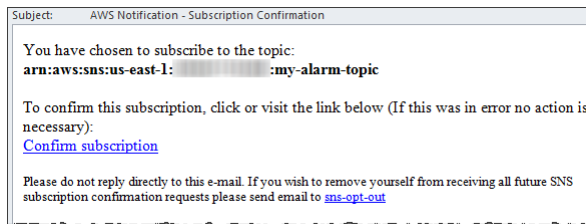
Back Save Alarm

b. After reviewing the alarm summary, click **Save Alarm**.

6. Confirm your subscription to the alarm topic.

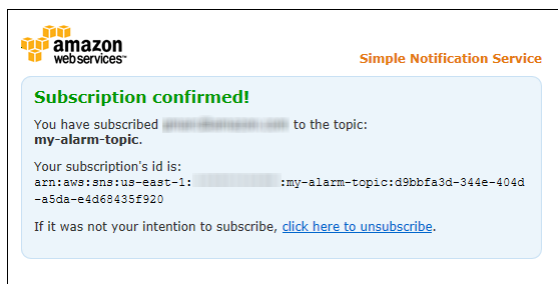
a. Open the Amazon Simple Notification Service (Amazon SNS) email that is sent to the email address that you specified when creating the topic.

The example below shows a notification.



b. Confirm your subscription by clicking the link in the email.

A subscription confirmation displays.



Monitoring Cache Storage

The following section discusses how to monitor a gateway's cache storage and how to create an alarm so that you get a notification when parameters of cache pass specified thresholds. This enables you to proactively add cache storage to a gateway.

Monitoring cache storage applies only to the gateway-cached architecture. For more information, see [How AWS Storage Gateway Works \(p. 3\)](#).

Item of Interest	How to Measure
Total usage of cache	Use the <code>CachePercentUsed</code> and <code>TotalCacheSize</code> metrics with the <code>Average</code> statistic. For example, use the <code>CachePercentageUsed</code> with the <code>Average</code> statistic to analyze the cache usage over a time. The <code>TotalCacheSize</code> metric changes only when you add cache to the gateway.
Percentage of read requests that are served from the cache	Use the <code>CacheHitPercent</code> metric with the <code>Average</code> statistic. Typically, you will want <code>CacheHitPercent</code> to remain high.
Percentage of cache that is dirty, that is, it contains content that has not been upload to AWS	Use the <code>CachePercentDirty</code> metrics with the <code>Average</code> statistic. Typically, you want <code>CachePercentDirty</code> to remain low.

The following tasks assume that you are starting in the Amazon CloudWatch console.

To measure cache percent dirty for a gateway and all its volumes

1	Select the StorageGateway: Gateway Metrics dimension and find the gateway that you want to work with.
2	Select the <code>CachePercentDirty</code> metric.
3	Select a Time Range .
4	Select the <code>Average</code> statistic.
5	Select a Period of 5 minutes to match the default reporting time.
6	The resulting time-ordered set of data points that contains the percent used of cache percent dirty.

To measure cache percent dirty for a volume

1	Select the StorageGateway: Volume Metrics dimension and find the volume that you want to work with.
2	Select the <code>CachePercentDirty</code> metric.
3	Select a Time Range .
4	Select the <code>Average</code> statistic.
5	Select a Period of 5 minutes to match the default reporting time.

6	The resulting time-ordered set of data points that contains the percent used of cache percent dirty.
---	--

Understanding AWS Storage Gateway Metrics

Topics

- [Gateway Metrics \(p. 272\)](#)
- [Storage Volume Metrics \(p. 274\)](#)

Gateway Metrics

For the discussion here, we define *gateway* metrics as metrics that are scoped to the gateway—that is, they measure something about the gateway. Since a gateway contains one or more volumes, a gateway-specific metric is representative of all volumes on the gateway. For example, the `CloudBytesUploaded` metric is the total number of bytes that the gateway sent to the cloud during the reporting period. This includes the activity of all the volumes on the gateway.

When working with gateway metric data, you will specify the unique identification of the gateway that you are interested in viewing metrics for. To do this, you can either specify the `GatewayId` or the `GatewayName`. When you want to work with metric for a gateway, you specify the gateway *dimension* in the metrics namespace, which distinguishes a gateway-specific metric from a volume-specific metric. For more information, see [Using the Amazon CloudWatch Console \(p. 261\)](#).

The following table describes the AWS Storage Gateway metrics that you can use to get information about your gateway. The entries in the table are grouped functionally by measure.

Metric	Description
<code>CacheHitPercent</code>	Percent of application reads served from the cache. This metric applies only to the gateway-cached volume setup. The sample is taken at the end of the reporting period. Units: Percent
<code>CachePercentageUsed</code>	Percent utilization of the gateway's cache storage. This metric applies only to the gateway-cached volume setup. The sample is taken at the end of the reporting period. Units: Percent
<code>CachePercentDirty</code>	Percent of the gateway's cache that has not been persisted to AWS. This metric applies only to the gateway-cached volume setup. The sample is taken at the end of the reporting period. Units: Percent
<code>CloudBytesDownloaded</code>	The total number of pre-compressed bytes that the gateway downloaded from AWS during the reporting period. Use this metric with the <code>Sum</code> statistic to measure throughput and with the <code>Samples</code> statistic to measure operations per second (IOPS). Units: Bytes

AWS Storage Gateway User Guide
Understanding AWS Storage Gateway Metrics

Metric	Description
CloudDownloadLatency	<p>The total number of milliseconds spent reading data from AWS during the reporting period.</p> <p>Use this metric with the <i>Average</i> statistic to measure latency.</p> <p>Units: Milliseconds</p>
CloudBytesUploaded	<p>The total number of pre-compressed bytes that the gateway uploaded to AWS during the reporting period.</p> <p>Use this metric with the <i>Sum</i> statistic to measure throughput and with the <i>Samples</i> statistic to measure operations per second (IOPS).</p> <p>Units: Bytes</p>
UploadBufferFree	<p>The total amount of unused space in the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Units: Bytes</p>
UploadBufferPercentUsed	<p>Percent utilization of the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Units: Percent</p>
UploadBufferUsed	<p>The total number of bytes being used in the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Units: Bytes</p>
QueuedWrites	<p>The number of bytes waiting to be written to AWS, sampled at the end of the reporting period for all volumes in the gateway. These bytes are kept in your gateway's working storage.</p> <p>Units: Bytes</p>
ReadBytes	<p>The total number of bytes read from your on-premises applications in the reporting period for all volumes in the gateway.</p> <p>Use this metric with the <i>Sum</i> statistic to measure throughput and with the <i>Samples</i> statistic to measure operations per second (IOPS).</p> <p>Units: Bytes</p>
ReadTime	<p>The total number of milliseconds spent to do reads from your on-premises applications in the reporting period for all volumes in the gateway.</p> <p>Use this metric with the <i>Average</i> statistic to measure latency.</p> <p>Units: Milliseconds</p>
TotalCacheSize	<p>The total size of the cache in bytes. This metric applies only to the gateway-cached volume setup. The sample is taken at the end of the reporting period.</p> <p>Units: Bytes</p>

Metric	Description
WriteBytes	<p>The total number of bytes written to your on-premises applications in the reporting period for all volumes in the gateway.</p> <p>Use this metric with the <code>Sum</code> statistic to measure throughput and with the <code>Samples</code> statistic to measure operations per second (IOPS).</p> <p>Units: Bytes</p>
WriteTime	<p>The total number of milliseconds spent to do writes from your on-premises applications in the reporting period for all volumes in the gateway.</p> <p>Use this metric with the <code>Average</code> statistic to measure latency.</p> <p>Units: Milliseconds</p>
WorkingStorageFree	<p>The total amount of unused space in the gateway's working storage. The sample is taken at the end of the reporting period.</p> <p>Note Working storage applies only to the gateway-stored volume setup. The upload buffer applies to both the gateway-stored and gateway-cached volume setups. If you are working with both types of gateway setups, you may find it more convenient to use just the corresponding upload buffer metric, <code>UploadBufferFree</code>.</p> <p>Units: Bytes</p>
WorkingStoragePercentageUsed	<p>Percent utilization of the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Note Working storage applies only to the gateway-stored volume setup. The upload buffer applies to both the gateway-stored and gateway-cached volume setups. If you are working with both types of gateway setups, you may find it more convenient to use just the corresponding upload buffer metric, <code>UploadBufferPercentUsed</code>.</p> <p>Units: Percent</p>
WorkingStorageUsed	<p>The total number of bytes being used in the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Note Working storage applies only to the gateway-stored volume setup. The upload buffer applies to both the gateway-stored and gateway-cached volume setups. If you are working with both types of gateway setups, you may find it more convenient to use just the corresponding upload buffer metric, <code>UploadBufferUsed</code>.</p> <p>Units: Bytes</p>

Storage Volume Metrics

In this section, we discuss the AWS Storage Gateway metrics that give you information about a storage volume of a gateway. Each volume of a gateway has a set of metrics associated with it. Note that some volume-specific metrics have the same name as a gateway-specific metric. These metrics represent the

same kinds of measurements, but are scoped to the volume instead of the gateway. You must always specify whether you want to work with either a gateway or a storage volume metric before working with a metric. Specifically, when working with volume metrics, you must specify the *VolumeId* of the storage volume for which you are interested in viewing metrics. For more information, see [Using the Amazon CloudWatch Console \(p. 261\)](#).

The following table describes the AWS Storage Gateway metrics that you can use to get information about your storage volumes.

Metric	Description
CacheHitPercent	<p>Percent of application reads from the volume that are served from cache. This metric applies only to cached volumes. The sample is taken at the end of the reporting period.</p> <p>When there is no application reads from the volume, this metric reports 100%.</p> <p>Units: Percent</p>
CachePercentageUsed	<p>The volume's contribution to the overall percent utilization of the gateway's cache storage. This metric applies only to cached volumes. The sample is taken at the end of the reporting period.</p> <p>Use the <code>CachePercentageUsed</code> metric of the gateway to view overall percent utilization of the gateway's cache storage. For more information, see Gateway Metrics (p. 272).</p> <p>Units: Percent</p>
CachePercentDirty	<p>The volume's contribution to the overall percentage of the gateway's cache that has not been persisted to AWS. This metric applies only to the cached-volumes. The sample is taken at the end of the reporting period.</p> <p>Use the <code>CachePercentDirty</code> metric of the gateway to view the overall percentage of the gateway's cache that has not been persisted to AWS. For more information, see Gateway Metrics (p. 272).</p> <p>Units: Percent</p>
ReadBytes	<p>The total number of bytes read from your on-premises applications in the reporting period.</p> <p>Use this metric with the <code>Sum</code> statistic to measure throughput and with the <code>Samples</code> statistic to measure operations per second (IOPS).</p> <p>Units: Bytes</p>
ReadTime	<p>The total number of milliseconds spent to do reads from your on-premises applications in the reporting period.</p> <p>Use this metric with the <code>Average</code> statistic to measure latency.</p> <p>Units: Milliseconds</p>
WriteBytes	<p>The total number of bytes written to your on-premises applications in the reporting period.</p> <p>Use this metric with the <code>Sum</code> statistic to measure throughput and with the <code>Samples</code> statistic to measure operations per second (IOPS).</p> <p>Units: Bytes</p>

Metric	Description
WriteTime	The total number of milliseconds spent to do writes from your on-premises applications in the reporting period. Use this metric with the <i>Average</i> statistic to measure latency. Units: Milliseconds
QueuedWrites	The number of bytes waiting to be written to AWS, sampled at the end of the reporting period. Units: Bytes

Related Section

- [API Reference for AWS Storage Gateway \(p. 283\)](#)

Access Control Using AWS Identity and Access Management (IAM)

AWS Identity and Access Management (IAM) helps you securely control access to Amazon Web Services and your account resources. With IAM, you can create multiple IAM users under the umbrella of your AWS account. To learn more about IAM and its features, go to [What Is IAM?](#)

Every user you create in the IAM system starts with no permissions. In other words, by default, users can do nothing. A *permission* is a general term we use to mean the ability to perform an action against a resource. The AWS Storage Gateway API (see [API Reference for AWS Storage Gateway \(p. 283\)](#)) enables a list of actions you can perform. However, unless you explicitly grant a user permissions, that user cannot perform any of these actions. You grant a permission to a user with a policy. A policy is a document that formally states one or more permissions. For more information about IAM policies, go to [Overview of Policies](#).

You write a policy using the access policy language that IAM uses. You then attach the policy to a user or a group in your AWS account. For more information about the policy language, go to [The Access Policy Language](#) in *Using AWS Identity and Access Management*.

The [Element Descriptions](#) section of *Using AWS Identity and Access Management* describes elements you can use in a policy. The following information about some of the policy elements is specific to AWS Storage Gateway:

- **Resource**—The object or objects the policy covers. You identify resources using the following Amazon Resource Name (ARN) format.

```
arn:aws:<vendor>:<region>:<namespace>:<relative-id>
```

In this format, *vendor* is the product name "storagegateway" and *namespace* is the account ID. In AWS Storage Gateway, there are three types of resources, *gateway*, *volume*, and *iSCSITarget*. For each type of resource, the following table shows example ARNs.

Resource	Description
Gateway ARN	arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway

Resource	Description
Volume ARN	arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB
Target ARN (name of an iSCSI target)	arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/qn1997-05.com:amazon:myvolume

- **Action**—The specific type or types of action allowed or denied. For a complete list of AWS Storage Gateway actions, see [Operations in AWS Storage Gateway \(p. 305\)](#).

Note

The Amazon EBS snapshots generated from AWS Storage Gateway are managed as Amazon EC2 resources and corresponding EC2 actions.

This section provides example IAM policies that illustrate how to grant a user permission to perform specific AWS Storage Gateway actions. You can then attach these policies to a user for whom you want to grant access permissions.

Example Policies

Example 1: Allow all actions

The following policy allows a user to perform all the AWS Storage Gateway actions. The policy also allows the user to perform Amazon EC2 actions ([DescribeSnapshots](#) and [DeleteSnapshot](#)) on the Amazon EBS snapshots generated from AWS Storage Gateway.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "*"
    },
    {
      "Action": [
        "ec2:DescribeSnapshots",
        "ec2:DeleteSnapshot"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

Example 2: Allow read-only access to a gateway

The following policy allows all `List*` and `Describe*` actions on all resources. Note that these actions are read actions. So the policy does not allow the user to change state of any resources—that is, the policy does not allow the user to perform the actions such as `DeleteGateway`, `ActivateGateway`, and `ShutdownGateway`.

The policy also allows the `DescribeSnapshots` Amazon EC2 action. For more information, go to [DescribeSnapshots](#) in the *Amazon Elastic Compute Cloud API Reference*.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:List*",
        "storagegateway:Describe*"
      ],
      "Effect": "Allow",
      "Resource": "*"
    },
    {
      "Action": [
        "ec2:DescribeSnapshots"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

In the preceding policy, instead of using a wild card, you could scope resources covered by the policy to a specific gateway. The policy would then allow the actions only on the specific gateway.

```
"Resource": "arn:aws:storagegateway:us-east-1:111122223333:gateway/[Gateway Name]/*"
```

Within a gateway, you can further restrict the scope of the resources to only the gateway volumes.

```
"Resource": "arn:aws:storagegateway:us-east-1:111122223333:gateway/[Gateway Name]/volume/*"
```

Example 3: Allow access to a specific gateway

The following policy allows all actions on a specific gateway. That is, the user is restricted from accessing other gateways you might have deployed.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:storagegateway:[AWS Region]:[AWS Account]:gateway/[Gateway Name]/*"
    }
  ]
}
```

The preceding policy works if the user to whom the policy is attached uses either the API or an AWS SDK to access the gateway. However, if this user plans to use the AWS Storage Gateway console, you must also grant permission to the `ListGateways` action.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:storagegateway:[AWS Region]:[AWS Account]:gateway/[Gateway Name]/*"
    },
    {
      "Action": [
        "storagegateway:ListGateways"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

Additionally, if the user plans to activate the specific gateway, you must also grant permission to the `ActivateGateway` action.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:storagegateway:[AWS Region]:[AWS Account]:gateway/[Gateway Name]/*"
    },
    {

```

```
    "Action": [
      "storagegateway:ListGateways",
      "storagegateway:ActivateGateway"
    ],
    "Effect": "Allow",
    "Resource": "*"
  }
]
```

Example 4: Grant permissions to access a specific volume

The following policy allows a user all actions to a specific volume on a gateway. Because a user does not get any permissions by default, the policy restricts the user to accessing only a specific volume.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:storagegateway:[AWS Region]:[AWS Account]:gateway/[Gateway Name]/volume/[Volume Name]"
    }
  ]
}
```

The preceding policy works if the user to whom the policy is attached uses either the API or an AWS SDK to access the volume. However, if this user plans to use the AWS Storage Gateway console, you must also grant permission to the `ListGateways` action.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:storagegateway:[AWS Region]:[AWS Account]:gateway/[Gateway Name]/volume/[Volume Name]"
    },
    {
      "Action": [
        "storagegateway:ListGateways"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

Example 5: Allow all actions on gateways with a specific prefix

The following policy allows a user to perform all action on gateways whose name starts with "DeptX". The policy also allows the `DescribeSnapshots` Amazon EC2 action.

```
{
  "Statement": [
    {
      "Action": [
        "storagegateway:*"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:storagegateway:[AWS Region]:[AWS Account]:gateway/[Gateway Name Prefix]*"
    },
    {
      "Action": [
        "ec2:DescribeSnapshots"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

The preceding policy works if the user to whom the policy is attached uses either the API or an AWS SDK to access the gateway. However, if this user plans to use the AWS Storage Gateway console, you must grant additional permissions as described in [3: Allow access to a specific gateway \(p. 280\)](#).

API Reference for AWS Storage Gateway

Topics

- [AWS Storage Gateway Required Request Headers \(p. 283\)](#)
- [Signing Requests \(p. 285\)](#)
- [Error Responses \(p. 287\)](#)
- [Operations in AWS Storage Gateway \(p. 305\)](#)

In addition to using the console, you can use the AWS Storage Gateway API to programmatically configure and manage your gateways. This section describes the AWS Storage Gateway operations, request signing for authentication and the error handling. For information about the regions and endpoints available for AWS Storage Gateway, see [Regions and Endpoints](#).

Note

You can also use the AWS SDKs when developing applications with AWS Storage Gateway. The AWS SDKs for Java, .NET and PHP wrap the underlying AWS Storage Gateway API, simplifying your programming tasks. For information about downloading the SDK libraries, go to [Sample Code Libraries](#).

AWS Storage Gateway Required Request Headers

This section describes the required headers that you must send with every POST request to AWS Storage Gateway. You include HTTP headers to identify key information about the request including the operation you want to invoke, the date of the request, and information that indicates the authorization of you as the sender of the request. Headers are case insensitive and the order of the headers is not important.

The following example shows headers that are used in the [ActivateGateway \(p. 307\)](#) operation.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
```

AWS Storage Gateway User Guide

Required Request Headers

```
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.ActivateGateway
```

The following are the headers that must include with your POST requests to AWS Storage Gateway. Headers shown below that begin with "x-amz" are AWS-specific headers. All other headers listed are common header used in HTTP transactions.

Header	Description
Authorization	<p>The authorization header contains several of pieces of information about the request that enable AWS Storage Gateway to determine if the request is a valid action for the requester. The format of this header is as follows (line breaks added for readability):</p> <pre>Authorization: AWS4-HMAC_SHA456 Credentials=<i>YourAccessKey</i>/<i>yyyymmdd</i>/<i>region</i>/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=<i>CalculatedSignature</i></pre> <p>In the preceding syntax, you specify <i>YourAccessKey</i>, the year, month, and day (<i>yyyymmdd</i>), the <i>region</i>, and the <i>CalculatedSignature</i>. The format of the authorization header is dictated by the requirements of the AWS V4 Signing process. The details of signing are discussed in the topic Signing Requests (p. 285).</p>
Content-Type	<p>Use <code>application/x-amz-json-1.1</code> as the content type for all requests to AWS Storage Gateway.</p> <pre>Content-Type: application/x-amz-json-1.1</pre>
Host	<p>Use the host header to specify the AWS Storage Gateway endpoint where you send your request. For example, <code>storagegateway.us-east-1.amazonaws.com</code> is the endpoint for the US East Region. For more information about the endpoints available for AWS Storage Gateway, see Regions and Endpoints.</p> <pre>Host: storagegateway.<i>region</i>.amazonaws.com</pre>
x-amz-date	<p>You must provide the time stamp in either the HTTP <code>Date</code> header or the AWS <code>x-amz-date</code> header. (Some HTTP client libraries don't let you set the <code>Date</code> header.) When an <code>x-amz-date</code> header is present, the AWS Storage Gateway ignores any <code>Date</code> header during the request authentication. The <code>x-amz-date</code> format must be ISO8601 Basic in the <code>YYYYMMDD'T'HHMMSS'Z'</code> format. If both the <code>Date</code> and <code>x-amz-date</code> header are used, the format of the <code>Date</code> header does not have to be ISO8601.</p> <pre>x-amz-date: <i>YYYYMMDD'T'HHMMSS'Z'</i></pre>

Header	Description
x-amz-target	This header specifies the version of the API and the operation that you are requesting. The target header values are formed by concatenating the API version with the API name and are in the following format. <code>x-amz-target: StorageGateway_<i>APIVersion</i>.<i>operationName</i></code> The <i>operationName</i> value (e.g. "ActivateGateway") can be found from the API list, API Reference for AWS Storage Gateway (p. 283) .

Signing Requests

AWS Storage Gateway requires that you authenticate every request you send by signing the request. To sign a request, you calculate a digital signature using a cryptographic hash function. A cryptographic hash is a function that returns a unique hash value based on the input. The input to the hash function includes the text of your request and your secret access key. The hash function returns a hash value that you include in the request as your signature. The signature is part of the `Authorization` header of your request.

After receiving your request, AWS Storage Gateway recalculates the signature using the same hash function and input that you used to sign the request. If the resulting signature matches the signature in the request, AWS Storage Gateway processes the request. Otherwise, the request is rejected.

AWS Storage Gateway supports authentication using [AWS Signature Version 4](#). The process for calculating a signature can be broken into three tasks:

- [Task 1: Create a Canonical Request](#)

Rearrange your HTTP request into a canonical format. Using a canonical form is necessary because AWS Storage Gateway uses the same canonical form when it recalculates a signature to compare with the one you sent.

- [Task 2: Create a String to Sign](#)

Create a string that you will use as one of the input values to your cryptographic hash function. The string, called the *string to sign*, is a concatenation of the name of the hash algorithm, the request date, a *credential scope* string, and the canonicalized request from the previous task. The *credential scope* string itself is a concatenation of date, region, and service information.

- [Task 3: Create a Signature](#)

Create a signature for your request by using a cryptographic hash function that accepts two input strings: your *string to sign* and a *derived key*. The *derived key* is calculated by starting with your secret access key and using the *credential scope* string to create a series of Hash-based Message Authentication Codes (HMACs).

Example Signature Calculation

The following example walks you through the details of creating a signature for [ListGateways \(p. 372\)](#). The example could be used as a reference to check your signature calculation method. Other reference calculations are included in the [Signature Version 4 Test Suite](#) of the Amazon Web Services Glossary.

The example assumes the following:

- The time stamp of the request is "Mon, 10 Sep 2012 00:00:00" GMT.

- The endpoint is the US East (Northern Virginia) Region.

The general request syntax (including the JSON body) is:

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
x-amz-Date: 20120910T000000Z
Authorization: SignatureToBeCalculated
Content-type: application/x-amz-json-1.1
x-amz-target: StorageGateway_20120630.ListGateways
{ }
```

The canonical form of the request calculated for [Task 1: Create a Canonical Request \(p. 285\)](#) is:

```
POST
/

content-type:application/x-amz-json-1.1
host:storagegateway.us-east-1.amazonaws.com
x-amz-date:20120910T000000Z
x-amz-target:StorageGateway_20120630.ListGateways

content-type;host;x-amz-date;x-amz-target
44136fa355b3678a1146ad16f7e8649e94fb4fc21fe77e8310c060f61caaff8a
```

The last line of the canonical request is the hash of the request body. Also, note the empty third line in the canonical request. This is because there are no query parameters for this API (or any AWS Storage Gateway APIs).

The *string to sign* for [Task 2: Create a String to Sign \(p. 285\)](#) is:

```
AWS4-HMAC-SHA256
20120910T000000Z
20120910/us-east-1/storagegateway/aws4_request
92c0effa6f9224ac752ca179a04cecbde3038b0959666a8160ab452c9e51b3e
```

The first line of the *string to sign* is the algorithm, the second line is the time stamp, the third line is the *credential scope*, and the last line is a hash of the canonical request from Task 1.

For [Task 3: Create a Signature \(p. 285\)](#), the *derived key* can be represented as:

```
derived key = HMAC(HMAC(HMAC(HMAC("AWS4" + YourSecretAccessKey, "20120910"), "us-east-1"), "storagegateway"), "aws4_request")
```

If the secret access key, wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY, is used, then the calculated signature is:

```
6d4c40b8f2257534dbdca9f326f147a0a7a419b63aff349d9d9c737c9a0f4c81
```

The final step is to construct the *Authorization* header. For the demonstration access key AKIAIOSFODNN7EXAMPLE, the header (with line breaks added for readability) is:

```
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120910/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=6d4c40b8f2257534dbdca9f326f147a0a7a419b63aff349d9d9c737c9a0f4c81
```

Error Responses

Topics

- [Exceptions \(p. 287\)](#)
- [Operation Error Codes \(p. 288\)](#)
- [Error Responses \(p. 303\)](#)

This section provides reference information about AWS Storage Gateway errors. These errors are represented by an error exception and an operation error code. For example, the error exception `InvalidSignatureException` is returned by any API response if there is a problem with the request signature. However, the operation error code `ActivationKeyInvalid` is returned only for the [ActivateGateway \(p. 307\)](#) API.

Depending on the type of error, AWS Storage Gateway may return only just an exception, or it may return both an exception and an operation error code. Examples of error responses are shown in the [Error Responses \(p. 303\)](#).

Exceptions

The following table lists AWS Storage Gateway API exceptions. When an AWS Storage Gateway operation returns an error response, the response body contains one of these exceptions. The `InternalServerError` and `InvalidGatewayRequestException` return one of the [Operation Error Codes \(p. 288\)](#) message codes that give the specific operation error code.

Exception	Message	HTTP Status Code
<code>IncompleteSignatureException</code>	The specified signature is incomplete.	400 Bad Request
<code>InternalFailure</code>	The request processing has failed due to some unknown error, exception or failure.	500 Internal Server Error
<code>InternalServerError</code>	One of the operation error code messages in Operation Error Codes (p. 288) .	500 Internal Server Error
<code>InvalidAction</code>	The requested action or operation is invalid.	400 Bad Request
<code>InvalidClientTokenId</code>	The X.509 certificate or AWS Access Key ID provided does not exist in our records.	403 Forbidden
<code>InvalidGatewayRequestException</code>	One of the operation error code messages in Operation Error Codes (p. 288) .	400 Bad Request

Exception	Message	HTTP Status Code
InvalidSignatureException	The request signature we calculated does not match the signature you provided. Check your AWS Access Key and signing method.	400 Bad Request
MissingAction	The request is missing an action or operation parameter.	400 Bad Request
MissingAuthenticationToken	The request must contain either a valid (registered) AWS Access Key ID or X.509 certificate.	403 Forbidden
RequestExpired	The request is past the expiration date or the request date (either with 15 minute padding), or the request date occurs more than 15 minutes in the future.	400 Bad Request
SerializationException	An error occurred during serialization. Check that your JSON payload is well-formed.	400 Bad Request
ServiceUnavailable	The request has failed due to a temporary failure of the server.	503 Service Unavailable
SubscriptionRequiredException	The AWS Access Key Id needs a subscription for the service.	400 Bad Request
ThrottlingException	Rate exceeded.	400 Bad Request
UnknownOperationException	An unknown operation was specified. Valid operations are listed in Operations in AWS Storage Gateway (p. 305) .	400 Bad Request
UnrecognizedClientException	The security token included in the request is invalid.	400 Bad Request
ValidationException	The value of an input parameter is bad or out of range.	400 Bad Request

Operation Error Codes

The following table shows the mapping between AWS Storage Gateway operation error codes and APIs that can return the codes. All operation error codes are returned with one of two general exceptions - `InternalServerError` and `InvalidGatewayRequestException` exception - described in [Exceptions \(p. 287\)](#).

Operation Error Code	Message	Operations That Return this Error Code
ActivationKeyExpired	The specified activation key has expired.	ActivateGateway (p. 307)
ActivationKeyInvalid	The specified activation key is invalid.	ActivateGateway (p. 307)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
ActivationKeyNotFound	The specified activation key was not found.	ActivateGateway (p. 307)
BandwidthThrottleScheduleNotFound	The specified bandwidth throttle was not found.	DeleteBandwidthRateLimit (p. 331)
CannotExportSnapshot	The specified snapshot cannot be exported.	CreateCachediSCSIVolume (p. 318) CreateStorediSCSIVolume (p. 327)
InitiatorNotFound	The specified initiator was not found.	DeleteChapCredentials (p. 333)
DiskAlreadyAllocated	The specified disk is already allocated.	AddCache (p. 310) AddUploadBuffer (p. 313) AddWorkingStorage (p. 315) CreateStorediSCSIVolume (p. 327)
DiskDoesNotExist	The specified disk does not exist.	AddCache (p. 310) AddUploadBuffer (p. 313) AddWorkingStorage (p. 315) CreateStorediSCSIVolume (p. 327)
DiskSizeNotGigAligned	The specified disk is not gigabyte-aligned.	CreateStorediSCSIVolume (p. 327)
DiskSizeGreaterThanVolumeMaxSize	The specified disk size is greater than the maximum volume size.	CreateStorediSCSIVolume (p. 327)
DiskSizeLessThanVolumeSize	The specified disk size is less than the volume size.	CreateStorediSCSIVolume (p. 327)
DuplicateCertificateInfo	The specified certificate information is a duplicate.	ActivateGateway (p. 307)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
GatewayInternalError	A gateway internal error occurred.	AddCache (p. 310) AddUploadBuffer (p. 313) AddWorkingStorage (p. 315) CreateCachediSCSIVolume (p.318) CreateSnapshot (p. 321) CreateStorediSCSIVolume (p.327) CreateSnapshotFromVolumeRecoveryPoint (p.324) DeleteBandwidthRateLimit (p. 331) DeleteChapCredentials (p. 333) DeleteVolume (p. 340) DescribeBandwidthRateLimit (p.343) DescribeCache (p. 345) DescribeCachediSCSIVolumes (p.348) DescribeChapCredentials (p. 352) DescribeGatewayInformation (p.354) DescribeMaintenanceStartTime (p.358) DescribeSnapshotSchedule (p.360) DescribeStorediSCSIVolumes (p.363) DescribeWorkingStorage (p. 370) ListLocalDisks (p. 375) ListVolumes (p. 381) ListVolumeRecoveryPoints (p.378) ShutdownGateway (p. 384) StartGateway (p. 387) UpdateBandwidthRateLimit (p.389) UpdateChapCredentials (p. 391) UpdateMaintenanceStartTime (p.399) UpdateGatewaySoftwareNow (p.396) UpdateSnapshotSchedule (p. 401)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
GatewayNotConnected	The specified gateway is not connected.	AddCache (p. 310) AddUploadBuffer (p. 313) AddWorkingStorage (p. 315) CreateCachediSCSIVolume (p.318) CreateSnapshot (p. 321) CreateStorediSCSIVolume (p.327) CreateSnapshotFromVolumeRecoveryPoint (p.324) DeleteBandwidthRateLimit (p. 331) DeleteChapCredentials (p. 333) DeleteVolume (p. 340) DescribeBandwidthRateLimit (p.343) DescribeCache (p. 345) DescribeCachediSCSIVolumes (p.348) DescribeChapCredentials (p. 352) DescribeGatewayInformation (p.354) DescribeMaintenanceStartTime (p.358) DescribeSnapshotSchedule (p.360) DescribeStorediSCSIVolumes (p.363) DescribeWorkingStorage (p. 370) ListLocalDisks (p. 375) ListVolumes (p. 381) ListVolumeRecoveryPoints (p.378) ShutdownGateway (p. 384) StartGateway (p. 387) UpdateBandwidthRateLimit (p.389) UpdateChapCredentials (p. 391) UpdateMaintenanceStartTime (p.399) UpdateGatewaySoftwareNow (p.396) UpdateSnapshotSchedule (p. 401)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
GatewayNotFound	The specified gateway was not found.	AddCache (p. 310) AddUploadBuffer (p. 313) AddWorkingStorage (p. 315) CreateCachediSCSIVolume (p.318) CreateSnapshot (p. 321) CreateSnapshotFromVolumeRecoveryPoint (p.324) CreateStorediSCSIVolume (p. 327) DeleteBandwidthRateLimit (p. 331) DeleteChapCredentials (p. 333) DeleteGateway (p. 336) DeleteVolume (p. 340) DescribeBandwidthRateLimit (p.343) DescribeCache (p. 345) DescribeCachediSCSIVolumes (p.348) DescribeChapCredentials (p. 352) DescribeGatewayInformation (p.354) DescribeMaintenanceStartTime (p.358) DescribeSnapshotSchedule (p. 360) DescribeStorediSCSIVolumes (p.363) DescribeWorkingStorage (p. 370) ListLocalDisks (p. 375) ListVolumes (p. 381) ListVolumeRecoveryPoints (p. 378) ShutdownGateway (p. 384) StartGateway (p. 387) UpdateBandwidthRateLimit (p. 389) UpdateChapCredentials (p. 391) UpdateMaintenanceStartTime (p.399) UpdateGatewaySoftwareNow (p.396) UpdateSnapshotSchedule (p. 401)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
GatewayProxyNetworkConnectionBusy	The specified gateway proxy network connection is busy.	AddCache (p. 310) AddUploadBuffer (p. 313) AddWorkingStorage (p. 315) CreateCachediSCSIVolume (p.318) CreateSnapshot (p. 321) CreateSnapshotFromVolumeRecoveryPoint (p.324) CreateStorediSCSIVolume (p. 327) DeleteBandwidthRateLimit (p. 331) DeleteChapCredentials (p. 333) DeleteVolume (p. 340) DescribeBandwidthRateLimit (p.343) DescribeCache (p. 345) DescribeCachediSCSIVolumes (p.348) DescribeChapCredentials (p. 352) DescribeGatewayInformation (p.354) DescribeMaintenanceStartTime (p.358) DescribeSnapshotSchedule (p.360) DescribeStorediSCSIVolumes (p.363) DescribeWorkingStorage (p. 370) ListLocalDisks (p. 375) ListVolumes (p. 381) ListVolumeRecoveryPoints (p.378) ShutdownGateway (p. 384) StartGateway (p. 387) UpdateBandwidthRateLimit (p.389) UpdateChapCredentials (p. 391) UpdateMaintenanceStartTime (p.399) UpdateGatewaySoftwareNow (p.396) UpdateSnapshotSchedule (p. 401)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
InternalError	An internal error occurred.	

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
		ActivateGateway (p. 307) AddCache (p. 310) AddUploadBuffer (p. 313) AddWorkingStorage (p. 315) CreateCachediSCSIVolume (p.318) CreateSnapshot (p. 321) CreateSnapshotFromVolumeRecoveryPoint (p.324) CreateStorediSCSIVolume (p. 327) DeleteBandwidthRateLimit (p.331) DeleteChapCredentials (p. 333) DeleteGateway (p. 336) DeleteVolume (p. 340) DescribeBandwidthRateLimit (p.343) DescribeCache (p. 345) DescribeCachediSCSIVolumes (p.348) DescribeChapCredentials (p. 352) DescribeGatewayInformation (p.354) DescribeMaintenanceStartTime (p.358) DescribeSnapshotSchedule (p.360) DescribeStorediSCSIVolumes (p.363) DescribeWorkingStorage (p. 370) ListLocalDisks (p. 375) ListGateways (p. 372) ListVolumes (p. 381) ListVolumeRecoveryPoints (p.378) ShutdownGateway (p. 384) StartGateway (p. 387) UpdateBandwidthRateLimit (p.389) UpdateChapCredentials (p. 391) UpdateMaintenanceStartTime (p.399)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
		UpdateGatewayInformation (p. 394) UpdateGatewaySoftwareNow (p. 396) UpdateSnapshotSchedule (p. 401)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
InvalidParameters	The specified request contains invalid parameters.	

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
		ActivateGateway (p. 307) AddCache (p. 310) AddUploadBuffer (p. 313) AddWorkingStorage (p. 315) CreateCachediSCSIVolume (p.318) CreateSnapshot (p. 321) CreateSnapshotFromVolumeRecoveryPoint (p.324) CreateStorediSCSIVolume (p. 327) DeleteBandwidthRateLimit (p. 331) DeleteChapCredentials (p. 333) DeleteGateway (p. 336) DeleteVolume (p. 340) DescribeBandwidthRateLimit (p.343) DescribeCache (p. 345) DescribeCachediSCSIVolumes (p.348) DescribeChapCredentials (p. 352) DescribeGatewayInformation (p.354) DescribeMaintenanceStartTime (p.358) DescribeSnapshotSchedule (p. 360) DescribeStorediSCSIVolumes (p.363) DescribeWorkingStorage (p. 370) ListLocalDisks (p. 375) ListGateways (p. 372) ListVolumes (p. 381) ListVolumeRecoveryPoints (p. 378) ShutdownGateway (p. 384) StartGateway (p. 387) UpdateBandwidthRateLimit (p. 389) UpdateChapCredentials (p. 391) UpdateMaintenanceStartTime (p.399)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
		UpdateGatewayInformation (p. 394) UpdateGatewaySoftwareNow (p. 396) UpdateSnapshotSchedule (p. 401)
LocalStorageLimitExceeded	The local storage limit was exceeded.	AddCache (p. 310) AddUploadBuffer (p. 313) AddWorkingStorage (p. 315)
LunInvalid	The specified LUN is invalid.	CreateStorediSCSIVolume (p. 327)
MaximumVolumeCountExceeded	The maximum volume count was exceeded.	CreateCachediSCSIVolume (p. 318) CreateStorediSCSIVolume (p. 327) DescribeCachediSCSIVolumes (p. 348) DescribeStorediSCSIVolumes (p. 363)
NetworkConfigurationChanged	The gateway network configuration has changed.	CreateCachediSCSIVolume (p. 318) CreateStorediSCSIVolume (p. 327)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
NotSupported	The specified operation is not supported.	

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
		ActivateGateway (p. 307) AddCache (p. 310) AddUploadBuffer (p. 313) AddWorkingStorage (p. 315) CreateCachediSCSIVolume (p.318) CreateSnapshot (p. 321) CreateSnapshotFromVolumeRecoveryPoint (p.324) CreateStorediSCSIVolume (p. 327) DeleteBandwidthRateLimit (p.331) DeleteChapCredentials (p. 333) DeleteGateway (p. 336) DeleteVolume (p. 340) DescribeBandwidthRateLimit (p.343) DescribeCache (p. 345) DescribeCachediSCSIVolumes (p.348) DescribeChapCredentials (p. 352) DescribeGatewayInformation (p.354) DescribeMaintenanceStartTime (p.358) DescribeSnapshotSchedule (p.360) DescribeStorediSCSIVolumes (p.363) DescribeWorkingStorage (p. 370) ListLocalDisks (p. 375) ListGateways (p. 372) ListVolumes (p. 381) ListVolumeRecoveryPoints (p.378) ShutdownGateway (p. 384) StartGateway (p. 387) UpdateBandwidthRateLimit (p.389) UpdateChapCredentials (p. 391) UpdateMaintenanceStartTime (p.399)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
		UpdateGatewayInformation (p. 394) UpdateGatewaySoftwareNow (p. 396) UpdateSnapshotSchedule (p. 401)
OutdatedGateway	The specified gateway is out of date.	ActivateGateway (p. 307)
SnapshotInProgressException	The specified snapshot is in progress.	DeleteVolume (p. 340)
SnapshotIdInvalid	The specified snapshot is invalid.	CreateCachediSCSIVolume (p. 318) CreateStorediSCSIVolume (p. 327)
StagingAreaFull	The staging area is full.	CreateCachediSCSIVolume (p. 318) CreateStorediSCSIVolume (p. 327)
TargetAlreadyExists	The specified target already exists.	CreateCachediSCSIVolume (p. 318) CreateStorediSCSIVolume (p. 327)
TargetInvalid	The specified target is invalid.	CreateCachediSCSIVolume (p. 318) CreateStorediSCSIVolume (p. 327) DeleteChapCredentials (p. 333) DescribeChapCredentials (p. 352) UpdateChapCredentials (p. 391)
TargetNotFound	The specified target was not found.	CreateCachediSCSIVolume (p. 318) CreateStorediSCSIVolume (p. 327) DeleteChapCredentials (p. 333) DescribeChapCredentials (p. 352) DeleteVolume (p. 340) UpdateChapCredentials (p. 391)

Operation Error Code	Message	Operations That Return this Error Code
UnsupportedOperationForGatewayType	The specified operation is not valid for the type of the gateway.	AddCache (p. 310) AddWorkingStorage (p. 315) CreateCachediSCSIVolume (p.318) CreateSnapshotFromVolumeRecoveryPoint (p.324) CreateStorediSCSIVolume (p. 327) DeleteSnapshotSchedule (p. 338) DescribeCache (p. 345) DescribeCachediSCSIVolumes (p.348) DescribeStorediSCSIVolumes (p.363) DescribeUploadBuffer (p. 367) DescribeWorkingStorage (p. 370) ListVolumeRecoveryPoints (p. 378)
VolumeAlreadyExists	The specified volume already exists.	CreateCachediSCSIVolume (p. 318) CreateStorediSCSIVolume (p. 327)
VolumeIdInvalid	The specified volume is invalid.	DeleteVolume (p. 340)
VolumeInUse	The specified volume is already in use.	DeleteVolume (p. 340)
VolumeNotFound	The specified volume was not found.	CreateSnapshot (p. 321) CreateSnapshotFromVolumeRecoveryPoint (p.324) DeleteVolume (p. 340) DescribeCachediSCSIVolumes (p.348) DescribeSnapshotSchedule (p.360) DescribeStorediSCSIVolumes (p.363) UpdateSnapshotSchedule (p. 401)
VolumeNotReady	The specified volume is not ready.	CreateSnapshot (p. 321) CreateSnapshotFromVolumeRecoveryPoint (p.324)

Error Responses

When there is an error, the response header information contains:

- Content-Type: application/x-amz-json-1.1
- An appropriate 4xx or 5xx HTTP status code

The body of an error response contains information about the error that occurred. The following sample error response shows the output syntax of response elements common to all error responses.

```
{
  "__type": "String",
  "message": "String",
  "error":
    { "errorCode": "String",
      "errorDetails": "String"
    }
}
```

The following table explains the JSON error response fields shown in the preceding syntax.

__type

One of the exceptions from [Exceptions \(p. 287\)](#).

Type: String

error

Contains API-specific error details. In general errors (i.e., not specific to any API), this error information is not shown.

Type: Collection

errorCode

One of the operation error codes from [Operation Error Codes \(p. 288\)](#).

Type: String

errorDetails

This field is not used in the current version of the API.

Type: String

message

One of the operation error code messages from [Operation Error Codes \(p. 288\)](#).

Type: String

Error Response Examples

The following JSON body is returned if you use the [DescribeStorediSCSIVolumes \(p. 363\)](#) API and specify a gateway ARN request input that does not exist.

```
{
  "__type": "InvalidGatewayRequestException",
  "message": "The specified volume was not found.",
  "error": {
    "errorCode": "VolumeNotFound"
  }
}
```

The following JSON body is returned if AWS Storage Gateway calculates a signature that does not match the signature sent with a request.

```
{
  "__type": "InvalidSignatureException",
```

```
"message": "The request signature we calculated does not match the signature  
you provided."  
}
```

Operations in AWS Storage Gateway

Topics

- [ActivateGateway](#) (p. 307)
- [AddCache](#) (p. 310)
- [AddUploadBuffer](#) (p. 313)
- [AddWorkingStorage](#) (p. 315)
- [CreateCachediSCSIVolume](#) (p. 318)
- [CreateSnapshot](#) (p. 321)
- [CreateSnapshotFromVolumeRecoveryPoint](#) (p. 324)
- [CreateStorediSCSIVolume](#) (p. 327)
- [DeleteBandwidthRateLimit](#) (p. 331)
- [DeleteChapCredentials](#) (p. 333)
- [DeleteGateway](#) (p. 336)
- [DeleteSnapshotSchedule](#) (p. 338)
- [DeleteVolume](#) (p. 340)
- [DescribeBandwidthRateLimit](#) (p. 343)
- [DescribeCache](#) (p. 345)
- [DescribeCachediSCSIVolumes](#) (p. 348)
- [DescribeChapCredentials](#) (p. 352)
- [DescribeGatewayInformation](#) (p. 354)
- [DescribeMaintenanceStartTime](#) (p. 358)
- [DescribeSnapshotSchedule](#) (p. 360)
- [DescribeStorediSCSIVolumes](#) (p. 363)
- [DescribeUploadBuffer](#) (p. 367)
- [DescribeWorkingStorage](#) (p. 370)
- [ListGateways](#) (p. 372)
- [ListLocalDisks](#) (p. 375)
- [ListVolumeRecoveryPoints](#) (p. 378)
- [ListVolumes](#) (p. 381)
- [ShutdownGateway](#) (p. 384)
- [StartGateway](#) (p. 387)
- [UpdateBandwidthRateLimit](#) (p. 389)
- [UpdateChapCredentials](#) (p. 391)
- [UpdateGatewayInformation](#) (p. 394)
- [UpdateGatewaySoftwareNow](#) (p. 396)
- [UpdateMaintenanceStartTime](#) (p. 399)
- [UpdateSnapshotSchedule](#) (p. 401)
- [Data Types](#) (p. 404)
- [Enumeration Types](#) (p. 411)

AWS Storage Gateway User Guide

Operations in AWS Storage Gateway

This section contains detailed descriptions of all AWS Storage Gateway operations, their request parameters, response elements, possible errors, and examples of requests and responses.

AWS Storage Gateway uses JSON to send and receive data. Returned JSON from AWS Storage Gateway APIs is subject to future expansion. You should build your client software to be forward compatible with AWS Storage Gateway by ignoring unknown JSON fields.

The following table summarizes the operations available in AWS Storage Gateway according to which [gateway architecture](#) (p. 3) (gateway-cached or gateway-stored volumes) to which they apply.

Operation	Used for Gateway-Cached Setup?	Used for Gateway-Stored Setup?
ActivateGateway (p. 307)	Yes	Yes
AddCache (p. 310)	Yes	No
AddUploadBuffer (p. 313)	Yes	Yes. For stored gateways, this operation is equivalent to using AddWorkingStorage (p. 315).
AddWorkingStorage (p. 315)	No	Yes
CreateCachediSCSIVolume (p.318)	Yes	No
CreateSnapshot (p. 321)	Yes	Yes
CreateSnapshotFromVolumeRecoveryPoint (p.324)	Yes	No
CreateStorediSCSIVolume (p.327)	No	Yes
DeleteBandwidthRateLimit (p.331)	Yes	Yes
DeleteChapCredentials (p. 333)	Yes	Yes
DeleteGateway (p. 336)	Yes	Yes
DeleteSnapshotSchedule (p. 338)	Yes	Yes
DeleteVolume (p. 340)	Yes	Yes
DescribeBandwidthRateLimit (p.343)	Yes	Yes
DescribeCache (p. 345)	Yes	No
DescribeCachediSCSIVolumes (p.348)	Yes	No
DescribeChapCredentials (p. 352)	Yes	Yes
DescribeGatewayInformation (p.354)	Yes	Yes
DescribeMaintenanceStartTime (p.358)	Yes	Yes
DescribeSnapshotSchedule (p.360)	Yes	Yes
DescribeStorediSCSIVolumes (p.363)	Yes	Yes
DescribeUploadBuffer (p. 367)	Yes	Yes. For stored gateways, this operation is equivalent to using DescribeWorkingStorage (p. 370).
DescribeWorkingStorage (p. 370)	No	Yes

Operation	Used for Gateway-Cached Setup?	Used for Gateway-Stored Setup?
ListGateways (p. 372)	Yes	Yes
ListLocalDisks (p. 375)	Yes	Yes
ListVolumeRecoveryPoints (p.378)	Yes	No
ListVolumes (p. 381)	Yes	Yes
ShutdownGateway (p. 384)	Yes	Yes
StartGateway (p. 387)	Yes	Yes
UpdateBandwidthRateLimit (p.389)	Yes	Yes
UpdateChapCredentials (p. 391)	Yes	Yes
UpdateGatewayInformation (p.394)	Yes	Yes
UpdateGatewaySoftwareNow (p.396)	Yes	Yes
UpdateMaintenanceStartTime (p.399)	Yes	Yes
UpdateSnapshotSchedule (p. 401)	Yes	Yes

ActivateGateway

Description

This operation activates the gateway you previously deployed on your VMware host. The activation process associates your gateway with your account. For more information, see [Downloading and Deploying AWS Storage Gateway VM](#) (p. 91). In the activation process, you specify information such as the type of gateway, the region you want to use for storing snapshots, the time zone for scheduled snapshots and the gateway schedule window, an activation key, and a name for your gateway. You can change the gateway's name and time zone after activation (see [UpdateGatewayInformation](#) (p. 394)).

Note

You must power on the gateway VM before you can activate your gateway.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.ActivateGateway

{
  "ActivationKey": "String",
  "GatewayName": "String",
  "GatewayTimezone": "String",
  "GatewayRegion": "String",
```

```
"GatewayType" : "String"  
}
```

JSON Fields

ActivationKey

Your gateway activation key. You can obtain the activation key by sending an HTTP GET request with redirects disabled to the gateway IP address (port 80).

The redirect URL returned in the response includes the activation key as part of the query string in the parameter `activationKey`. It may also include other activation-related parameters. However, these are merely defaults—the arguments you pass to the `ActivateGateway` API call determine the actual configuration of your gateway.

Required: Yes

Type: String

GatewayName

A unique identifier for your gateway. This name becomes part of the gateway ARN, which is what you use as an input to other operations. Gateway names are unique per AWS account, but not globally.

Length: Minimum length of 2. Maximum length of 255.

Required: Yes

Type: String. ASCII characters only, and the name cannot be all spaces and cannot contain a forward (/) or backward slash (\).

GatewayRegion

One of the [Regions \(p. 412\)](#) values that indicates the region where you want to store the snapshot backups. The gateway region specified must be the same region as the region in your `Host` header in the request.

Required: Yes

Type: String

GatewayTimezone

One of the [GatewayTimezone \(p. 411\)](#) values that indicates the time zone you want to set for the gateway. The time zone is used, for example, for scheduling snapshots and your gateway's maintenance schedule.

Required: Yes

Type: String

GatewayType

One of the [GatewayType \(p. 412\)](#) values that defines the type of gateway to activate. The type specified is critical to all later functions of the gateway and cannot be changed after activation. The default value is `STORED`.

Required: No

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

AWS Storage Gateway returns the ARN of the activated gateway. It is a string made of information such as your account, gateway name, and region. This ARN is used to reference the gateway in other API operations as well as resource-based authorization.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- ActivationKeyExpired
- ActivationKeyInvalid
- ActivationKeyNotFound
- DuplicateCertificateInfo
- InternalError
- InvalidParameters
- NotSupported
- OutdatedGateway

Examples

Example Request

The following example shows a request that activates a cached gateway.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
```

```
x-amz-target: StorageGateway_20120630.ActivateGateway

{
  "ActivationKey": "29AV1-30FV9-VVIUB-NKT0I-LR06V",
  "GatewayName": "mygateway",
  "GatewayTimezone": "GMT-12:00",
  "GatewayRegion": "us-east-1",
  "GatewayType": "CACHED"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glllelqeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 80

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [StartGateway](#) (p. 387)
- [ListGateways](#) (p. 372)
- [ShutdownGateway](#) (p. 384)
- [DescribeGatewayInformation](#) (p. 354)
- [DeleteGateway](#) (p. 336)

AddCache

Description

This operation configures one or more gateway local disks as cache storage for a specified gateway. This operation is supported only for the gateway-cached volume architecture (see [How AWS Storage Gateway Works](#) (p. 3)).

In the request, you specify the gateway Amazon Resource Name (ARN) to which you want to add cache storage, and one or more disk IDs that you want to configure as cache storage.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.AddCache
```

```
{
  "GatewayARN": "String",
  "DiskIds": [
    "String",
    ...
  ]
}
```

JSON Fields

DiskIds

An array of strings that identify disks that are to be configured as cache. Each string in the array must be minimum length of 1 and maximum length of 300. You can get the disk IDs from the [ListLocalDisks \(p. 375\)](#) API.

Required: Yes

Type: Array

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The ARN of the gateway for which cache storage was configured.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- DiskAlreadyAllocated

- DiskDoesNotExist
- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- LocalStorageLimitExceeded
- NotSupported
- UnsupportedOperationForGatewayType

Examples

Example Request

The following example shows a request that specifies that two local disks of a gateway are to be configured as cache.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.AddCache

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"

  "DiskIds": [
    "pci-0000:03:00.0-scsi-0:0:0:0",
    "pci-0000:03:00.0-scsi-0:0:1:0"
  ]
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 85

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [DescribeCache](#) (p. 345)

- [ListLocalDisks](#) (p. 375)

AddUploadBuffer

Description

This operation configures one or more gateway local disks as upload buffer space for a specified gateway. This operation is supported for both the gateway-stored and gateway-cached volume architectures (see [How AWS Storage Gateway Works](#) (p. 3)).

In the request, you specify the gateway Amazon Resource Name (ARN) to which you want to add upload buffer space, and one or more disk IDs that you want to configure as an upload buffer.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.AddUploadBuffer

{
  "GatewayARN": "String",
  "DiskIds": [
    "String",
    ...
  ]
}
```

JSON Fields

DiskIds

An array of strings that identify disks that are to be configured as upload buffer space. Each string in the array must be minimum length of 1 and maximum length of 300. You can get disk IDs from the [ListLocalDisks](#) (p. 375) API.

Required: Yes

Type: Array

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 372) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The ARN of the gateway for which the upload buffer was configured.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- DiskAlreadyAllocated
- DiskDoesNotExist
- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- LocalStorageLimitExceeded
- NotSupported

Examples

Example Request

The following example shows a request that specifies that two local disks of a gateway are to be configured as upload buffer.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
```



```
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.AddUploadBuffer

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"

  "DiskIds": [
    "pci-0000:03:00.0-scsi-0:0:0:0",
    "pci-0000:03:00.0-scsi-0:0:1:0"
  ]
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 85

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [DescribeUploadBuffer](#) (p. 367)
- [ListLocalDisks](#) (p. 375)

AddWorkingStorage

Description

This operation configures one or more gateway local disks as working storage for a gateway. This operation is supported only for the gateway-stored volume architecture (see [How AWS Storage Gateway Works](#) (p. 3)).

Note

Working storage is also referred to as the upload buffer. You can also use the [AddUploadBuffer](#) (p. 313) operation to add an upload buffer to a stored-volume gateway.

In the request, you specify the gateway Amazon Resource Name (ARN) to which you want to add working storage, and one or more disk IDs that you want to configure as working storage.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
```

```
x-amz-date: date
x-amz-target: StorageGateway_20120630.AddWorkingStorage

{
  "GatewayARN": "String",
  "DiskIds": [
    "String",
    ...
  ]
}
```

JSON Fields

DiskIds

An array of strings that identify disks that are to be configured as working storage. Each string in the array must be minimum length of 1 and maximum length of 300. You can get the disk IDs from the [ListLocalDisks \(p. 375\)](#) API.

Required: Yes

Type: Array

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The ARN of the gateway for which working storage was configured.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- DiskAlreadyAllocated
- DiskDoesNotExist
- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- LocalStorageLimitExceeded
- NotSupported
- UnsupportedOperationForGatewayType

Examples

Example Request

The following example shows a request that specifies that two local disks of a gateway are to be configured as working storage.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.AddWorkingStorage

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"

  "DiskIds": [
    "pci-0000:03:00.0-scsi-0:0:0:0",
    "pci-0000:03:00.0-scsi-0:0:1:0"
  ]
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 85

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [DescribeWorkingStorage](#) (p. 370)
- [ListLocalDisks](#) (p. 375)

CreateCachediSCSIVolume

Description

This operation creates a cached volume on a specified cached gateway. This operation is supported only for the gateway-cached volume architecture (see [How AWS Storage Gateway Works](#) (p. 3)).

Note

Cache storage must be allocated to the gateway before you can create a cached volume. Use the [AddCache](#) (p. 310) operation to add cache storage to a gateway.

In the request, you must specify the gateway, size of the volume in bytes, the iSCSI target name, an IP address on which to expose the target, and a unique client token. In response, AWS Storage Gateway creates the volume and returns information about it such as the volume Amazon Resource Name (ARN), its size, and the iSCSI target ARN that initiators can use to connect to the volume target.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.CreateCachediSCSIVolume

{
  "GatewayARN": "String",
  "VolumeSizeInBytes": "String",
  "SnapshotId": "String",
  "TargetName": "String",
  "NetworkInterfaceId": "String",
  "ClientToken": "String"
}
```

JSON Fields

ClientToken

A unique identifying string for the cached volume.

Length: Minimum length of 5. Maximum length of 100.

Required: Yes

Type: String

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 372) operation to return a list of gateways for your account and region.

Required: yes

Type: String

NetworkInterfaceId

The network interface of the gateway on which to expose the iSCSI target. Only IPv4 addresses are accepted. Use the [DescribeGatewayInformation \(p. 354\)](#) operation to get a list of the network interfaces available on the gateway.

Valid Values: A valid IP address.

Required: Yes

Type: String

SnapshotId

The snapshot ID (e.g., "snap-1122aabb") of the snapshot to restore as the new stored volume. Specify this field if you want to create the iSCSI cached volume from a snapshot; otherwise, do not include this field. To list snapshots for your account, use [DescribeSnapshots](#) in *Amazon Elastic Compute Cloud API Reference*.

Length: 13

Valid Values: Must be a valid snapshot ID, "snap-" followed by eight hexadecimal characters.

Required: No

Type: String

TargetName

The name of the iSCSI target used by initiators to connect to the target and as a suffix for the target ARN. For example, specifying **TargetName** as *myvolume* results in the target ARN of *arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/iqn.1997-05.com.amazon:myvolume*. The target name must be unique across all volumes of a gateway.

Length: Minimum length of 1. Maximum length of 200.

Constraints: The name can contain lowercase letters, numbers, periods (.), and hyphens (-).

Required: Yes

Type: String

VolumeSizeInBytes

The size of the cached volume.

Constraints: The size must be GiB aligned.

Required: Yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date
```

```
{  
  "TargetARN" : "String",  
  "VolumeARN" : "String"  
}
```

JSON Fields

TargetARN

The ARN of the volume target that includes the iSCSI name that initiators can use to connect to the target.

Type: String

VolumeARN

The ARN of the configured volume.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- CannotExportSnapshot
- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- MaximumVolumeCountExceeded
- NetworkConfigurationChanged
- NotSupported
- SnapshotIdInvalid
- StagingAreaFull
- TargetAlreadyExists
- TargetInvalid
- TargetNotFound
- UnsupportedOperationForGatewayType
- VolumeAlreadyExists

Examples

Example Request

The following example shows a request that specifies that a local disk of a gateway be configured as a cached volume.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.CreateCachediSCSIVolume

{
  "ClientToken": "cachedvol112233",
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
  "NetworkInterfaceId": "10.1.1.1",
  "TargetName": "myvolume",
  "VolumeSizeInBytes": 536870912000
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 263

{
  "TargetARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/iqn.1997-05.com.amazon:myvolume",
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB"
}
```

Related Actions

- [ListVolumes](#) (p. 381)
- [ListLocalDisks](#) (p. 375)
- [DeleteVolume](#) (p. 340)
- [DescribeCachediSCSIVolumes](#) (p. 348)

CreateSnapshot

Description

This operation initiates a snapshot of a volume.

AWS Storage Gateway provides the ability to back up point-in-time snapshots of your data to Amazon Simple Storage (Amazon S3) for durable off-site recovery, as well as import the data to an Amazon Elastic Block Store (EBS) volume in Amazon Elastic Compute Cloud (Amazon EC2). You can take snapshots of your gateway volume on a scheduled or ad-hoc basis. This API enables you to take an ad-hoc snapshot. For more information, see [Working with Snapshots](#) (p. 199).

In the `CreateSnapshot` request you identify the volume by providing its Amazon Resource Name (ARN). You must also provide description for the snapshot. When AWS Storage Gateway takes the snapshot of specified volume, the snapshot and its description appear in the AWS Storage Gateway console. In response, AWS Storage Gateway returns you a snapshot ID. You can use this snapshot ID to check the snapshot progress or later use it when you want to create a volume from a snapshot.

Note

To list or delete a snapshot, you must use the Amazon EC2 API. For more information, go to [DeleteSnapshot](#) and [DescribeSnapshots](#) in *Amazon Elastic Compute Cloud API Reference*.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.CreateSnapshot

{
  "VolumeARN": "String",
  "SnapshotDescription": "String"
}
```

JSON Fields

SnapshotDescription

Textual description of the snapshot that appears in the Amazon EC2 console, Elastic Block Store snapshots panel in the **Description** field, and in the AWS Storage Gateway snapshot **Details** pane, **Description** field

Length: Minimum length of 1. Maximum length of 255.

Required: yes

Type: String

VolumeARN

The Amazon Resource Name (ARN) of the volume. Use the [ListVolumes](#) (p. 381) operation to return a list of gateway volumes.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
```



```
Content-length: payloadLength  
Date: date
```

```
{  
  "VolumeARN": "String",  
  "SnapshotId": "String"  
}
```

JSON Fields

SnapshotId

The snapshot ID that is used to refer to the snapshot in future operations such as describing snapshots (Amazon Elastic Compute Cloud API DescribeSnapshots) or creating a volume from a snapshot ([CreateStorediSCSIVolume \(p. 327\)](#)).

Type: String

VolumeARN

The Amazon Resource Name (ARN) of the volume of which the snapshot was taken. Use [ListVolumes \(p. 381\)](#) to get volume ARNs of a gateway.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- VolumeNotFound
- VolumeNotReady

Examples

Example Request

The following example sends a CreateSnapshot request to take snapshot of the specified an example volume.

```
POST / HTTP/1.1  
Host: storagegateway.us-east-1.amazonaws.com  
Content-Type: application/x-amz-json-1.1  
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066ab
```

```
dd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.CreateSnapshot

{
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygate
way/volume/vol-1122AABB",
  "SnapshotDescription": "snapshot description"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllelqeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 137

{
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygate
way/volume/vol-1122AABB",
  "SnapshotId": "snap-78e22663"
}
```

Related Actions

- [UpdateSnapshotSchedule](#) (p. 401)
- [DescribeSnapshotSchedule](#) (p. 360)

CreateSnapshotFromVolumeRecoveryPoint

Description

This operation initiates a snapshot of a volume from a volume recovery point. This operation is supported only for the gateway-cached volume architecture (see [How AWS Storage Gateway Works](#) (p. 3)).

A volume recovery point is a point in time at which all data of the volume is consistent and from which you can create a snapshot. To get a list of volume recovery point for gateway-cached volumes, use [ListVolumeRecoveryPoints](#) (p. 378).

In the `CreateSnapshotFromVolumeRecoveryPoint` request, you identify the volume by providing its Amazon Resource Name (ARN). You must also provide a description for the snapshot. When AWS Storage Gateway takes a snapshot of the specified volume, the snapshot and its description appear in the AWS Storage Gateway console. In response, AWS Storage Gateway returns you a snapshot ID. You can use this snapshot ID to check the snapshot progress or later use it when you want to create a volume from a snapshot.

Note

To list or delete a snapshot, you must use the Amazon EC2 API. For more information, go to [DeleteSnapshot](#) and [DescribeSnapshots](#) in *Amazon Elastic Compute Cloud API Reference*.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.CreateSnapshotFromRecoveryPoint

{
  "VolumeARN": "String",
  "SnapshotDescription": "String"
}
```

JSON Fields

SnapshotDescription

A textual description of the snapshot that appears in the Amazon EC2 console, Elastic Block Store snapshots panel in the **Description** field, and in the AWS Storage Gateway snapshot **Details** pane, **Description** field.

Length: Minimum length of 1. Maximum length of 255.

Required: yes

Type: String

VolumeARN

The Amazon Resource Name (ARN) of the volume. Use the [ListVolumes \(p. 381\)](#) operation to return a list of gateway volumes.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "SnapshotId": "String",
  "VolumeARN": "String",
  "VolumeRecoveryPointTime": "String"
}
```

JSON Fields

SnapshotId

The snapshot ID that is used to refer to the snapshot in future operations such as describing snapshots (Amazon Elastic Compute Cloud API `DescribeSnapshots`) or creating a volume from a snapshot (`CreateStorediSCSIVolume` (p. 327)).

Type: String

VolumeARN

The ARN of the volume of which the snapshot was taken. Obtain volume ARNs from the `ListVolumes` (p. 381) operation.

Type: String

VolumeRecoveryPointTime

The time of the recovery point. Data up to this recovery point are included in the snapshot.

Type: String format of a date in the ISO8601 extended YYYY-MM-DD'T'HH:MM:SS'Z' format.

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses](#) (p. 287).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- UnsupportedOperationForGatewayType
- VolumeNotFound
- VolumeNotReady

Examples

Example Request

The following example sends a `CreateSnapshotFromVolumeRecoveryPoint` request to create snapshot from the recovery point of a volume.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.CreateSnapshotFromVolumeRecoveryPoint

{
```

```
"VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB",  
"SnapshotDescription": "snapshot description"  
}
```

Example Response

```
HTTP/1.1 200 OK  
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0  
Date: Wed, 12 Sep 2012 12:00:02 GMT  
Content-Type: application/x-amz-json-1.1  
Content-length: 137  
  
{  
  "SnapshotId": "snap-78e22663",  
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB",  
  "VolumeRecoveryPointTime": "2012-06-30T10:10:10.000Z"  
}
```

Related Actions

- [ListVolumeRecoveryPoints](#) (p. 378)

CreateStorediSCSIVolume

Description

This operation creates a volume on a specified gateway. This operation is supported only for the gateway-stored volume architecture (see [How AWS Storage Gateway Works](#) (p. 3)).

The size of the volume to create is inferred from the disk size. You can choose to preserve existing data on the disk, create a volume from an existing snapshot, or create an empty volume. If you choose to create an empty volume, any existing data on the disk is erased.

In the request, you must specify the gateway and the disk information on which you are creating the volume. In response, AWS Storage Gateway creates the volume and returns information about it such as the volume Amazon Resource Name (ARN), its size, and the iSCSI target ARN that initiators can use to connect to the volume target.

Request

Syntax

```
POST / HTTP/1.1  
Host: storagegateway.region.amazonaws.com  
Authorization: authorization  
Content-Type: application/x-amz-json-1.1  
x-amz-date: date  
x-amz-target: StorageGateway_20120630.CreateStorediSCSIVolume  
  
{
```

```
"GatewayARN": "String",  
"DiskId": "String",  
"SnapshotId": "String",  
"PreserveExistingData": Boolean,  
"TargetName": "String",  
"NetworkInterfaceId": "String"  
}
```

JSON Fields

DiskId

The unique identifier of the gateway local disk that is configured as a stored volume. Use [ListLocalDisks \(p. 375\)](#) to list disk IDs for a gateway.

Required: Yes

Type: String

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

NetworkInterfaceId

The network interface of the gateway on which to expose the iSCSI target. Only IPv4 addresses are accepted. Use the [DescribeGatewayInformation \(p. 354\)](#) to get a list of the network interfaces available on the gateway.

Valid Values: A valid IP address.

Required: Yes

Type: String

PreserveExistingData

Specify this field as true if you want to preserve the data on the local disk. Otherwise, specifying this field as false creates an empty volume.

Valid Values: true | false

Required: Yes

Type: Boolean

SnapshotId

The snapshot ID (e.g. "snap-1122aabb") of the snapshot to restore as the new stored volume. Specify this field if you want to create the iSCSI storage volume from a snapshot; otherwise, do not include this field. To list snapshots for your account use [DescribeSnapshots](#) in *Amazon Elastic Compute Cloud API Reference*.

Length: 13

Valid Values: Must be a valid snapshot ID, "snap-" followed by eight hexadecimal characters.

Required: No

Type: String

TargetName

The name of the iSCSI target used by initiators to connect to the target and as a suffix for the target ARN. For example, specifying **TargetName** as *myvolume* results in the target ARN of *arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/iqn.1997-05.com.amazon:myvolume*. The target name must be unique across all volumes of a gateway.

Length: Minimum length of 1. Maximum length of 200.

Constraints: The name can contain lowercase letters, numbers, periods (.), and hyphens (-).

Required: Yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "VolumeARN": "String",
  "VolumeSizeInBytes": Number,
  "TargetARN": "String"
}
```

JSON Fields

TargetARN

The ARN of the volume target that includes the iSCSI name that initiators can use to connect to the target.

Type: String

VolumeARN

The ARN of the configured volume.

Type: String

VolumeSizeInBytes

The size of the volume in bytes.

Type: Number

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- CannotExportSnapshot
- DiskAlreadyAllocated
- DiskDoesNotExist

- DiskSizeNotGigAligned
- DiskSizeGreaterThanVolumeMaxSize
- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- LunInvalid
- MaximumVolumeCountExceeded
- NetworkConfigurationChanged
- NotSupported
- SnapshotIdInvalid
- StagingAreaFull
- TargetAlreadyExists
- TargetInvalid
- TargetNotFound
- UnsupportedOperationForGatewayType
- VolumeAlreadyExists

Examples

Example Request

The following example shows a request that specifies that a local disk of a gateway be configured as a volume.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.CreateStorediSCSIVolume

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
  "DiskId": "pci-0000:03:00.0-scsi-0:0:0:0",
  "PreserveExistingData": "true",
  "TargetName": "myvolume",
  "NetworkInterfaceId": "10.1.1.1"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
```



```
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 263
```

```
{
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygate
way/volume/vol-1122AABB",
  "VolumeSizeInBytes": 1099511627776,
  "TargetARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygate
way/target/iqn.1997-05.com.amazon:myvolume"
}
```

Related Actions

- [ListVolumes](#) (p. 381)
- [ListLocalDisks](#) (p. 375)
- [DeleteVolume](#) (p. 340)
- [DescribeStorediSCSIVolumes](#) (p. 363)

DeleteBandwidthRateLimit

Description

This operation deletes the bandwidth rate limits of a gateway. You can delete either the upload and download bandwidth rate limit, or you can delete both. If you delete only one of the limits, the other limit remains unchanged. To specify which gateway to work with, use the Amazon Resource Name (ARN) of the gateway in your request.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DeleteBandwidthRateLimit

{
  "GatewayARN": "String",
  "BandwidthType": "String"
}
```

JSON Fields

BandwidthType

One of the [BandwidthType](#) (p. 411) values that indicates the gateway bandwidth rate limit to delete.

Valid Values: UPLOAD | DOWNLOAD | ALL

Required: Yes

Type: String

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway whose bandwidth rate information was deleted.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- BandwidthThrottleScheduleNotFound
- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported

Examples

The following example shows a request that deletes both of the bandwidth rate limits of a gateway.

Example Request

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DeleteBandwidthRateLimit

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
  "BandwidthType": "ALL"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 85

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [DescribeBandwidthRateLimit](#) (p. 343)
- [UpdateBandwidthRateLimit](#) (p. 389)

DeleteChapCredentials

Description

This operation deletes Challenge-Handshake Authentication Protocol (CHAP) credentials for a specified iSCSI target and initiator pair.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DeleteChapCredentials
```

```
{  
  "TargetARN": "String",  
  "InitiatorName": "String"  
}
```

JSON Fields

InitiatorName

The iSCSI initiator that connects to the target.

Length: Minimum length of 1. Maximum length of 255.

Valid Values: The initiator name can contain lowercase letters, numbers, periods (.), and hyphens (-).

Required: Yes

Type: String

TargetARN

The Amazon Resource Name (ARN) of the iSCSI volume target. Use the [DescribeStorediSCSIVolumes \(p. 363\)](#) operation to return to retrieve the TargetARN for specified VolumeARN.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK  
x-amzn-RequestId: x-amzn-RequestId  
Content-Type: application/x-amz-json-1.1  
Content-length: payloadLength  
Date: date  
  
{  
  "TargetARN": "String",  
  "InitiatorName": "String"  
}
```

JSON Fields

InitiatorName

The iSCSI initiator that connects to the target.

Type: String

TargetARN

The Amazon Resource Name (ARN) of the target.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- InitiatorNotFound
- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- TargetInvalid
- TargetNotFound

Examples

Example Request

The following example shows a request that deletes the CHAP credentials for an iSCSI target myvolume.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DeleteChapCredentials

{
  "TargetARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/iqn.1997-05.com.amazon:myvolume",
  "InitiatorName": "iqn.1991-05.com.microsoft:computername.domain.example.com"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 203

{
  "TargetARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/iqn.1997-05.com.amazon:myvolume",
  "InitiatorName": "iqn.1991-05.com.microsoft:computername.domain.example.com"
}
```

Related Actions

- [UpdateChapCredentials](#) (p. 391)
- [DescribeChapCredentials](#) (p. 352)

DeleteGateway

Description

This operation deletes a gateway. To specify which gateway to delete, use the Amazon Resource Name (ARN) of the gateway in your request. The operation deletes the gateway; however, it does not delete the gateway virtual machine (VM) from your host computer.

After you delete a gateway, you cannot reactivate it. Completed snapshots of the gateway volumes are not deleted upon deleting the gateway, however, pending snapshots will not complete. After you delete a gateway, your next step is to remove it the VM from your environment. Reusing the VM for a new gateway is not supported.

Important

You no longer pay software charges after the gateway is deleted; however, your existing Amazon EBS snapshots persist and you will continue to be billed for these snapshots. You can choose to remove all remaining Amazon EBS snapshots by canceling your Amazon EC2 subscription. If you prefer not to cancel your Amazon EC2 subscription, you can delete your snapshots using the Amazon EC2 console. For more information, see the [AWS Storage Gateway Detail Page](#).

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DeleteGateway

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 372) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the deleted gateway.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayNotFound
- InternalError
- InvalidParameters
- NotSupported

Examples

Example Request

The following example shows a request that deletes a gateway.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DeleteGateway

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 85

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [ListGateways](#) (p. 372)
- [ShutdownGateway](#) (p. 384)

DeleteSnapshotSchedule

Description

This operation deletes a snapshot schedule of a volume.

You can take snapshots of your gateway volumes on a scheduled or ad-hoc basis. This API enables you to delete a snapshot schedule for a volume. For more information, see [Working with Snapshots](#) (p. 199).

In the `DeleteSnapshotSchedule` request, you identify the volume by providing its Amazon Resource Name (ARN).

Note

To list or delete a snapshot, you must use the Amazon EC2 API. For more information, go to [DeleteSnapshot](#) and [DescribeSnapshots](#) in *Amazon Elastic Compute Cloud API Reference*.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DeleteSnapshotSchedule

{
  "VolumeARN": "String"
}
```

JSON Fields

VolumeARN

The Amazon Resource Name (ARN) of the volume. Use the [ListVolumes](#) (p. 381) operation to return a list of gateway volumes.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "VolumeARN": "String"
}
```

JSON Fields

VolumeARN

The ARN of the volume of which the snapshot was taken.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- UnsupportedOperationForGatewayType
- VolumeNotFound
- VolumeNotReady

Examples

Example Request

The following example shows a request that deletes a volume.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
```

```
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DeleteSnapshotSchedule

{
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 137

{
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB"
}
```

Related Actions

- [UpdateSnapshotSchedule](#) (p. 401)
- [DescribeSnapshotSchedule](#) (p. 360)

DeleteVolume

Description

This operation deletes the specified gateway volume that you created previously using either the [CreateCachediSCSIVolume](#) (p. 318) or [CreateStorediSCSIVolume](#) (p. 327) operation. For gateway-stored volumes, the local disk that was configured as the storage volume is not deleted. You can reuse the local disk to create another storage volume.

Before you delete a gateway volume, make sure there are no iSCSI connections to the volume you are deleting. You should also make sure there is no snapshot in progress. You can use the Amazon Elastic Compute Cloud (Amazon EC2) API to query snapshots on the volume you are deleting and check the snapshot status. For more information, go to [DescribeSnapshots](#) in *Amazon Elastic Compute Cloud API Reference*.

In the request, you must provide the Amazon Resource Name (ARN) of the storage volume you want to delete.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DeleteVolume

{
  "VolumeARN": "String"
}
```

JSON Fields

VolumeARN

The Amazon Resource Name (ARN) of the volume. Use the [ListVolumes \(p. 381\)](#) operation to return a list of gateway volumes.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "VolumeARN": "String"
}
```

JSON Fields

VolumeARN

The Amazon Resource Name (ARN) of the storage volume that was deleted. It is the same ARN you provided in the request.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError

- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- SnapshotInProgressException
- TargetNotFound
- VolumeIdInvalid
- VolumeInUse
- VolumeNotFound

Examples

Example Request

The following example shows a request that deletes a volume.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DeleteVolume

{
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 104

{
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB"
}
```

Related Actions

- [CreateStorediSCSIVolume \(p. 327\)](#)
- [ListLocalDisks \(p. 375\)](#)

DescribeBandwidthRateLimit

Description

This operation returns the bandwidth rate limits of a gateway. By default, these limits are not set, which means no bandwidth rate limiting is in effect.

This operation only returns a value for a bandwidth rate limit only if the limit is set. If no limits are set for the gateway, then this operation returns only the gateway ARN in the response body. To specify which gateway to describe, use the Amazon Resource Name (ARN) of the gateway in your request.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DescribeBandwidthRateLimit

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String",
  "AverageUploadRateLimitInBitsPerSec": Number,
  "AverageDownloadRateLimitInBitsPerSec": Number
}
```

JSON Fields

AverageDownloadRateLimitInBitsPerSec

The average download bandwidth rate limit in bits per second. This field does not appear in the response if the download rate limit is not set.

Type: Number

AverageUploadRateLimitInBitsPerSec

The average upload bandwidth rate limit in bits per second. This field does not appear in the response if the upload rate limit is not set.

Type: Number

GatewayARN

The Amazon Resource Name (ARN) of the gateway whose rate bandwidths are described.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported

Examples

Example Request

The following example shows a request that returns the bandwidth throttle properties of a gateway.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DescribeBandwidthRateLimit

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 182

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygate
way",
  "AverageUploadRateLimitInBitsPerSec": 102400,
  "AverageDownloadRateLimitInBitsPerSec": 51200
}
```

Related Actions

- [UpdateBandwidthRateLimit](#) (p. 389)
- [DeleteBandwidthRateLimit](#) (p. 331)

DescribeCache

Description

This operation returns information about the cache of a gateway. This operation is supported only for the gateway-cached volume architecture (see [How AWS Storage Gateway Works](#) (p. 3)).

The response includes disk IDs that are configured as cache, and it includes the amount of cache allocated and used.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DescribeCache

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 372) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "CacheAllocationInBytes": Number,
  "CacheDirtyPercentage": Number,
  "CacheHitPercentage": Number,
  "CacheMissPercentage": Number,
  "CacheUsedPercentage": Number,
  "DiskIds":
    [ String,
      ...
    ],
  "GatewayARN": String
}
```

JSON Fields

CacheAllocatedInBytes

The size allocated, in bytes, for the cache. If no cache is defined for the gateway, this field returns 0.

Type: Number

CacheDirtyPercentage

The percentage (0 to 100) of the cache that contains data that has not yet been persisted to Amazon S3. If no cache is defined for the gateway, this field returns 0.

Type: Number

CacheHitPercentage

The percentage (0 to 100) of data read from the storage volume that was read from cache. If no cache is defined for the gateway, this field returns 0.

Type: Number

CacheMissPercentage

The percentage (0 to 100) of data read from the storage volume that was not read from the cache, but was read from Amazon S3. If no cache is defined for the gateway, this field returns 0.

Type: Number

CacheUsedPercentage

The percentage (0 to 100) of the cache storage in use. If no cache is defined for the gateway, this field returns 0.

Type: Number

DiskIds

An array of the gateway's local disk IDs that are configured as cache. Each local disk ID is specified as a string (minimum length of 1 and maximum length of 300). If no local disks are configured as cache, then the `DiskIds` array is empty.

Type: Array

GatewayARN

In response, AWS Storage Gateway returns the ARN of the activated gateway. If you don't remember the ARN of a gateway, you can use the List Gateways operations to return a list of gateways for your account and region.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- UnsupportedOperationForGatewayType

Examples

Example Request

The following example shows a request to obtain a description of a gateway's working storage.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DescribeCache

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
```

```
Content-Type: application/x-amz-json-1.1
Content-length: 271

{
  "CacheAllocationInBytes": 2199023255552,
  "CacheDirtyPercentage": 0.07,
  "CacheHitPercentage": 99.68,
  "CacheMissPercentage": 0.32,
  "CacheUsedPercentage": 0.07,
  "DiskIds": [
    "pci-0000:03:00.0-scsi-0:0:0:0",
    "pci-0000:04:00.0-scsi-0:1:0:0"
  ],
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [AddCache](#) (p. 310)
- [ListLocalDisks](#) (p. 375)

DescribeCachediSCSIVolumes

Description

This operation returns a description of the gateway volumes specified in the request. This operation is supported only for the gateway-cached volume architecture (see [How AWS Storage Gateway Works](#) (p. 3)).

The list of gateway volumes in the request must be from one gateway. In the response Amazon Storage Gateway returns volume information sorted by volume Amazon Resource Name (ARN).

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DescribeCachediSCSIVolumes

{
  "VolumeARNs": [ "String", ... ]
}
```

JSON Fields

VolumeARNs

An array of strings, where each string represents the ARN of a cached volume. All of the specified cached volumes must be from the same gateway. Use [ListVolumes](#) (p. 381) to get volume ARNs of a gateway.

Required: Yes

Type: Array

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "CachediSCSIVolumes":
  [
    {
      "VolumeiSCSIAttributes":
      {
        "ChapEnabled": Boolean,
        "LunNumber": Number,
        "NetworkInterfaceId": "String",
        "NetworkInterfacePort": Number,
        "TargetARN": "String"
      },
      "VolumeARN": "String",
      "VolumeId": "String",
      "VolumeSizeInBytes": Number,
      "VolumeStatus": "String",
      "VolumeType": "String",
      "SourceSnapshotId": "String"
    },
    ...
  ]
}
```

JSON Fields

CachediSCSIVolumes

An array of [CachediSCSIVolume \(p. 404\)](#) objects where each object contains metadata about one cached volume.

Type: Object

ChapEnabled

Indicates whether mutual CHAP is enabled for the iSCSI target.

Type: String

LunNumber

The logical disk number.

Type: String

NetworkInterfaceId

The network interface ID of the cached volume that initiators use to map the cached volume as an iSCSI target.

Type: String

NetworkInterfacePort

The port used to communicate with iSCSI targets.

Type: Number

SourceSnapshotId

If the cached volume was created from a snapshot, this field contains the snapshot ID used, for example, snap-1122aabb. Otherwise, this field is not included.

Type: String

TargetARN

The ARN of the volume target.

Type: String

VolumeARN

The ARN of the stored volume.

Type: String

VolumeId

The unique identifier of the storage volume, e.g. vol-1122AABB.

Type: String

VolumeiSCSIAttributes

An [VolumeiSCSIAttributes](#) (p. 410) object that represents a collection of iSCSI attributes for one stored volume.

Type: Object

VolumeSizeInBytes

The size of the volume in bytes that was specified in the [CreateCachediSCSIVolume](#) (p. 318) operation.

Type: Number

VolumeStatus

One of the [VolumeStatus](#) (p. 412) values that indicates the state of the volume.

Type: String

VolumeType

One of the enumeration values describing the type of volume. Currently, only STORED iSCSI volumes are supported.

Type: [VolumeType](#) (p. 412)

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses](#) (p. 287).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- MaximumVolumeCountExceeded
- NotSupported
- UnsupportedOperationForGatewayType

- VolumeNotFound

Examples

Example Request

The following example shows a request that returns a description of a volume.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DescribeCachediSCSIVolumes

{
  "VolumeARNs": [ "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB" ]
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 664

{
  "CachediSCSIVolumes": [
    {
      "VolumeiSCSIAttributes": {
        "ChapEnabled": true,
        "LunNumber": 0,
        "NetworkInterfaceId": "10.243.43.207",
        "NetworkInterfacePort": 3260,
        "TargetARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/iqn.1997-05.com.amazon:myvolume"
      },
      "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB",
      "VolumeDiskId": "pci-0000:03:00.0-scsi-0:0:0:0",
      "VolumeId": "vol-1122AABB",
      "VolumeSizeInBytes": 1099511627776,
      "VolumeStatus": "AVAILABLE",
      "VolumeType": "CACHED iSCSI"
    }
  ]
}
```

Related Actions

- [CreateCachediSCSIVolume](#) (p. 318)
- [DescribeUploadBuffer](#) (p. 367)
- [ListLocalDisks](#) (p. 375)

DescribeChapCredentials

Description

This operation returns an array of Challenge-Handshake Authentication Protocol (CHAP) credentials information for a specified iSCSI target, one for each target-initiator pair.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DescribeChapCredentials

{
  "TargetARN": "String"
}
```

JSON Fields

TargetARN

The Amazon Resource Name (ARN) of the iSCSI volume target. Use the [DescribeStorediSCSIVolumes](#) (p. 363) operation to return to retrieve the TargetARN for specified VolumeARN.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "ChapCredentials": [
```

```
{
  {
    "InitiatorName": "String",
    "SecretToAuthenticateInitiator": "String",
    "SecretToAuthenticateTarget": "String",
    "TargetARN": "String"
  },
  ...
}
```

JSON Fields

ChapCredentials

An array of [ChapInfo](#) (p. 406) objects that represents CHAP credentials. Each object in the array contains CHAP credential information for one target-initiator pair. If no CHAP credentials are set, an empty array is returned.

Type: Array

InitiatorName

The iSCSI initiator that connects to the target.

Type: String

SecretToAuthenticateInitiator

The secret key that the initiator (e.g. Windows client) must provide to participate in mutual CHAP with the target.

Type: String

SecretToAuthenticateTarget

The secret key that the target must provide to participate in mutual CHAP with the initiator (e.g. Windows client).

Type: String

TargetARN

The Amazon Resource Name (ARN) of the storage volume.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses](#) (p. 287).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- TargetInvalid
- TargetNotFound

Examples

Example Request

The following example shows a request that returns the CHAP credentials of an iSCSI target.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DescribeChapCredentials

{
  "TargetARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/iqn.1997-05.com.amazon:myvolume"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 267

{
  "ChapCredentials": {
    "TargetName": "iqn.1997-05.com.amazon:myvolume",
    "SecretToAuthenticateInitiator": "111111111111",
    "InitiatorName": "iqn.1991-05.com.microsoft:computername.domain.example.com",
    "SecretToAuthenticateTarget": "222222222222"
  }
}
```

Related Actions

- [DeleteChapCredentials](#) (p. 333)
- [UpdateChapCredentials](#) (p. 391)

DescribeGatewayInformation

Description

This operation returns metadata about a gateway such as its name, network interfaces, configured time zone, and the state (whether the gateway is running or not). To specify which gateway to describe, use the Amazon Resource Name (ARN) of the gateway in your request.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DescribeGatewayInformation

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String",
  "GatewayId": "String",
  "GatewayNetworkInterfaces": [
    { "MacAddress": "String",
      "IPv4Address": "String",
      "IPv6Address": "String"
    },
    ...
  ],
  "GatewayState": "String",
  "GatewayTimezone": "String",
  "GatewayType": "String",
  "NextUpdateAvailabilityDate": "String"
}
```

JSON Fields

GatewayARN

The ARN of the gateway that is described in the response. It is the same gateway ARN you send with the request.

Type: String

GatewayId

The gateway ID.

Type: String

GatewayNetworkInterfaces

A [NetworkInterface](#) (p. 408) array that contains descriptions of the gateway network interfaces.

Type: Array

GatewayState

One of the [GatewayState](#) (p. 411) values that indicates the operating state of the gateway.

Type: String

GatewayTimezone

One of the [GatewayTimezone](#) (p. 411) values that indicates the time zone configured for the gateway.

Type: String

GatewayType

The type of gateway, such as `CACHED` or `STORED`.

Type: String

Ipv4Address

The Internet Protocol version 4 (IPv4) address of an interface of the gateway.

Type: String

Ipv6Address

The Internet Protocol version 6 (IPv6) address of an interface of the gateway. Currently not supported.

Type: String

MacAddress

The Media Access Control address (MAC address) of a gateway network interface. Currently not supported.

Type: String

NextUpdateAvailabilityDate

The date at which an update to the gateway is available. This date is in the time zone of the gateway. If the gateway is not available for an update, this field is not returned in the response.

Type: String format of a date in the ISO8601 extended YYYY-MM-DD'T'HH:MM:SS'Z' format.

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses](#) (p. 287).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy

- [InvalidParameters](#)

Examples

Example Request

The following example shows a request for describing a gateway.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DescribeGatewayInformation

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 268

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
  "GatewayId": "sgw-AABB1122",
  "GatewayNetworkInterfaces": [
    {
      "Ipv4Address": "10.35.69.216"
    }
  ],
  "GatewayState": "RUNNING",
  "GatewayTimezone": "GMT-8:00",
  "GatewayType": "CACHED"
}
```

Related Actions

- [ListGateways](#) (p. 372)

DescribeMaintenanceStartTime

Description

This operation returns your gateway's weekly maintenance start time including the day and time of the week. Note that values are in terms of the gateway's time zone.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DescribeMaintenanceStartTime

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String",
  "HourOfDay": Number,
  "MinuteOfHour": Number,
  "DayOfWeek": Number,
  "Timezone": "String"
}
```

JSON Fields

DayOfWeek

The ordinal number that represents the day of the week, where 0 represents Sunday and 6 represents Saturday. The day of week is in the time zone of the gateway.

Type: Number. Between 0 and 6.

GatewayARN

The Amazon Resource Name (ARN) of the gateway for which the maintenance time is described.

Type: String

HourOfDay

The hour component of the maintenance start time represented as *hh*, where *hh* is the hour (0 to 23). The hour of the day is in the time zone of the gateway.

Type: Number

MinuteOfHour

The minute component of the maintenance start time represented as *mm*, where *mm* is the minute (0 to 59). The minute of the hour is in the time zone of the gateway.

Type: Number

Timezone

One of the [GatewayTimezone \(p. 411\)](#) values that indicates the time zone that is set for the gateway. The start time and day of week specified should be in the time zone of the gateway.

Type: String.

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported

Examples

Example Request

The following example shows a request that describes a gateway's maintenance window.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
```

```
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DescribeMaintenanceStartTime

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 173

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
  "HourOfDay": 15,
  "MinuteOfHour": 25,
  "DayOfWeek": 2,
  "Timezone": "GMT+7:00"
}
```

Related Actions

- [UpdateMaintenanceStartTime](#) (p. 399)

DescribeSnapshotSchedule

Description

This operation describes the snapshot schedule of a specified gateway volume. The snapshot schedule information includes intervals at which snapshots are automatically initiated on the volume.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DescribeSnapshotSchedule

{
  "VolumeARN": "String"
}
```

JSON Fields

VolumeARN

The Amazon Resource Name (ARN) of the volume. Use the [ListVolumes \(p. 381\)](#) operation to return a list of gateway volumes.

Required: yes

Type: String

Response

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date

{
  "VolumeARN": "String",
  "StartAt": Number,
  "RecurrenceInHours": Number,
  "Description": "String",
  "Timezone": "String"
}
```

JSON Fields

Description

The snapshot description.

Type: String

RecurrenceInHours

The number of hours between snapshots.

Type: Number. One of the values 1 | 2 | 4 | 8 | 12 | 24.

StartAt

The hour of the day at which the snapshot schedule begins represented as *hh*, where *hh* is the hour (0 to 23). The hour of the day is in the time zone of the gateway.

Type: Number.

Timezone

One of the [GatewayTimezone \(p. 411\)](#) values that indicates the time zone of the gateway.

Type: String

VolumeARN

The Amazon Resource Name (ARN) of the volume that was specified in the request.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- VolumeNotFound

Examples

The following example shows a request that retrieves the snapshot schedule for a volume.

Example Request

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DescribeSnapshotSchedule

{
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 230

{
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB",
  "StartAt": 6,
  "RecurrenceInHours": 24,
  "Description": "sgw-AABB1122:vol-AABB1122:Schedule",
  "Timezone": "GMT+7:00"
}
```


Related Actions

- [UpdateSnapshotSchedule](#) (p. 401)

DescribeStorediSCSIVolumes

Description

This operation returns a description of the gateway volumes specified in the request. This operation is supported only for the gateway-stored volume architecture (see [How AWS Storage Gateway Works](#) (p. 3)).

The list of gateway volumes in the request must be from one gateway. In the response, AWS Storage Gateway returns volume information sorted by volume Amazon Resource Name (ARN).

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DescribeStorediSCSIVolumes

{
  "VolumeARNs": [ "String", ... ]
}
```

JSON Fields

VolumeARNs

An array of strings, where each string represents the ARN of a stored volume. All of the specified stored volumes must be from the same gateway. Use [ListVolumes](#) (p. 381) to get volume ARNs of a gateway.

Required: Yes

Type: Array

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "StorediSCSIVolumes":
```

```
[
  {
    "VolumeiSCSIAttributes": {
      "ChapEnabled": Boolean,
      "NetworkInterfaceId": "String",
      "NetworkInterfacePort": Number,
      "TargetARN": "String",
      "LunNumber": Number
    },
    "PreservedExistingData": Boolean,
    "VolumeARN": "String",
    "VolumeDiskId": "String",
    "VolumeId": "String",
    "VolumeType": "String",
    "VolumeStatus": "String",
    "VolumeSizeInBytes": Number,
    "VolumeProgress": Number,
    "SourceSnapshotId": "String"
  },
  ...
]
```

JSON Fields

ChapEnabled

Indicates whether mutual CHAP is enabled for the iSCSI target.

Type: String

LunNumber

The logical disk number.

Type: String

NetworkInterfaceId

The network interface ID of the stored volume that initiators use to map the stored volume as an iSCSI target.

Type: String

NetworkInterfacePort

The port used to communicate with iSCSI targets.

Type: Number

PreservedExistingData

Indicates if when the stored volume was created, existing data on the underlying local disk was preserved.

Valid Values: true | false

Type: Boolean

SourceSnapshotId

If the stored volume was created from a snapshot, this field contains the snapshot ID used, for example, snap-1122aabb. Otherwise, this field is not included.

Type: String

StorediSCSIVolumes

An array of [StorediSCSIVolume](#) (p. 408) objects where each object contains metadata about one stored volume.

Type: Array

TargetARN

The ARN of the volume target.

Type: String

VolumeARN

The ARN of the stored volume.

Type: String

VolumeDiskId

The disk ID of the local disk that was specified in the [CreateStorediSCSIVolume \(p. 327\)](#) operation.

Type: String

VolumeId

The unique identifier of the storage volume, e.g. vol-1122AABB.

Type: String

VolumeiSCSIAttributes

An [VolumeiSCSIAttributes \(p. 410\)](#) object that represents a collection of iSCSI attributes for one stored volume.

Type: Object

VolumeProgress

Represents the percentage complete if the volume is restoring or bootstrapping that represents the percent of data transferred. This field does not appear in the response if the stored volume is not restoring or bootstrapping.

Type: Number (double)

VolumeSizeInBytes

The size of the volume in bytes.

Type: Number

VolumeStatus

One of the [VolumeStatus \(p. 412\)](#) values that indicates the state of the volume.

Type: String

VolumeType

One of the enumeration values describing the type of volume. Currently, only STORED iSCSI volumes are supported.

Type: [VolumeType \(p. 412\)](#)

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- MaximumVolumeCountExceeded

- NotSupported
- UnsupportedOperationForGatewayType
- VolumeNotFound

Examples

Example Request

The following example shows a request that returns a description of a volume.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DescribeStorediSCSIVolumes

{
  "VolumeARNs": ["arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB"]
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glllelqeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 664

{
  "StorediSCSIVolumes": [
    {
      "VolumeiSCSIAttributes": {
        "ChapEnabled": true,
        "LunNumber": 0,
        "NetworkInterfaceId": "10.243.43.207",
        "NetworkInterfacePort": 3260,
        "TargetARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/ign.1997-05.com.amazon:myvolume"
      },
      "PreservedExistingData": false,
      "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB",
      "VolumeDiskId": "pci-0000:03:00.0-scsi-0:0:0:0",
      "VolumeId": "vol-1122AABB",
      "VolumeProgress": 23.7,
      "VolumeSizeInBytes": 1099511627776,
      "VolumeStatus": "BOOTSTRAPPING",
      "VolumeType": "STORED iSCSI"
    }
  ]
}
```

```
]
}
```

Related Actions

- [CreateStorediSCSIVolume](#) (p. 327)
- [DescribeWorkingStorage](#) (p. 370)
- [ListLocalDisks](#) (p. 375)

DescribeUploadBuffer

Description

This operation returns information about the upload buffer of a gateway. This operation is supported for both the gateway-stored and gateway-cached volume architectures (see [How AWS Storage Gateway Works](#) (p. 3)).

The response includes disk IDs that are configured as upload buffer space, and it includes the amount of upload buffer space allocated and used.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DescribeUploadBuffer

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 372) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "DiskIds":
  [ "String",
    ...
  ],
  "GatewayARN": "String",
  "UploadBufferAllocatedInBytes": Number,
  "UploadBufferUsedInBytes": Number
}
```

JSON Fields

DiskIds

An array of the gateway's local disk IDs that are configured as working storage. Each local disk ID is specified as a string (minimum length of 1 and maximum length of 300). If no local disks are configured as working storage, then the `DiskIds` array is empty.

Type: Array

GatewayARN

In response, AWS Storage Gateway returns the ARN of the activated gateway. If you don't remember the ARN of a gateway, you can use the List Gateways operations to return a list of gateways for your account and region.

Type: String

UploadBufferAllocatedInBytes

The total upload buffer in bytes allocated for the gateway. If no upload buffer is configured for the gateway, this field returns 0.

Type: Number

UploadBufferUsedInBytes

The total upload buffer in bytes in use by the gateway. If no upload buffer is configured for the gateway, this field returns 0.

Type: Number

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy

- InternalError
- InvalidParameters
- NotSupported

Examples

Example Request

The following example shows a request to obtain a description of a gateway's working storage.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DescribeUploadBuffer

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 271

{
  "DiskIds": [
    "pci-0000:03:00.0-scsi-0:0:0:0",
    "pci-0000:04:00.0-scsi-0:1:0:0"
  ],
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
  "UploadBufferAllocatedInBytes": 161061273600,
  "UploadBufferUsedInBytes": 0
}
```

Related Actions

- [AddUploadBuffer](#) (p. 313)
- [ListLocalDisks](#) (p. 375)

DescribeWorkingStorage

Description

This operation returns information about the working storage of a gateway. This operation is supported only for the gateway-stored volume architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)).

Note

Working storage is also referred to as the upload buffer. You can also use the [DescribeUploadBuffer \(p. 367\)](#) operation to add an upload buffer to a stored-volume gateway.

The response includes disk IDs that are configured as working storage, and it includes the amount of working storage allocated and used.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.DescribeWorkingStorage

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "DiskIds":
    [ "String",
      ...
    ]
}
```



```
  ],  
  "GatewayARN": "String",  
  "WorkingStorageUsedInBytes": Number,  
  "WorkingStorageAllocatedInBytes": Number  
}
```

JSON Fields

DiskIds

An array of the gateway's local disk IDs that are configured as working storage. Each local disk ID is specified as a string (minimum length of 1 and maximum length of 300). If no local disks are configured as working storage, then the `DiskIds` array is empty.

Type: Array

GatewayARN

In response, AWS Storage Gateway returns the ARN of the activated gateway. If you don't remember the ARN of a gateway, you can use the List Gateways operations to return a list of gateways for your account and region.

Type: String

WorkingStorageAllocatedInBytes

The total working storage in bytes allocated for the gateway. If no working storage is configured for the gateway, this field returns 0.

Type: Number

WorkingStorageUsedInBytes

The total working storage in bytes in use by the gateway. If no working storage is configured for the gateway, this field returns 0.

Type: Number

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- UnsupportedOperationForGatewayType

Examples

Example Request

The following example shows a request to obtain a description of a gateway's working storage.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.DescribeWorkingStorage

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 271

{
  "DiskIds": [
    "pci-0000:03:00.0-scsi-0:0:0:0",
    "pci-0000:04:00.0-scsi-0:1:0:0"
  ],
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
  "WorkingStorageAllocatedInBytes": 2199023255552,
  "WorkingStorageUsedInBytes": 789207040
}
```

Related Actions

- [AddWorkingStorage](#) (p. 315)
- [ListLocalDisks](#) (p. 375)

ListGateways

Description

This operation lists gateways owned by an AWS account in a region specified in the request. The returned list is ordered by gateway Amazon Resource Name (ARN).

By default, the operation returns a maximum of 100 gateways. This operation supports pagination that allows you to optionally reduce the number of gateways returned in a response.

If you have more gateways than are returned in a response—that is, the response returns only a truncated list of your gateways—the response contains a marker that you can specify in your next request to fetch the next page of gateways.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.ListGateways

{
  "Marker": "String",
  "Limit": Number
}
```

JSON Fields

Limit

Specifies that the list of gateways returned be limited to the specified number of items.

Constraints: Minimum value of 1. Maximum value of 100.

Required: No

Type: Number

Marker

An opaque string that indicates the position at which to begin the returned list of gateways.

Valid Values: A marker obtained from the response of a previous List Gateways request.

Required: No

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "Gateways" : [
    {
      "GatewayARN": "String"
    },
    ...
  ],
  "Marker": "String"
}
```

JSON Fields

Gateways

An array of gateway objects composed of a `GatewayARN` and `GatewayName`.

Type: Array of [GatewayInfo](#) (p. 407) objects.

GatewayARN

The Amazon Resource Name (ARN) of a gateway.

Type: String

Marker

Use the marker in your next request to fetch the next set of gateways in the list. If there are no more gateways to list, this field does not appear in the response.

Type: String | null

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses](#) (p. 287).

- `InternalServerError`
- `InvalidParameters`
- `NotSupported`

Examples

List gateways

The following example does not specify any criteria for the returned list. Note that the request body is "{}". The response returns gateways (or up to the first 100) in the specified region owned by the AWS account.

Example Request

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.ListGateways
{}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 212
```

```
{
  "GatewayList": [
    {
      "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
      "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway2"
    }
  ]
}
```

Related Actions

- [DescribeGatewayInformation](#) (p. 354)

ListLocalDisks

Description

This operation returns a list of the local disks of a gateway.

To specify which gateway to describe for this operation, you use the Amazon Resource Name (ARN) of the gateway in the body of the request.

The request returns all disks, specifying which are configured as working storage, stored volume, or not configured at all.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.ListLocalDisks
```

```
{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 372) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "Disks": [
    {
      "DiskAllocationResource": "String",
      "DiskAllocationType": "String",
      "DiskId": "String",
      "DiskNode": "String",
      "DiskPath": "String",
      "DiskSizeInBytes": Number
    },
    ...
  ],
  "GatewayARN": "String"
}
```

JSON Fields

DiskAllocationResource

If the disk is configured as a volume, then this field contains information about the volume, including volume ID and target name. This field is included in the response only if the local disk is configured as a volume. The format of this field is *targetIdqn::LUNNumber::region-volumeId*.

Type: String

DiskAllocationType

One of the [DiskAllocationType](#) (p. 411) enumeration values that identifies how the local disk is used.

Type: String

DiskId

The unique device ID or other distinguishing data that identify the local disk.

Type: String

DiskNode

The device node of the local disk as assigned by the virtualization environment. You can use this value, for example, in the VMware vSphere client or Microsoft Hyper-V Manager to identify specific disks you want to work with.

Type: String

DiskPath

The path of the local disk in the gateway virtual machine (VM).

Type: String

Disks

An array of [Disk](#) (p. 407) objects.

Type: Array

DiskSizeInBytes

The size of the local disk in bytes

Type: Number

GatewayARN

The Amazon Resource Name (ARN) of the activated gateway whose local disk information is returned.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported

Examples

Example Request

The following example shows a request that returns information about a gateway's local disks.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.ListLocalDisks

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 604

{
```

```
"Disks": [
  {
    "DiskAllocationType": "WORKING STORAGE",
    "DiskId": "pci-0000:03:00.0-scsi-0:0:0:0",
    "DiskNode": "SCSI(0:0)",
    "DiskPath": "/dev/sda",
    "DiskSizeInBytes": 1099511627776
  },
  {
    "DiskAllocationResource": "iqn.1997-05.com.amazon:myvolume::0::us-east-1-vol-1122AABB",
    "DiskAllocationType": "STORED iSCSI VOLUME",
    "DiskId": "pci-0000:03:00.0-scsi-0:0:1:0",
    "DiskNode": "SCSI(0:1)",
    "DiskPath": "/dev/sdb",
    "DiskSizeInBytes": 1099511627776
  }
],
"GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [ListVolumes](#) (p. 381)
- [DescribeGatewayInformation](#) (p. 354)

ListVolumeRecoveryPoints

Description

This operation lists the recovery points for a specified gateway. This operation is supported only for the gateway-cached volume architecture (see [How AWS Storage Gateway Works](#) (p. 3)).

Each gateway-cached volume has one recovery point. A volume recovery point is a point in time at which all data of the volume is consistent and from which you can create a snapshot. To create a snapshot from a volume recovery point, use [CreateSnapshotFromVolumeRecoveryPoint](#) (p. 324).

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.ListVolumeRecoveryPoints

{
  "GatewayARN": "String"
}
```


JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String",
  "VolumeRecoveryPointInfos": [
    {
      "VolumeARN": "String",
      "VolumeRecoveryPointTime": "String",
      "VolumeSizeInBytes": Number,
      "VolumeUsageInBytes": Number
    },
    ...
  ]
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the activated gateway whose local disk information is returned.

Type: String

VolumeARN

The Amazon Resource Name (ARN) of the volume associated with the recovery point.

Type: String

VolumeRecoveryPointInfos

An array of [VolumeRecoveryPointInfo \(p. 410\)](#) objects, where each object describes a recovery point. If no recovery points are defined for the volume, then `VolumeRecoveryPointInfos` is an empty array `[]`.

Type: Array

VolumeRecoveryPointTime

The time of the recovery point.

Type: String format of a date in the ISO8601 extended YYYY-MM-DD'T'HH:MM:SS'Z' format.

VolumeSizeInBytes

The size, in bytes, of the volume to which the recovery point is associated.

Type: Number

VolumeUsageInBytes

The size, in bytes, of the volume in use at the time of the recovery point.

Type: Number

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- UnsupportedOperationForGatewayType

Examples

Example Request

The following example sends a `ListVolumeRecoveryPoints` request to take a snapshot of the specified example volume.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.ListVolumeRecoveryPoints

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 137

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
}
```

```
"VolumeRecoveryPointInfos": [  
  {  
    "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/myg  
ateway/volume/vol-1122AABB",  
    "VolumeRecoveryPointTime": "2012-09-04T21:08:44.627Z",  
    "VolumeSizeInBytes": 536870912000,  
    "VolumeUsageInBytes": 6694048  
  }  
]
```

Related Actions

- [CreateSnapshotFromVolumeRecoveryPoint](#) (p. 324)

ListVolumes

Description

This operation lists the volumes of a gateway. Results are sorted by volume ARN. The response includes only the volume ARNs. If you want additional volume information, use the [DescribeStorediSCSIVolumes](#) (p. 363) API.

The operation supports pagination. By default, the operation returns a maximum of up to 100 volumes. You can optionally specify the `Limit` field in the body to limit the number of volumes in the response. If the number of volumes returned in the response is truncated, the response includes a `Marker` field. You can use this `Marker` value in your subsequent request to retrieve the next set of volumes.

Request

Syntax

```
POST / HTTP/1.1  
Host: storagegateway.region.amazonaws.com  
Authorization: authorization  
Content-Type: application/x-amz-json-1.1  
x-amz-date: date  
x-amz-target: StorageGateway_20120630.ListVolumes  
  
{  
  "GatewayARN": "String",  
  "Marker": "String",  
  "Limit": Number  
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 372) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Limit

Specifies that the list of volumes returned be limited to the specified number of items.

Constraint: Minimum value of 1. Maximum value of 100.

Required: No

Type: Number

Marker

A string that indicates the position at which to begin the returned list of volumes. Obtain the marker from the response of a previous List iSCSI Volumes request.

Required: No

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String",
  "VolumeInfos": [
    { "VolumeARN": "String",
      "VolumeType": "String"
    },
    ...
  ],
  "Marker": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of a gateway.

Type: String

VolumeInfos

An array of [VolumeInfo \(p. 409\)](#) objects, where each object describes an iSCSI volume. If no volumes are defined for the gateway, then `VolumeInfos` is an empty array `[]`.

Type: Array

Marker

Use the marker in your next request to continue pagination of iSCSI volumes. If there are no more volumes to list, this field does not appear in the response body.

Type: String

VolumeARN

The Amazon Resource Name (ARN) of the storage volume.

Type: String

VolumeType

One of the [VolumeType](#) (p. 412) values.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses](#) (p. 287).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported

Examples

Example Request

The List iSCSI Volumes request in this example does not specify a `limit` or `marker` field in the response body. The response returns the volumes (up to the first 100) of the gateway.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.ListVolumes

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 421

{
```

```
"GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygate  
way",  
"VolumeInfos": [  
  {  
    "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/myg  
ateway/volume/vol-1122AABB",  
    "VolumeType": "STORED iSCSI"  
  },  
  {  
    "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/myg  
ateway/volume/vol-3344CCDD",  
    "VolumeType": "STORED iSCSI"  
  }  
]
```

Related Actions

- [ListLocalDisks](#) (p. 375)
- [DescribeStorediSCSIVolumes](#) (p. 363)
- [CreateStorediSCSIVolume](#) (p. 327)

ShutdownGateway

Description

This operation shuts down a gateway. To specify which gateway to shut down, use the Amazon Resource Name (ARN) of the gateway in the body of your request.

The operation shuts down the gateway service component running in the storage gateway's virtual machine (VM) and not the VM itself.

Note

If you want to shut down the VM, it is recommended that you first shut down the gateway component in the VM to avoid unpredictable conditions.

After the gateway is shut down, you cannot call any other API except [StartGateway](#) (p. 387), [DescribeGatewayInformation](#) (p. 354), and [ListGateways](#) (p. 372). For more information, see [ActivateGateway](#) (p. 307). Your applications cannot read from or write to the gateway's storage volumes, and there are no snapshots taken.

Note

When you make a shutdown request, you get a 200 OK success response immediately. However, it might take some time for the gateway to shut down. You can call the Describe Gateway API to check the status. For more information, see [ActivateGateway](#) (p. 307).

If do not intend to use the gateway again, you must delete the gateway ([DeleteGateway](#) (p. 336)) to no longer pay software charges associated with the gateway.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.ShutdownGateway

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway that was shut down.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected

- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported

Examples

Example Request

The following example shows a request that shuts down a gateway.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.ShutdownGateway

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 85

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [StartGateway](#) (p. 387)
- [DeleteGateway](#) (p. 336)
- [ActivateGateway](#) (p. 307)

StartGateway

Description

This operation starts a gateway that you previously shut down (see [ShutdownGateway \(p. 384\)](#)). After the gateway starts, you can now make other API calls, your applications can read from or write to the gateway's storage volumes and you will be able to take snapshot backups.

Note

When you make a request, you get a 200 OK success response immediately. However, it might take some time for the gateway to be ready. You should call `Describe Gateway` and check the status before making any additional API calls. For more information, see [ActivateGateway \(p. 307\)](#).

To specify which gateway to start, use the Amazon Resource Name (ARN) of the gateway in your request.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.StartGateway

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway that was restarted.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported

Examples

Example Request

The following example shows a request that starts a gateway.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.StartGateway

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 85

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [ShutdownGateway](#) (p. 384)
- [DeleteGateway](#) (p. 336)

UpdateBandwidthRateLimit

Description

This operation updates the bandwidth rate limits of a gateway. You can update both the upload and download bandwidth rate limit or specify only one of the two. If you don't set a bandwidth rate limit, the existing rate limit remains.

By default, a gateway's bandwidth rate limits are not set. If you don't set any limit, the gateway does not have any limitations on its bandwidth usage and could potentially use the maximum available bandwidth.

To specify which gateway to update, use the Amazon Resource Name (ARN) of the gateway in your request.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.UpdateBandwidthRateLimit

{
  "GatewayARN": "String",
  "AverageUploadRateLimitInBitsPerSec": Number,
  "AverageDownloadRateLimitInBitsPerSec": Number
}
```

JSON Fields

AverageDownloadRateLimitInBitsPerSec

The average download bandwidth rate limit in bits per second.

Constraint: Minimum value of 102400.

Required: Yes, if `AverageUploadRateLimitInBitsPerSec` is not specified, otherwise, not required.

Type: Number

AverageUploadRateLimitInBitsPerSec

The average upload bandwidth rate limit in bits per second.

Constraint: Minimum value of 51200.

Required: Yes, if `AverageDownloadRateLimitInBitsPerSec` is not specified, otherwise, not required.

Type: Number

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway whose throttle information was updated.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported

Examples

Example Request

The following example shows a request that returns the bandwidth throttle properties of a gateway.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
```

```
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.UpdateBandwidthRateLimit

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
  "AverageUploadRateLimitInBitsPerSec": 51200,
  "AverageDownloadRateLimitInBitsPerSec": 102400
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 85

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [DescribeBandwidthRateLimit](#) (p. 343)
- [DeleteBandwidthRateLimit](#) (p. 331)

UpdateChapCredentials

Description

This operation updates the Challenge-Handshake Authentication Protocol (CHAP) credentials for a specified iSCSI target. By default, a gateway does not have CHAP enabled; however, for added security, you might use it.

Important

When you update CHAP credentials, all existing connections on the target are closed and initiators must reconnect with the new credentials.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
```

```
x-amz-target: StorageGateway_20120630.UpdateChapCredentials

{
  "TargetARN": "String",
  "SecretToAuthenticateInitiator": "String",
  "InitiatorName": "String",
  "SecretToAuthenticateTarget": "String"
}
```

JSON Fields

InitiatorName

The iSCSI initiator that connects to the target.

Length: Minimum length of 1. Maximum length of 255.

Valid Values: The initiator name can contain lowercase letters, numbers, periods (.), and hyphens (-).

Required: Yes

Type: String

SecretToAuthenticateInitiator

The secret key that the initiator (e.g. Windows client) must provide to participate in mutual CHAP with the target.

Length: Minimum length of 12. Maximum length of 16.

Required: Yes

Type: String

SecretToAuthenticateTarget

The secret key that the target must provide to participate in mutual CHAP with the initiator (e.g. Windows client).

Length: Minimum length of 12. Maximum length of 16.

Required: No

Type: String

TargetARN

The Amazon Resource Name (ARN) of the iSCSI volume target. Use the [DescribeStorediSCSIVolumes \(p. 363\)](#) operation to return to retrieve the TargetARN for specified VolumeARN.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
```

```
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "TargetARN": "String",
  "InitiatorName": "String"
}
```

JSON Fields

InitiatorName

The iSCSI initiator that connects to the target. This is the same initiator name specified in the request.

Type: String

TargetARN

The Amazon Resource Name (ARN) of the target. This is the same target specified in the request.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- TargetInvalid
- TargetNotFound

Examples

Example Request

The following example shows a request that updates CHAP credentials for an iSCSI target.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.UpdateChapCredentials

{
```

```
"TargetARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/iqn.1997-05.com.amazon:myvolume",
"SecretToAuthenticateInitiator": "111111111111",
"InitiatorName": "iqn.1991-05.com.microsoft:computername.domain.example.com",

"SecretToAuthenticateTarget": "222222222222"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 203

{
  "TargetARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/iqn.1997-05.com.amazon:myvolume",
  "InitiatorName": "iqn.1991-05.com.microsoft:computername.domain.example.com"
}
```

Related Actions

- [DeleteChapCredentials](#) (p. 333)
- [DescribeChapCredentials](#) (p. 352)
- [Configuring CHAP Authentication for Your Storage Volume](#) (p. 167)

UpdateGatewayInformation

Description

This operation updates a gateway's metadata, which includes the gateway's name and time zone. To specify which gateway to update, use the Amazon Resource Name (ARN) of the gateway in your request.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.UpdateGatewayInformation

{
  "GatewayARN": "String",
  "GatewayName": "String",
  "GatewayTimezone": "String"
}
```


JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

GatewayName

The name of the gateway.

Length: Minimum length of 2. Maximum length of 255.

Required: No

Type: String. Unicode characters with no slashes.

GatewayTimezone

One of the [GatewayTimezone \(p. 411\)](#) values that represents the time zone for your gateway. The time zone is used, for example, when a time stamp is given to a snapshot.

Required: No

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway that was updated.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- InternalError
- InvalidParameters
- NotSupported

Examples

Example Request

The following example shows a request that updates the name of a gateway.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.UpdateGatewayInformation

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
  "GatewayName": "mygateway2"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 85

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway2"
}
```

Related Actions

- [DescribeGatewayInformation](#) (p. 354)
- [ListGateways](#) (p. 372)

UpdateGatewaySoftwareNow

Description

This operation updates the gateway virtual machine (VM) software if an update is available. The request immediately triggers the software update.

Note

When you make this request, you get a 200 OK success response immediately. However, it might take some time for the update to complete. You can call [DescribeGatewayInformation](#) (p. 354) to verify the gateway is in the STATE_RUNNING state. For more information, see [DescribeGatewayInformation](#) (p. 354).

Important

A software update forces a system restart of your gateway. You can minimize the chance of any disruption to your applications by increasing your iSCSI Initiators' timeouts. For more information about increasing iSCSI Initiator timeouts for Windows and Linux, see [Customizing Your Windows iSCSI Settings \(p. 163\)](#) and [Customizing Your Linux iSCSI Settings \(p. 166\)](#), respectively.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.UpdateGatewaySoftwareNow

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses](#) (p. 287).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported

Examples

Example Request

The following example shows a request that initiates a gateway VM update.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.UpdateGatewaySoftwareNow

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 85

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [DescribeMaintenanceStartTime](#) (p. 358)

UpdateMaintenanceStartTime

Description

This operation updates a gateway's weekly maintenance start time information, including day and time of the week. The maintenance time is the time in your gateway's time zone.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.UpdateMaintenanceStartTime

{
  "GatewayARN": "String",
  "HourOfDay": "Number",
  "MinuteOfHour": "Number",
  "DayOfWeek": Number
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 372\)](#) operation to return a list of gateways for your account and region.

Required: yes

Type: String

HourOfDay

The hour component of the maintenance start time represented as *hh*, where *hh* is the hour (00 to 23). The hour of the day is in the time zone of the gateway.

Length: 2

Valid Values: *hh*, where *hh* is the hour (00 to 23).

Required: Yes

Type: Number

MinuteOfHour

The minute component of the maintenance start time represented as *mm*, where *mm* is the minute (00 to 59). The minute of the hour is in the time zone of the gateway.

Length: 2

Valid Values: *mm*, where *mm* is the minute (00 to 59).

Required: Yes

Type: Number

DayOfWeek

The maintenance start time day of the week.

Length: 1

Valid Values An integer between 0 and 6, where 0 represents Sunday and 6 represents Saturday.

Required: Yes

Type: Number

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway whose maintenance start time is updated.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported

Examples

Example Request

The following example shows a request that updates the maintenance start time of `mygateway`.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.UpdateMaintenanceStartTime

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
  "TimeOfDay": 0,
  "MinuteOfHour": 30
  "DayOfWeek": 2
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 85

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Related Actions

- [DescribeMaintenanceStartTime](#) (p. 358)

UpdateSnapshotSchedule

Description

This operation updates a snapshot schedule configured for a gateway volume.

The default snapshot schedule for stored volumes is once every 24 hours, starting at the creation time of the volume. You can use this API to change the snapshot schedule configured for the volume.

In the request you must identify the gateway volume whose snapshot schedule you want to update, and the schedule information, including when you want the snapshot to begin on a day and the frequency (in hours) of snapshots.

Request

Syntax

```
POST / HTTP/1.1
Host: storagegateway.region.amazonaws.com
Authorization: authorization
Content-Type: application/x-amz-json-1.1
x-amz-date: date
x-amz-target: StorageGateway_20120630.UpdateSnapshotSchedule

{
  "VolumeARN": "String",
  "StartAt": Number,
  "RecurrenceInHours": "Number",
  "Description": "String"
}
```

JSON Fields

Description

Optional description of the snapshot that overwrites the existing description.

Length: Minimum length of 1. Maximum length of 255.

Required: No

Type: String

RecurrenceInHours

Frequency of snapshots. Specify the number of hours between snapshots.

Valid Values: One of the values 1 | 2 | 4 | 8 | 12 | 24.

Required: Yes

Type: Number

StartAt

The hour of the day at which the snapshot schedule begins represented as *hh*, where *hh* is the hour (0 to 23). The hour of the day is in the time zone of the gateway.

Length: 2

Valid Values: An integer between 0 and 23.

Required: Yes

Type: Number

VolumeARN

The Amazon Resource Name (ARN) of the volume. Use the [ListVolumes \(p. 381\)](#) operation to return a list of gateway volumes.

Required: yes

Type: String

Response

Syntax

```
HTTP/1.1 200 OK
x-amzn-RequestId: x-amzn-RequestId
Content-Type: application/x-amz-json-1.1
Content-length: payloadLength
Date: date

{
  "VolumeARN": "String"
}
```

JSON Fields

VolumeARN

The Amazon Resource Name (ARN) of the storage volume whose snapshot schedule was updated. It is the same value you provided in your request.

Type: String

Errors

This operation returns the following error codes in addition to exceptions common to all operations. For information about these errors and common exceptions, see [Error Responses \(p. 287\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- VolumeNotFound

Examples

The following example shows a request that updates a snapshot schedule.

Example Request

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.UpdateSnapshotSchedule
```

```
{
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygate
way/volume/vol-1122AABB",
  "StartAt": 0,
  "RecurrenceInHours": 1,
  "Description": "hourly snapshot"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 12 Sep 2012 12:00:02 GMT
Content-Type: application/x-amz-json-1.1
Content-length: 104

{
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygate
way/volume/vol-1122AABB"
}
```

Related Actions

- [DescribeSnapshotSchedule](#) (p. 360)

Data Types

The AWS Storage Gateway API contains several data types that various actions use. This section describes each data type in detail.

Note

The order of each field in the response is not guaranteed. Applications should not assume a particular order.

The following data types are supported:

- [CachediSCSIVolume](#) (p. 404)
- [ChapInfo](#) (p. 406)
- [Disk](#) (p. 407)
- [GatewayInfo](#) (p. 407)
- [NetworkInterface](#) (p. 408)
- [StorédiSCSIVolume](#) (p. 408)
- [VolumeInfo](#) (p. 409)
- [VolumeiSCSIAttributes](#) (p. 410)
- [VolumeRecoveryPointInfo](#) (p. 410)

CachediSCSIVolume

Describes a cached storage volume.

```
{ "VolumeiSCSIAttributes":  
  { "ChapEnabled": Boolean,  
    "LunNumber": Number,  
    "NetworkInterfaceId": "String",  
    "NetworkInterfacePort": Number,  
    "TargetARN": "String"  
  },  
  "VolumeARN": "String",  
  "VolumeId": "String",  
  "VolumeProgress": Number,  
  "VolumeSizeInBytes": Number,  
  "VolumeStatus": "String",  
  "VolumeType": "String",  
  "SourceSnapshotId": "String"  
}
```

ChapEnabled

Indicates whether mutual CHAP is enabled for the iSCSI target.

Type: String

LunNumber

The logical disk number.

Type: String

NetworkInterfaceId

The network interface ID of the cached volume that initiators use to map the cached volume as an iSCSI target.

Type: String

NetworkInterfacePort

The port used to communicate with iSCSI targets.

Type: Number

SourceSnapshotId

If the cached volume was created from a snapshot, this field contains the snapshot ID used, e.g. snap-1122aabb. Otherwise, this field is not included.

Type: String

TargetARN

The Amazon Resource Name (ARN) of the volume target.

Type: String

VolumeARN

The ARN of the storage volume.

Type: String

VolumeId

The unique identifier of the storage volume, e.g. vol-1122AABB.

Type: String

VolumeiSCSIAttributes

An [VolumeiSCSIAttributes](#) (p. 410) object that represents a collection of iSCSI attributes for one stored volume.

Type: Object

VolumeProgress

The percentage complete if the volume is restoring or bootstrapping that represents the percent of data transferred. This field does not appear in the response if the stored volume is not restoring or bootstrapping.

Type: Number

VolumeSizeInBytes

The size of the volume in bytes that was specified in the [CreateCachediSCSIVolume \(p. 318\)](#) operation.

Type: Number

VolumeStatus

One of the [VolumeStatus \(p. 412\)](#) values that indicates the state of the volume.

Type: String

VolumeType

One of the enumeration values describing the type of volume.

Type: [VolumeType \(p. 412\)](#)

ChapInfo

Describes Challenge-Handshake Authentication Protocol (CHAP) information that supports authentication between your gateway and iSCSI initiators.

```
{
  "InitiatorName": "String",
  "SecretToAuthenticateInitiator": "String",
  "SecretToAuthenticateTarget": "String",
  "TargetARN": "String"
}
```

InitiatorName

The iSCSI initiator that connects to the target.

Length: Minimum length of 1. Maximum length of 255.

Valid Values: The target name can contain lowercase letters, numbers, periods (.), and hyphens (-).

Type: String

SecretToAuthenticateInitiator

The secret key that the initiator (e.g. Windows client) must provide to participate in mutual CHAP with the target.

Length: Minimum length of 12. Maximum length of 16.

Type: String

SecretToAuthenticateTarget

The secret key that the target must provide to participate in mutual CHAP with the initiator (e.g. Windows client).

Length: Minimum length of 12. Maximum length of 16.

Type: String

TargetARN

The ARN of the volume.

Length: Minimum length of 50. Maximum length of 500.

Valid Values: The target name can contain lowercase letters, numbers, periods (.), and hyphens (-).

Type: String

Disk

Describes a gateway local disk.

```
{ "DiskId": "String",  
  "DiskPath": "String",  
  "DiskNode": "String",  
  "DiskSizeInBytes": Number,  
  "DiskAllocationType": "String",  
  "DiskAllocationResource": "String"  
}
```

DiskAllocationResource

The iSCSI Qualified Name (IQN) that is defined for the disk. This field is not included in the response if the local disk is not defined as an iSCSI target. The format of this field is *targetIqn::LUNNumber::region-volumeId*.

Type: String

DiskAllocationType

One of the [DiskAllocationType](#) (p. 411) enumeration values that identifies how the local disk is used.

Type: String

DiskId

The unique device ID or other distinguishing data that identify the local disk.

Type: String

DiskNode

The device node of the local disk as assigned by the virtualization environment.

Type: String

DiskPath

The path of the local disk in the gateway virtual machine (VM).

Type: String

DiskSizeInBytes

The local disk size in bytes.

Type: Number

GatewayInfo

Describes a gateway.

```
{ "GatewayARN": "String"  
}
```

GatewayARN

The ARN of a gateway.

Type: String

NetworkInterface

Describes a gateway's network interface.

```
{
  "Ipv4Address": "String",
  "MacAddress": "String",
  "Ipv6Address": "String"
}
```

Ipv4Address

The Internet Protocol version 4 (IPv4) address of the interface.

Type: String.

Ipv6Address

The Internet Protocol version 6 (IPv6) address of the interface. Currently not supported.

Type: String

MacAddress

The Media Access Control (MAC) address of the interface.

Type: String

StorediSCSIVolume

Describes an iSCSI stored volume.

```
{
  "VolumeARN": "String",
  "VolumeId": "String",
  "VolumeType": "String",
  "VolumeStatus": "String",
  "VolumeSizeInBytes": Number,
  "VolumeProgress": Number,
  "VolumeDiskId": "String",
  "SourceSnapshotId": "String",
  "PreservedExistingData": Boolean,
  "iSCSIAttributes": Array
}
```

VolumeiSCSIAttributes

An [VolumeiSCSIAttributes](#) (p. 410) object that represents a collection of iSCSI attributes for one stored volume.

Type: Object

PreservedExistingData

Indicates if when the stored volume was created, existing data on the underlying local disk was preserved.

Valid Values: true | false

Type: Boolean

SourceSnapshotId

If the stored volume was created from a snapshot, this field contains the snapshot ID used, for example, snap-78e22663. Otherwise, this field is not included.

Type: String

VolumeARN

The ARN of the storage volume.

Length: Minimum length of 50. Maximum length of 500.

Type: String

VolumeDiskId

The disk ID of the local disk that was specified in the [CreateStorediSCSIVolume \(p. 327\)](#) operation.

Type: String

VolumeId

The unique identifier of the volume, for example, vol-AE4B946D.

Type: String

VolumeProgress

Represents the percentage complete if the volume is restoring or bootstrapping that represents the percent of data transferred. This field does not appear in the response if the stored volume is not restoring or bootstrapping.

Type: String

VolumeStatus

One of the [VolumeStatus \(p. 412\)](#) values that indicates the state of the storage volume.

Type: String

VolumeSizeInBytes

The size of the volume in bytes.

Type: Number

VolumeType

One of the [VolumeType \(p. 412\)](#) enumeration values describing the type of the volume.

Type: String

VolumeInfo

Describes a storage volume.

```
{ "VolumeARN" : "String",  
  "VolumeType" : "String"  
}
```

VolumeARN

The ARN for the storage volume, for example, the following is a valid ARN
arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB".

Length: Minimum length of 50. Maximum length of 500.

Type: String

VolumeType

One of the [VolumeType](#) (p. 412) values that indicates the configuration of the storage volume, for example, as a storage volume.

Type: String

VolumeiSCSIAttributes

Lists iSCSI information about a volume.

```
{ "ChapEnabled": Boolean,  
  "LunNumber": Number,  
  "NetworkInterfaceId": "String",  
  "NetworkInterfacePort": Number,  
  "TargetARN": "String"  
}
```

ChapEnabled

Indicates whether mutual CHAP is enabled for the iSCSI target.

Type: Boolean

NetworkInterfaceId

The network interface identifier.

Type: String

NetworkInterfacePort

The port used to communicate with iSCSI targets.

Type: Number

LunNumber

The logical disk number.

Type: Number (positive integer).

TargetARN

The ARN of the volume target.

Length: Minimum length of 50. Maximum length of 800.

Valid Values: The target name can contain lowercase letters, numbers, periods (.), and hyphens (-).

Type: String

VolumeRecoveryPointInfo

Lists information about the recovery points of a cached volume.

```
{ "VolumeARN": "String",  
  "VolumeSizeInBytes": Number,  
  "VolumeUsageInBytes": Number,  
  "VolumeRecoveryPointTime": "String"  
}
```

VolumeARN

The ARN of the volume associated with the recovery point.

Type: String

VolumeSizeInBytes

The size, in bytes, of the volume to which the recovery point is associated.

Type: Number

VolumeUsageInBytes

The size, in bytes, of the volume in use at the time of the recovery point.

Type: Number

VolumeRecoveryPointTime

The time of the recovery point.

Type: String format of a date in the ISO8601 extended YYYY-MM-DD'T'HH:MM:SS'Z' format.

Enumeration Types

The AWS Storage Gateway API contains several enumeration types that various actions use. This section describes each enumeration.

The following enumeration values are supported:

- [BandwidthType](#) (p. 411)
- [DiskAllocationType](#) (p. 411)
- [GatewayState](#) (p. 411)
- [GatewayTimezone](#) (p. 411)
- [GatewayType](#) (p. 412)
- [Regions](#) (p. 412)
- [VolumeStatus](#) (p. 412)
- [VolumeType](#) (p. 412)

BandwidthType

The bandwidth rate limit type.

Valid Values: UPLOAD | DOWNLOAD | ALL

DiskAllocationType

The configuration of a gateway local disk.

Valid Values: AVAILABLE | CACHE STORAGE | STORED iSCSI VOLUME | UPLOAD BUFFER

GatewayState

The state of a gateway.

Valid Values: RUNNING | SHUTDOWN

GatewayTimezone

The time zone for your gateway. The time zone is used, for example, when a time stamp is given to a snapshot.

Valid Values: GMT-12:00 | GMT-11:00 | GMT-10:00 | GMT-9:00 | GMT-8:00 | GMT-7:00 | GMT-6:00 | GMT-5:00 | GMT-4:00 | GMT-3:30 | GMT-3:00 | GMT-2:00 | GMT-1:00 | GMT | GMT+1:00 | GMT+2:00 | GMT+3:00 | GMT+3:30 | GMT+4:00 | GMT+4:30 | GMT+5:00 | GMT+5:30 | GMT+5:45 | GMT+6:00 | GMT+7:00 | GMT+8:00 | GMT+9:00 | GMT+9:30 | GMT+10:00 | GMT+11:00 | GMT+12:00

GatewayType

The type of a gateway.

Valid Values: CACHED | STORED

Regions

The region your gateway is activated in and where your snapshots are stored.

Valid Values: us-east-1 | us-west-1 | us-west-2 | eu-west-1 | ap-northeast-1 | ap-southeast-1 | sa-east-1

VolumeStatus

The status of the storage volume.

Valid Values: AVAILABLE | BOOTSTRAPPING | CREATING | DELETED | IRRECOVERABLE | PASS THROUGH | RESTORING | RESTORE AND PASS THROUGH | UPLOAD BUFFER NOT CONFIGURED

VolumeType

The type of storage volume. Currently only STORED iSCSI is supported.

Valid Values: CACHED iSCSI | STORED iSCSI

Document History for AWS Storage Gateway

This Document History describes the important changes since the last release of the *AWS Storage Gateway User Guide*.

Relevant Dates to this History:

- **Current product version**—2012-06-30
- **Last document update**—May 6, 2013

Change	Description	Release Date
Support for Microsoft Hyper-V	AWS Storage Gateway now provides the ability to deploy an on-premises gateway on the Microsoft Hyper-V virtualization platform. Gateways deployed on Microsoft Hyper-V have all the same functionality and features as the existing on-premises Storage Gateway. To get started deploying a gateway with Microsoft Hyper-V, see Set Up and Activate (Hyper-V Host) (p. 33).	In this release.
Support for deploying a gateway on Amazon EC2	AWS Storage Gateway now provides the ability to deploy a gateway in Amazon Elastic Compute Cloud (Amazon EC2). You can launch a gateway instance in Amazon EC2 using the AWS Storage Gateway AMI available in AWS Marketplace . To get started deploying a gateway using the AWS Storage Gateway AMI, go to Launching and Activating an Amazon EC2 Gateway AMI (p. 139).	15 Jan 2013

Change	Description	Release Date
Support for gateway-cached volumes and introduction of API Version 2012-06-30	<p>In this release, AWS Storage Gateway introduces support for gateway-cached volumes. Gateway-cached volumes minimize the need to scale your on-premises storage infrastructure, while still providing your applications with low-latency access to their active data. You can create storage volumes up to 32 TiB in size and mount them as iSCSI devices from your on-premises application servers. Data written to your gateway-cached volumes is stored in Amazon Simple Storage Service (Amazon S3), with only a cache of recently written and recently read data stored locally on your on-premises storage hardware. Gateway-cached volumes allow you to utilize Amazon S3 for data where higher retrieval latencies are acceptable, such as for older, infrequently accessed data, while maintaining storage on-premises for data where low-latency access is required.</p> <p>In this release, AWS Storage Gateway also introduces a new API version that, in addition to supporting the current operations, provides new operations to support gateway-cached volumes.</p> <p>For more information on the two AWS Storage Gateway solutions, see How AWS Storage Gateway Works (p. 3).</p> <p>You can also try a test setup. For instructions, go to Getting Started with AWS Storage Gateway (p. 7).</p>	29 Oct 2012
API and IAM Support	<p>In this release, AWS Storage Gateway introduces API support as well as support for AWS Identity and Access Management (IAM).</p> <ul style="list-style-type: none"> • API support—You can now programmatically configure and manage your AWS Storage Gateway resources. For more information about the APIs, see API Reference for AWS Storage Gateway (p. 283) in <i>AWS Storage Gateway User Guide</i>. • IAM Support—AWS Identity and Access Management (IAM) enables you create users and manage user access to your AWS Storage Gateway resources by means of IAM policies. For examples of IAM policies, go to Access Control Using AWS Identity and Access Management (IAM) (p. 277). For more information about IAM, go to AWS Identity and Access Management (IAM) detail page. 	09 May 2012
Static IP Support	<p>You can now specify a static IP for your local gateway. For more information, see Configuring Your AWS Storage Gateway to Use a Static IP Address (p. 239).</p>	05 Mar 2012
New Guide	<p>This is the first release of <i>AWS Storage Gateway User Guide</i>.</p>	24 Jan 2012

Appendices for AWS Storage Gateway

This appendix includes the following sections.

Topics

- [Appendix A: The Components in Your vSphere Environment for AWS Storage Gateway \(p. 415\)](#)
- [Appendix B: Configuring a VMware ESXi Host for AWS Storage Gateway \(p. 417\)](#)
- [Appendix C: The Components in Your Hyper-V Environment for AWS Storage Gateway \(p. 421\)](#)
- [Appendix D: Configuring a Microsoft Hyper-V Host for AWS Storage Gateway \(p. 422\)](#)
- [Appendix E: About AWS Storage Gateway \(p. 433\)](#)

Appendix A: The Components in Your vSphere Environment for AWS Storage Gateway

You use VMware to create an on-premises virtual machine that hosts an AWS Storage Gateway. You use a VMware client to interact with a VMware server and create your virtual machines. A gateway virtual machine definition—or template—that contains all the files and data for creating a new gateway is available from the [AWS Storage Gateway Detail Page](#). The template is distributed as a single `.ova` file which is deployed on the VMware server. In this section, the components of the VMware vSphere environment that you need to know to use the AWS Storage Gateway service are discussed.

The following table describes the subset of vSphere components that you typically work with when using the AWS Storage Gateway service.

Component	Description
VMware vSphere	The VMware virtualization platform for managing its virtual computing infrastructure including the client and server.

AWS Storage Gateway User Guide
Appendix A: The Components in Your vSphere
Environment

Component	Description
VMware ESXi hypervisor OS (vSphere Server)	The VMware server OS that hosts the gateway virtual machine. You interact with the OS through the vSphere client GUI. To provision an AWS Storage Gateway, you only need to access the host during the activation of the gateway. For all other management and maintenance-related functions, you use the AWS Management Console.
VMware vSphere Client (vSphere Client)	The VMware software that you use on your computer to access and manage your VMware environment. You manage your virtual machine (that contains the gateway) using the client.
VMware High Availability	VMware High Availability (HA) is a component of vSphere that can provide protection from failures in your infrastructure layer supporting a gateway VM. VMware HA does this by using multiple hosts configured as a cluster so that if one host running a gateway VM fails, the gateway VM can be restarted automatically on another host within the cluster. AWS Storage Gateway can be used with VMware HA. For more information about VMware HA, go to VMware HA: Concepts and Best Practices . For more information about using VM HA with AWS Storage Gateway, see Using AWS Storage Gateway with VMware High Availability (p. 92).
Virtual machine	The software implementation of a computer that contains the components of AWS Storage Gateway. The virtual machine (VM) runs on the VMware vSphere platform.
OVA, OVF	A template that represents a customized virtual machine. The AWS Storage Gateway appliance is an Open Virtualization Format (OVF) package that is distributed in an Open Virtualization Application (OVA). The OVA template contains all the information needed to configure and start a gateway. You deploy the template using the client connected to a VMware server. For instructions about downloading the OVA template for AWS Storage Gateway, go to AWS Storage Gateway Detail Page .
Datastore	The storage on the vSphere server where the files that define a virtual machine are stored. These files come from the OVA file provided as part of the service. When you deploy the OVA, you select a datastore on which to store the file if there is more than one datastore for the VMware server.

Appendix B: Configuring a VMware ESXi Host for AWS Storage Gateway

This section provides basic information for you to set up your virtualization host. Following the basic setup, we also cover some optional host configuration.

The AWS Storage Gateway service includes an on-premises software appliance that communicates with AWS's cloud storage infrastructure. The appliance is packaged as a virtual machine that you deploy on a host running the VMware ESX/ESXi virtualization software. For more information on the VMware virtualization software, go to [VMware vSphere Hypervisor](#). For requirements that your VMware environment must meet to run AWS Storage Gateway, see [Requirements \(p. 6\)](#).

To install the VMware vSphere hypervisor OS on your host

1. Insert the VMware vSphere hypervisor disk in the disk drive.
2. Restart the computer.

Depending on your computer bios settings, the computer might automatically boot off your disk. If not, check the relevant settings to boot the computer from the hypervisor disk.

3. Follow the instructions on the monitor to install the VMware hypervisor OS.

This installation wipes any existing content on the disk and installs the hypervisor.

Tip

After a successful VMware hypervisor host installation, the monitor displays the IP address of the host computer. Note down this IP address. You use the IP address to connect to the host.

4. Set the time on the host.

For instructions, see [Synchronize VM Time with Host Time \(p. 15\)](#).

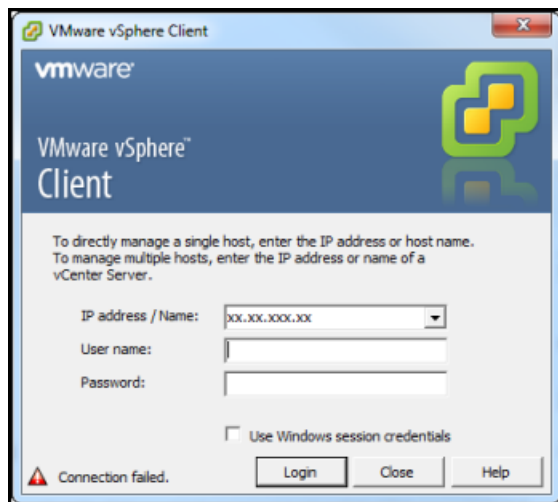
In the preceding steps, you provisioned a host with VMware hypervisor. The hypervisor is aware of host computer configuration, such as available processors, memory, and local hard disks. The host provides these resources to the AWS Storage Gateway.

You can optionally configure this host by adding more storage, such as additional direct-attached disks or SAN disks. The following steps illustrate how you can add one or more SAN disks to this host.

To connect to the hypervisor host

1. Start the VMware vSphere client and connect to the host using the host IP address.

The VMware vSphere Client dialog box appears.



2. Enter the IP address of the host in the **IP Address** field.
3. Enter the credentials in the **User Name** and **Password** fields.
4. Click **Login**.

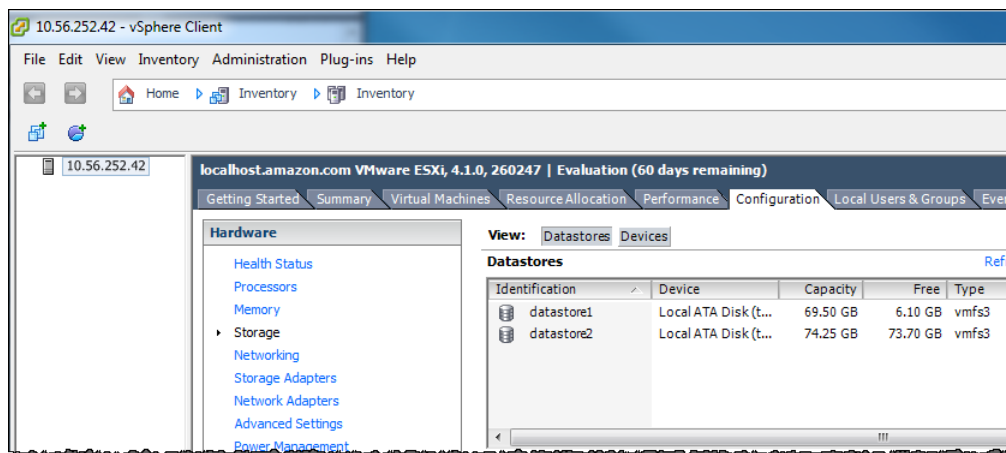
This connects your client to the host. You are now ready to configure the host.

To add a new iSCSI target


1. After you have connected to your remote device through the hypervisor, go to the **Configuration** tab of the host and click **Storage** in the **Hardware** list.

The **Datastores** pane shows the available data stores.

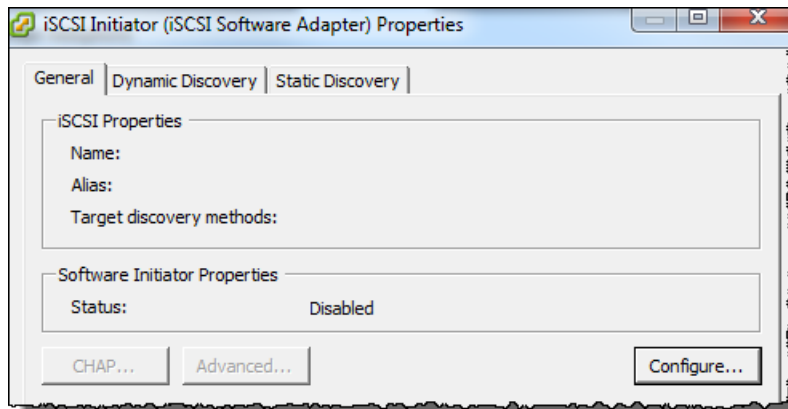
For example, the following example shows that the host has two local hard drives, datastore1, and datastore2 available.




2. In the **Hardware** list, click **Storage Adapters**.
3. In the **Storage Adapters** pane, select **iSCSI Software Adapter**, and then click the **Properties** link in the **Details** pane.




4. In the **iSCSI Initiator (iSCSI Software Adapter) Properties** dialog box, click **Configure**.



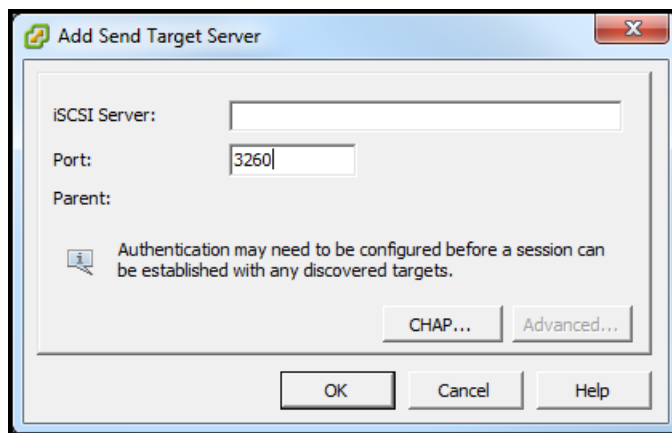
5. In the **General Properties** dialog box, select **Enabled** to set the software initiator status to enabled and click **OK**.



6. In the **iSCSI Initiator (iSCSI Software Adapter) Properties** dialog box, select the **Dynamic Discovery** tab, and click **Add** to add an iSCSI target.



7. In the **Add Send Target Server** dialog box, enter a name in the **iSCSI Server** field and a port in the **Port** field and click **OK**.



Enter the IP address or DNS name of the storage system.

The new iSCSI server location that is entered here appears in the **Sends Target** list on the **Dynamic Discovery** tab.

8. Click **Close** to close the **iSCSI Initiator (iSCSI Software Adapter) Properties** dialog box.

At this time, you have added a new iSCSI target in the host configuration.

Appendix C: The Components in Your Hyper-V Environment for AWS Storage Gateway

You use Microsoft Hyper-V to create an on-premises virtual machine that hosts an AWS Storage Gateway. You use the Hyper-V Manager to interact with a Hyper-V server and create your virtual machines. A gateway virtual machine definition—or template—that contains all the files and data for creating a new gateway is available from the AWS Storage Gateway console. The template is distributed as a single .zip file which you import into the Hyper-V server. In this section, the components of the Microsoft Hyper-V environment that you need to know to use the AWS Storage Gateway service are discussed.

The following table describes the subset of Hyper-V components that you typically work with when using the AWS Storage Gateway service.

Component	Description
Microsoft Hyper-V	The Microsoft virtualization platform for managing a virtual computing infrastructure including the client and server.
Hyper-V hypervisor OS	The Hyper-V server OS that hosts the gateway virtual machine. You interact with the OS through the Microsoft Hyper-V Manager GUI. To provision an AWS Storage Gateway you need to only access the host during the activation of the gateway. For all other management and maintenance-related functions, you use the AWS Management Console.
Hyper-V Manager	The Hyper-V client software that you use on your computer to access and manage your Hyper-V environment. You manage your virtual machine (that contains the gateway) using the client.
Virtual machine	The software implementation of a computer that contains the components of AWS Storage Gateway. The virtual machine (VM) runs on the Microsoft Hyper-V platform.
Import files (packaging of VM)	<p>The AWS Storage Gateway appliance is distributed as a compressed directory containing the following:</p> <ul style="list-style-type: none">• <code>Snapshots</code> folder, which will be empty for the AWS Storage Gateway .• <code>Virtual Hard Disks</code> folder, which contains one virtual hard disk file called <code>AWS-Storage-Gateway.vhd</code>.• <code>Virtual Machines</code> folder, which contains an exported configuration files <code>GUID.exp</code>, where <code>GUID</code> is the virtual machine ID.• <code>config.xml</code>, which contains configuration information used for importing. <p>You deploy AWS Storage Gateway to Hyper-V by first uncompressing the directory and then importing the uncompressed folder using the Hyper-V Manager. .</p>

Appendix D: Configuring a Microsoft Hyper-V Host for AWS Storage Gateway

This appendix provides basic information for you to set up, configure, and troubleshoot your [Microsoft Hyper-V 2008 R2](#) virtualization host including:

- Setting Up and Configuring
 - [Installing Microsoft Hyper-V](#) (p. 422)
 - [Connecting to Microsoft Hyper-V Host](#) (p. 424)
 - [Configuring Virtual Network Settings](#) (p. 425)
 - [Configuring a Share on a Microsoft Hyper-V Host](#) (p. 427)
 - [Adding a Virtual Disk Backed by a Hard Disk](#) (p. 428)
- Troubleshooting
 - [Troubleshooting Your Microsoft Hyper-V Setup](#) (p. 429)

Setting Up and Configuring a Microsoft Hyper-V Host

The AWS Storage Gateway service includes an on-premises software appliance that communicates with AWS's cloud storage infrastructure. The appliance is packaged as a virtual machine that you can deploy on a host running Microsoft Hyper-V virtualization software. For more information on the Microsoft Hyper-V software, go to [Microsoft Server Virtualization](#). For requirements that your Hyper-V environment must meet to run AWS Storage Gateway, see [Requirements](#) (p. 6).

Installing Microsoft Hyper-V

The section describe a procedure for installing Microsoft Hyper-V. If you already have a Microsoft Hyper-V virtualization environment or your environment will be set up by an administrator familiar with the platform, then you do not need to understand these steps in detail.

Refer to the [Hyper-V Getting Started Guide](#) on the *Microsoft TechNet* site for more information about the installation of Hyper-V.

To install the Microsoft Hyper-V hypervisor OS on your host

1. Insert the Microsoft Hyper-V disk in the disk drive.
2. Restart the computer.

Depending on your computer's BIOS settings, the computer might automatically boot off your disk. If not, check the relevant settings to boot the computer from the hypervisor disk.

3. Follow the instructions on the monitor to install the Hyper-V hypervisor OS.

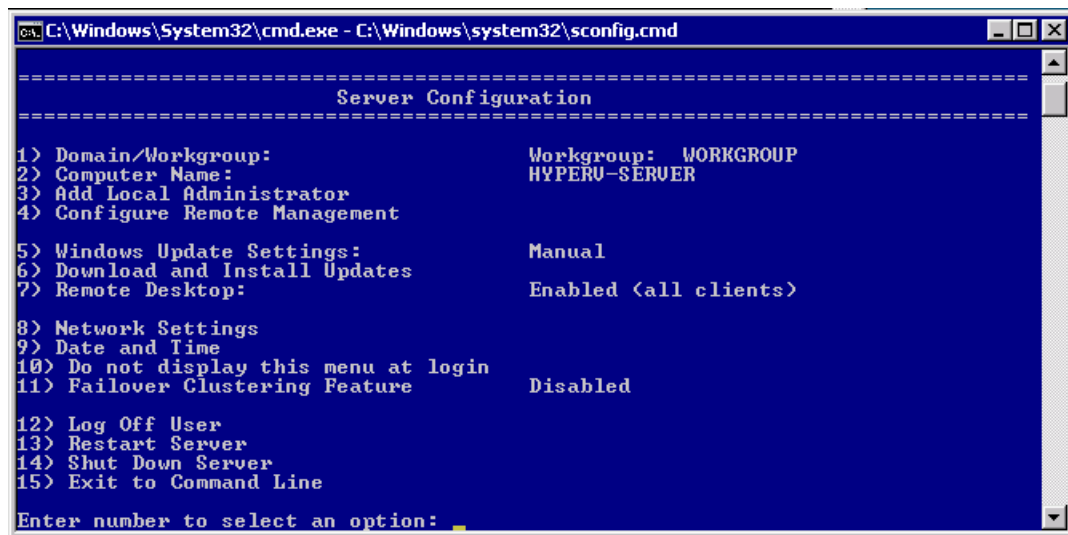
This installation wipes any existing content on the disk and installs the hypervisor.

After a successful Hyper-V hypervisor host installation, you will be prompted to create an Administrator account password. After creating this account, the monitor displays a **Server Configuration** menu where you will do further configuration of the host.

4. In the **Server Configuration** menu, configure the host. We recommend the following:

To...	Do This...
Configure remote management.	Select option 4 and then enable the following: <ul style="list-style-type: none"> option 1, Allow MMC Remote Management option 2, Enable Windows PowerShell option 3, Allow Server Manager Remote Management
Find the network address of the host.	Select option 8 and follow the prompts. Note the IP address for use later.
Set the date and time.	Select option 9 and follow the prompts.
(Optional) Change the computer name.	Select option 2 and follow the prompts. Since this requires a reboot, you may want to make this configuration change last.
(Optional) Add local administrators.	Select option 3 and follow the prompts.
(Optional) Enable remote desktop.	Select option 7 and follow the prompts.

The following example shows a Server Configuration menu.



- (Optional) You may need to put the IP address of the hypervisor host in your hosts file of client computers that connect to the hypervisor host.

For example, in Windows 7 and 8, the hosts file can be found at this location:

```
%SystemRoot%\system32\drivers\etc\hosts
```

Connecting to Microsoft Hyper-V Host

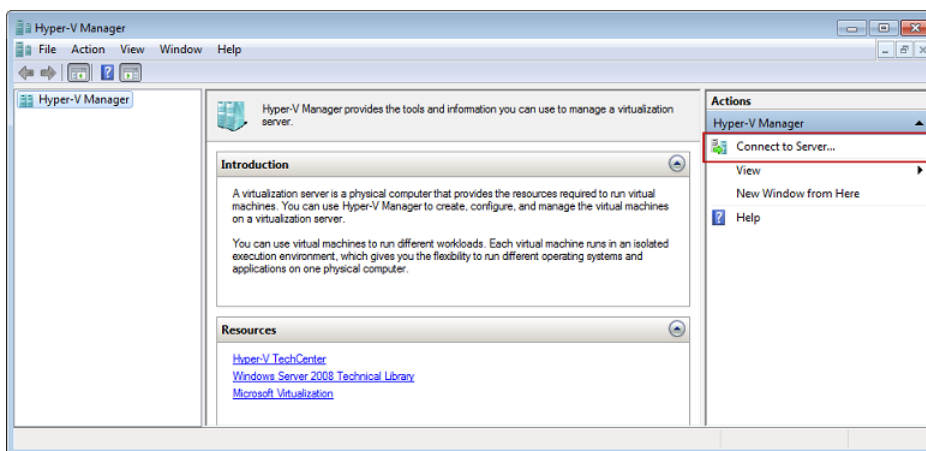
The Hyper-V Manager runs on your client computer and connects to the hypervisor host. You use the Microsoft Hyper-V Manager to import, configure, and start the AWS Storage Gateway VM.

To connect to the hypervisor host

1. Start the Microsoft Hyper-V Manager (virtmgmt.msc).

Note

The Hyper-V Manager is a feature that you enable for your client computer. For more information about enabling it, go to [Install and Configure Hyper-V Tools for Remote Administration](#).



2. In the **Actions** pane, select **Connect to Server...**
3. In the **Select Computer** dialog box, select **Another computer** and either type the IP address of the hypervisor host or the hostname and click **OK**.

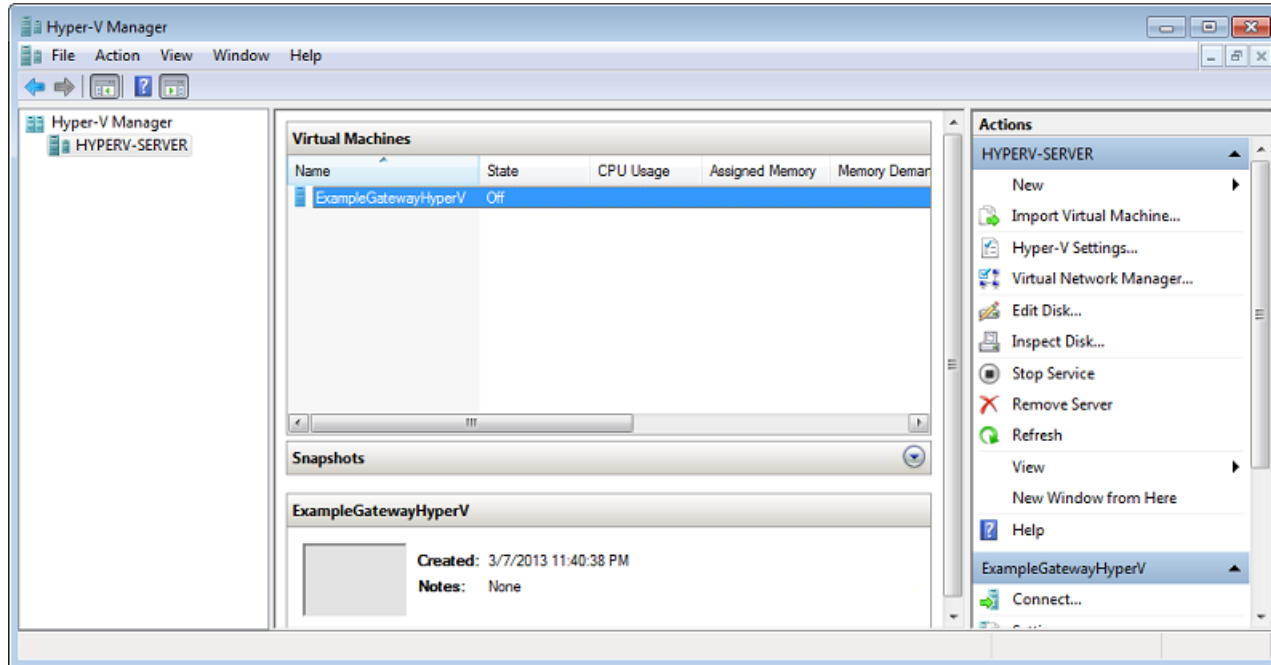
Note

To connect to a hypervisor host using the hostname, you may need to make an entry in your `hosts` file so that the hostname can be mapped to the correct IP address.

Note

If you have not been added to the local administrators for the hypervisor host, you may be prompted for credentials.

The following example shows Hyper-V Manager connected to a hypervisor host called `HYPERV-SERVER` with one gateway VM.

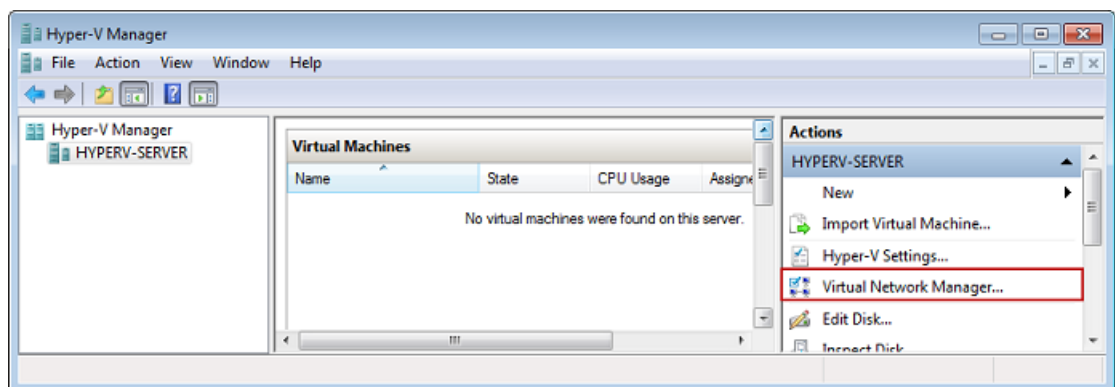


Configuring Virtual Network Settings


After you install and configure a Microsoft Hyper-V host, we recommend that you set up virtual networks by creating a new virtual network and associating it with a network interface of the host. Later, when you configure your gateway VM, you must associate it with one or more virtual networks so that the VM has connectivity.

To configure virtual network settings for the your Hyper-V host

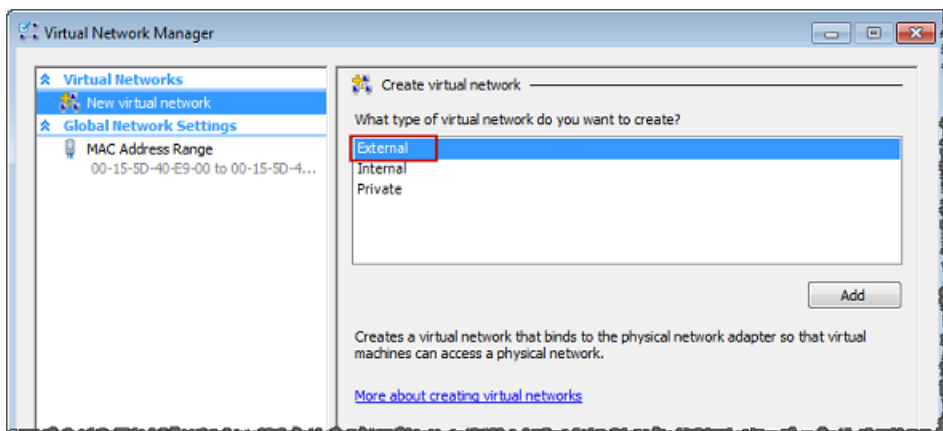
1. Start the Microsoft Hyper-V Manager (virtmgmt.msc).
2. In the hypervisor host list (left pane), select your hypervisor.
3. In the **Actions** menu, under the hypervisor host name (e.g., `HYPERV-SERVER`), click **Virtual Network Manager**.



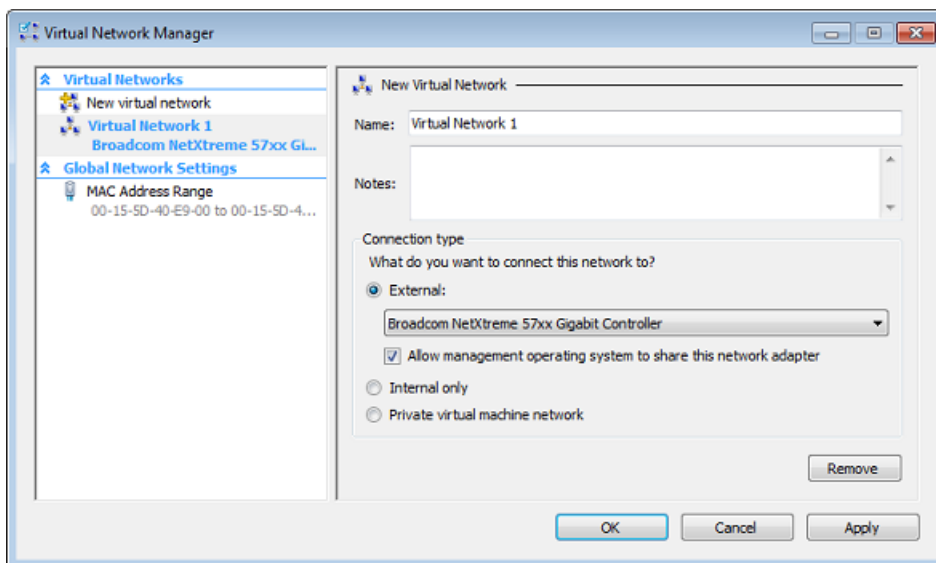
4. In the **Virtual Network Manager** dialog box, select **New virtual network**.



5. Select **External** as the virtual network type and click **Add**.



6. Provide a name for the network, and click **OK**.



When you configure your gateway virtual machine, you can use this virtual network.

Configuring a Share on a Microsoft Hyper-V Host

When you deploy AWS Storage Gateway to a Microsoft Hyper-V hypervisor you must copy over the gateway source files to the hypervisor so that you can import them into the hypervisor. The import process imports only from the local disk of the host. Depending on your virtualization environment and who needs to copy the source gateway files to the host, you might find it easier to create a share on the hypervisor host so that it can be mapped by client computers. This section describes how to create a share.

To configure a share on a Microsoft Hyper-V host

1. Access the host's **Server Configuration** menu directly by using the host's console or using Remote Desktop Connection to connect to the host.

You should see the **Server Configuration** menu as shown in [Installing Microsoft Hyper-V \(p. 422\)](#).

2. Create a share.
 - a. In a command window, enter the following firewall commands to allow file sharing:

```
netsh advfirewall firewall add rule name="File Sharing" dir=in action=allow protocol=TCP localport=445
netsh advfirewall firewall add rule name="File Sharing" dir=in action=allow protocol=TCP localport=139
```

- b. Create a directory for the share.

In the following example command, we create a share folder at C:\Users\Administrator\Share. Your location and drive letter may be different.

```
mkdir C:\Users\Administrator\Share
```

- c. Create the share.

In the following example command, the share name is *sharename* and *user* is a user authorized to access the hypervisor host.

```
net share sharename=C:\Users\Administrator\Share /grant:user,FULL
```

- d. Confirm the share was created by listing the shares.

The following command lists the shares defined. *sharename* should be included in this list.

```
net share
```

3. Test the share from a client computer by mapping the share.


```
\\hypervisor\sharename
```

Adding a Virtual Disk Backed by a Hard Disk

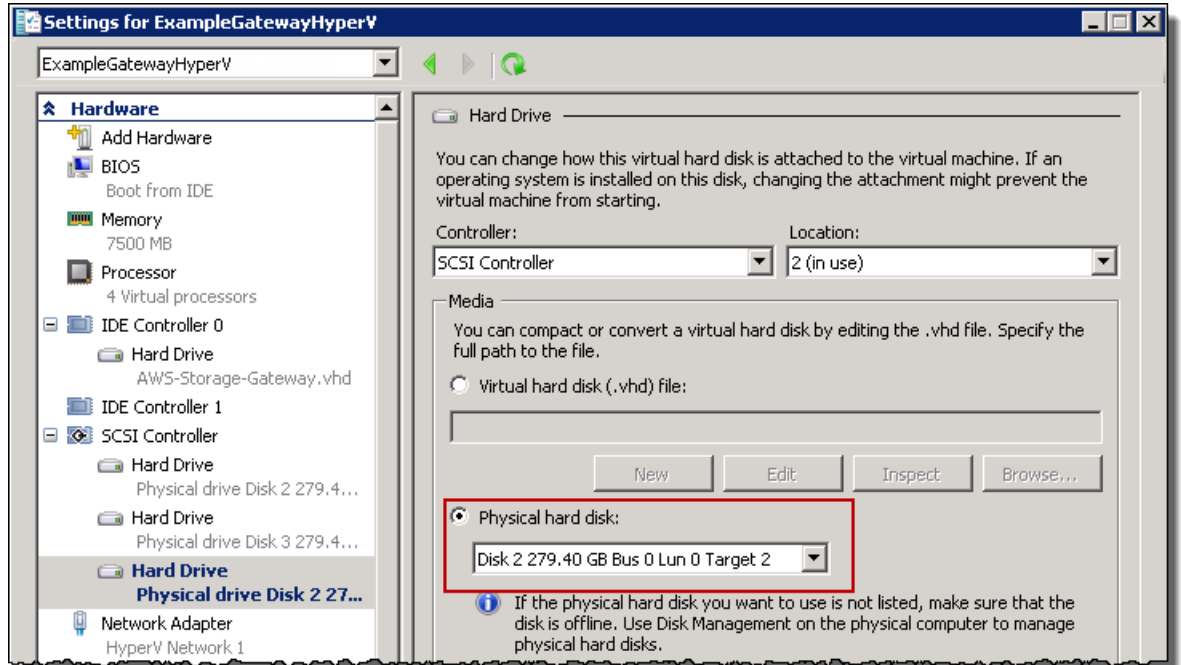
In a preceding section, you provisioned a host with Hyper-V hypervisor. The hypervisor is aware of host computer configuration, such as available processors, memory, and local hard disks. The host provides these resources to AWS Storage Gateway. You can optionally configure this host by adding more storage, such as additional direct-attached disks or SAN disks. In this section, we show you how to add a virtual disk backed by a direct-attached disk.

To add a virtual disk backed by a physical hard disk

1. Start the Microsoft Hyper-V Manager (virtmgmt.msc).
2. Select the VM.
3. In the **Actions** list for the VM, click **Settings....**
4. In the **Hardware** list, click **SCSI Controller**.
5. Select **Hard Drive** in the **SCSI Controller** pane and click **Add**.



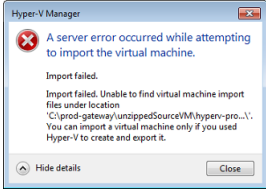
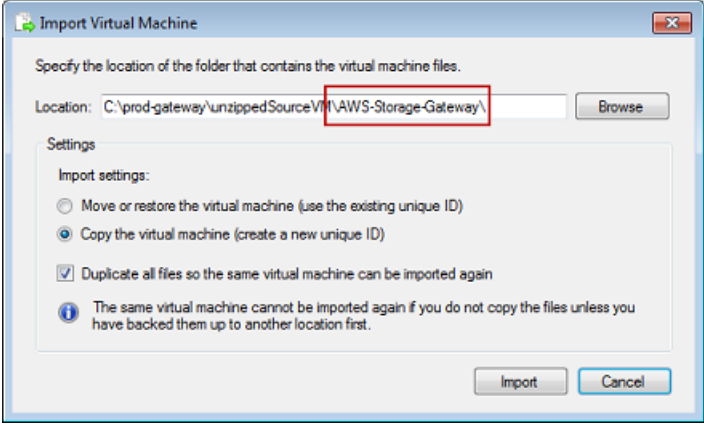
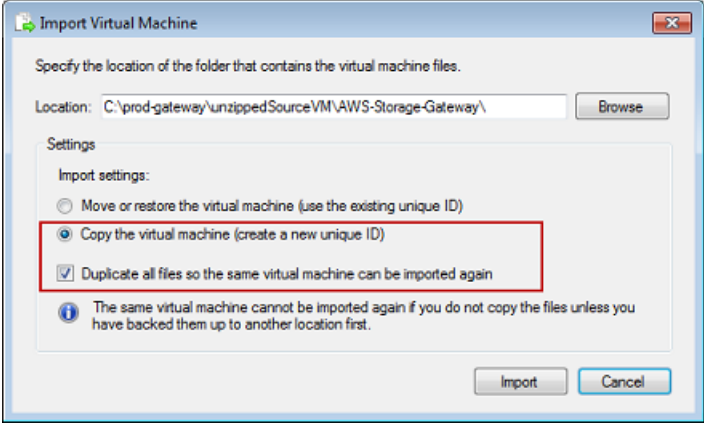
6. In the **Hard Drive** pane, select **Physical hard disk**.

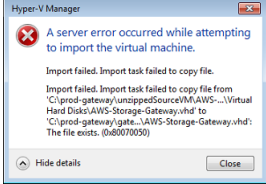
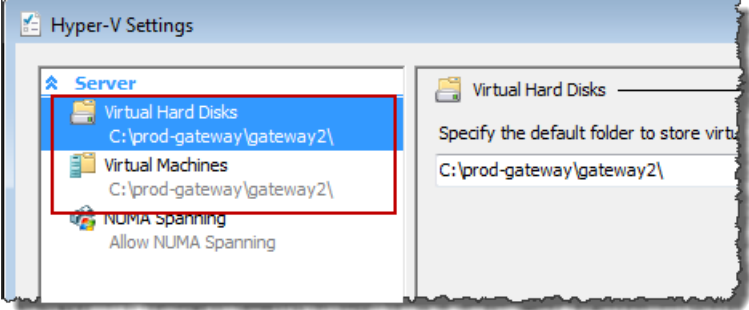
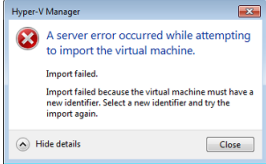
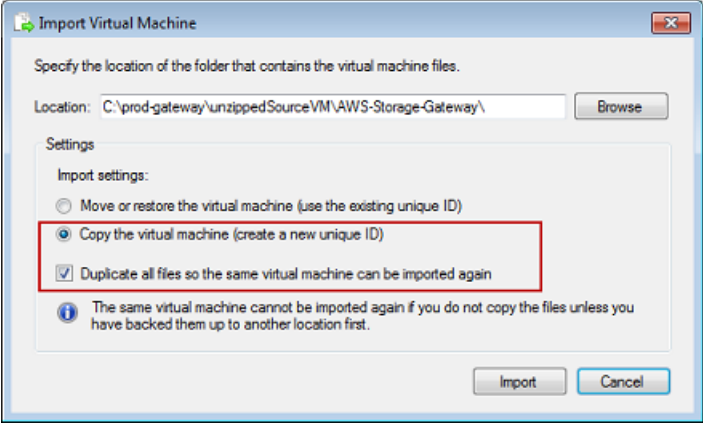


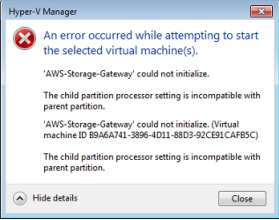
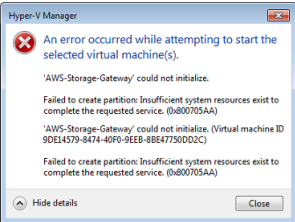
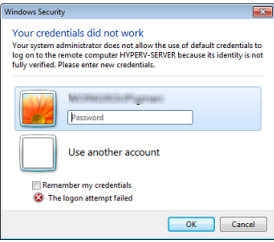
7. Click **OK**.

Troubleshooting Your Microsoft Hyper-V Setup

The following table lists typical issues that you might encounter when deploying AWS Storage Gateway on the Microsoft Hyper-V platform.

Issue	Action to Take
<p>You try to import a gateway and receive the error message: "Import failed. Unable to find virtual machine import file under location ...".</p>  <p>The screenshot shows an error dialog box titled "Hyper-V Manager" with a red 'X' icon. The text reads: "A server error occurred while attempting to import the virtual machine. Import failed. Import failed. Unable to find virtual machine import files under location 'C:\prod-gateway\unzippedSourceVM\hyper-pro...'. You can import a virtual machine only if you used Hyper-V to create and export it." There are "Hide details" and "Close" buttons at the bottom.</p>	<p>This error can occur for the following reasons:</p> <ul style="list-style-type: none"> • If you are not pointing to the root of the unzipped gateway source files. The last part of the location you specify in the Import Virtual Machine dialog box should be <code>AWS-Storage-Gateway\</code>, as the following example shows:  <p>The screenshot shows the "Import Virtual Machine" dialog box. The "Location" field contains the path "C:\prod-gateway\unzippedSourceVM\AWS-Storage-Gateway\". A red box highlights the "AWS-Storage-Gateway\" part of the path. The "Settings" section has "Copy the virtual machine (create a new unique ID)" selected, and "Duplicate all files so the same virtual machine can be imported again" checked. There are "Import" and "Cancel" buttons at the bottom.</p> <ul style="list-style-type: none"> • If you have already deployed a gateway and you did not select the Copy the virtual machine option and check the Duplicate all files... option in the Import Virtual Machine dialog box, then the VM was created in the location where you have the unzipped gateway files and you can not import from this location again. To fix this problem, get a fresh copy of the unzipped gateway source files and copy to a new location. Use the new location as the source of the import. The following example shows the options that you must check if you plan on creating multiple gateways from one unzipped source files location.  <p>The screenshot shows the "Import Virtual Machine" dialog box with the same path as above. A red box highlights the "Copy the virtual machine (create a new unique ID)" radio button and the "Duplicate all files so the same virtual machine can be imported again" checked checkbox. There are "Import" and "Cancel" buttons at the bottom.</p> <p>For more information about deploying the gateway, see Download and Deploy the AWS Storage Gateway VM on Your Host (p. 34) in the Getting Started exercise for Microsoft Hyper-V.</p>

Issue	Action to Take
<p>You try to import a gateway and receive the error message: "Import failed. Import task failed to copy file."</p> 	<p>If you have already deployed a gateway and you try to reuse the default folders that store the virtual hard disk files and virtual machine configuration files, then this error will occur. To fix this problem, specify new locations in the Hyper-V Settings dialog box.</p>  <p>For more information about deploying the gateway, see To import the VM (p. 37) in the Getting Started exercise for Microsoft Hyper-V.</p>
<p>You try to import a gateway and receive an error message: "Import failed. Import failed because the virtual machine must have a new identifier. Select a new identifier and try the import again."</p> 	<p>When you import the gateway make sure you select the Copy the virtual machine option and check the Duplicate all files... option in the Import Virtual Machine dialog box to create a new unique ID for the VM. The following example shows the options in the Import Virtual Machine dialog box that you should use.</p>  <p>For more information about importing the gateway, see To import the VM (p. 37) in the Getting Started exercise for Microsoft Hyper-V.</p>

Issue	Action to Take
<p>You try to start a gateway VM and receive an error message "The child partition processor setting is incompatible with parent partition."</p> 	<p>This error is likely caused by a CPU discrepancy between the required CPUs for the gateway and the available CPUs on the host. Ensure that the VM CPU count is supported by the underlying hypervisor. For more information about the requirements for AWS Storage Gateway, see Requirements (p. 6).</p>
<p>You try to start a gateway VM and receive an error message "Failed to create partition: Insufficient resources exist to complete the requested service."</p> 	<p>This error is likely caused by a RAM discrepancy between the required RAM for the gateway and the available RAM on the host. For more information about the requirements for AWS Storage Gateway, see Requirements (p. 6).</p>
<p>Your snapshots and gateway software updates are occurring at slightly different times than expected.</p>	<p>The gateway VM's clock may be offset from the actual time, known as clock drift. Check and correct the VM's time using local gateway console's time synchronization option. For more information, see Synchronizing Your Gateway VM Time (p. 243).</p>
<p>You need to put the unzipped Microsoft Hyper-V AWS Storage gateway files on the host file system.</p>	<p>Access the host as you would a typical Microsoft Windows server. For example, if the hypervisor host is name <code>hyperv-server</code>, then you can use the following UNC path <code>\\hyperv-server\c\$</code>, which assumes that the name <code>hyperv-server</code> can be resolved or is defined in your local hosts file. You can also create a share on the host (see Configuring a Share on a Microsoft Hyper-V Host (p. 427)).</p>
<p>You are prompted for credentials when connecting to hypervisor.</p> 	<p>Add your user credentials as a local administrator for the hypervisor host by using the <code>Sconfig.cmd</code> tool. For more information, see Setting Up and Configuring a Microsoft Hyper-V Host (p. 422).</p>

Appendix E: About AWS Storage Gateway

The source code for certain open source software components that are included with the AWS Storage Gateway software is available for download at:

- <https://s3.amazonaws.com/aws-storage-gateway-terms/sources.tar> for gateways deployed on VMware ESXi
- https://s3.amazonaws.com/aws-storage-gateway-terms/sources_hyperv.tar for gateways deployed on Microsoft Hyper-V

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).

The packages comprising the AWS Storage Gateway VM are tracked and monitored for security vulnerabilities. When updates are issued, they are applied to each gateway and the updated packages will increment their version number although the major version number of the Linux distribution may not increment. For more information about managing updates, see [Managing Gateway Updates Using the AWS Storage Gateway Console](#) (p. 226).