
AWS Storage Gateway

User Guide

API Version 2012-04-30



AWS Storage Gateway: User Guide

Copyright © 2012 Amazon Web Services LLC or its affiliates. All rights reserved.

The following are trademarks or registered trademarks of Amazon: Amazon, Amazon.com, Amazon.com Design, Amazon DevPay, Amazon EC2, Amazon Web Services Design, AWS, CloudFront, EC2, Elastic Compute Cloud, Kindle, and Mechanical Turk. 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
Service Highlights	2
How AWS Storage Gateway Works	3
Requirements	4
Pricing	5
Getting Started Tutorial	6
Step 1: Sign Up	7
Step 2: Set Up and Activate Gateway	7
Provision Host	8
Download and Deploy the VM	9
Provision Local Disk Storage	18
Configure the VM to Use Paravirtualized Disk Controllers	22
Activate Gateway	23
Step 3: Create Storage Volumes	27
Step 4: Configure Working Storage for the VM	29
Step 5: Access Your Volumes	32
Step 6: Test the Setup	39
Step 7: Sizing Your Work Storage	43
Where Do I Go from Here?	44
Setting Up AWS Storage Gateway	45
Downloading and Deploying the VM	46
Provisioning Local Disk Storage	47
Adding Local Disks for Your Application Data	47
Adding Local Disks for the Gateway Working Storage	51
Configure VM to Use Paravirtualization	56
Activating a Gateway	57
Managing Your Activated Gateway	61
Managing Storage Volumes	61
Storage Volume Status	67
Configuring Working Storage	70
Ongoing Management of Working Storage for a Gateway	73
Managing Your Application Access to Storage Volumes	76
Connecting from a Windows Client to Your Storage Volume	77
Connecting from a Red Hat Client to Your Storage Volume	78
Configuring CHAP Authentication for Your Storage Volume	80
Working With Snapshots	88
Finding a Snapshot	90
Editing a Snapshot Schedule	91
Creating an Adhoc Snapshot	92
Deleting a Snapshot	93
Restoring a Snapshot	94
Restoring a Snapshot to an AWS Storage Gateway Volume	95
Restoring a Snapshot to an Amazon EBS Volume	97
Performing Maintenance Tasks	97
Shutting Down and Turning On a Gateway	99
Managing Gateway Updates	100
Deleting a Gateway	101
Logging into Your Gateway Local Console	102
Routing Your Gateway Through a Proxy	104
Configuring Your Gateway to Use a Static IP Addresses	105
Testing Your Gateway Connectivity to the Internet	108
Configuring Your Gateway for Multiple Network Adapters (NICs)	109
Creating a Storage Volume on a Gateway with Multiple Network Adapters	112
Troubleshooting	113
Optimizing Gateway Performance	117
Monitoring Your AWS Storage Gateway	118
Using the Amazon CloudWatch Console	119
Measuring Performance Between Your Application and Gateway	120

Measuring Performance Between Your Gateway and AWS	122
Monitoring Working Storage	124
Understanding AWS Storage Gateway Metrics	128
Access Control	131
API Reference	137
Required Request Headers	137
Signing Requests	139
Error Responses	141
Operations in AWS Storage Gateway	153
ActivateGateway	154
AddWorkingStorage	156
CreateSnapshot	159
CreateStorediSCSIVolume	162
DeleteBandwidthRateLimit	166
DeleteChapCredentials	168
DeleteGateway	170
DeleteVolume	172
DescribeBandwidthRateLimit	175
DescribeChapCredentials	177
DescribeGatewayInformation	180
DescribeMaintenanceStartTime	183
DescribeSnapshotSchedule	185
DescribeStorediSCSIVolumes	188
DescribeWorkingStorage	192
ListGateways	194
ListLocalDisks	197
ListVolumes	200
ShutdownGateway	203
StartGateway	206
UpdateBandwidthRateLimit	208
UpdateChapCredentials	210
UpdateGatewayInformation	213
UpdateGatewaySoftwareNow	216
UpdateMaintenanceStartTime	218
UpdateSnapshotSchedule	220
Data Types	223
ChapInfo	223
Disk	224
GatewayInfo	225
NetworkInterface	225
StorediSCSIVolume	226
VolumeInfo	227
VolumeiSCSIAttributes	227
Enumeration Types	228
BandwidthType	228
DiskAllocationType	228
GatewayState	228
GatewayTimezone	228
Regions	229
VolumeStatus	229
VolumeType	229
Document History	230
Appendices	232
Appendix A: The Components in Your vSphere Environment	232
Appendix B: Configuring a VMware ESXi Host	234
Appendix C: About AWS Storage Gateway	238

What Is AWS Storage Gateway?

Topics

- [AWS Storage Gateway API \(p. 2\)](#)
- [Service Highlights of AWS Storage Gateway \(p. 2\)](#)
- [How AWS Storage Gateway Works \(p. 3\)](#)
- [Requirements \(p. 4\)](#)
- [Pricing \(p. 5\)](#)

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 an organization's on-premises IT environment and AWS's storage infrastructure. The service enables you to securely upload data to the AWS cloud for cost effective backup and rapid disaster recovery.

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

- **Disaster Recovery**—AWS Storage Gateway enables you to run your applications on-premises while transparently backing up data off-site to Amazon S3 in the form of Amazon EBS snapshots. Using Amazon EC2, you can configure virtual machine images of your application servers in AWS, and only pay for these servers when you need them. In the event 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 compute instances, and you have your environment 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, for new projects, or as a more cost-effective way to run your normal workloads, you can use the AWS Storage Gateway to mirror your on-premises data to Amazon EC2 instances.
- **Off-site Backup**—AWS Storage Gateway enables your existing on-premises applications to store data backups off-site on Amazon S3's scalable, reliable, secure, and cost-effective storage service. All data is securely transferred to AWS over SSL and stored encrypted in Amazon S3 using AES 256-bit encryption. The AWS Storage Gateway provides an attractive alternative to the traditional choice of either maintaining costly hardware in multiple datacenters, or dealing with the longer recovery times and operational burden of managing off-site tape storage.

For more information and pricing, go to the [AWS Storage Gateway 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, deployment overview, and the requirements for deploying the AWS Storage Gateway virtual machine (VM).
- **Getting Started Tutorial for AWS Storage Gateway (p. 6)**—The Getting Started section provides you with step-by-step 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 tutorial, you'll probably want to learn more about how to use AWS Storage Gateway. The following sections covers the fundamentals of setting up, managing, troubleshooting, and monitoring your gateway.

- **Setting Up AWS Storage Gateway (p. 45)**—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. 61)**—After you deploy and activate your gateway, this section provides you with information about how to manage your AWS Storage 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 the default and an adhoc snapshot, how to find information about a snapshot, and how to schedule a snapshot. For more information, see [Working With Snapshots in the AWS Storage Gateway Console \(p. 88\)](#). 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 Elastic Block Store (EBS) volume and attach that to an Amazon EC2 instance. For more information, see [Restoring a Snapshot Using the AWS Storage Gateway Console \(p. 94\)](#).

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 view metrics for throughput, latency, and a number of input/output operations. For more information, see [Monitoring Your AWS Storage Gateway \(p. 118\)](#).

AWS Storage Gateway API

All the preceding sections use 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 \(p. 137\)](#). 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](#).

Service Highlights of AWS Storage Gateway

AWS Storage Gateway is a service that connects a software appliance (gateway) on your premises with cloud-based AWS storage providing seamless integration between your organization's on-premise IT

environment and AWS's storage infrastructure. It enables you to securely upload data to the AWS cloud for cost effective backup and rapid disaster recovery.

AWS Storage Gateway provides low-latency performance by maintaining data on your on-premises storage hardware while asynchronously uploading this data to AWS, where it is encrypted and securely stored in the Amazon Simple Storage Service (Amazon S3).

AWS Storage Gateway enables you to back up point-in-time snapshots to the AWS cloud. These snapshots are stored in Amazon S3 as Amazon Elastic Block Store (Amazon EBS) snapshots. Whenever you need it, you can restore these backups and access them locally or restore the backup snapshots to Amazon EBS volumes and attach these volumes to your Amazon Elastic Compute Cloud (Amazon EC2) instances.

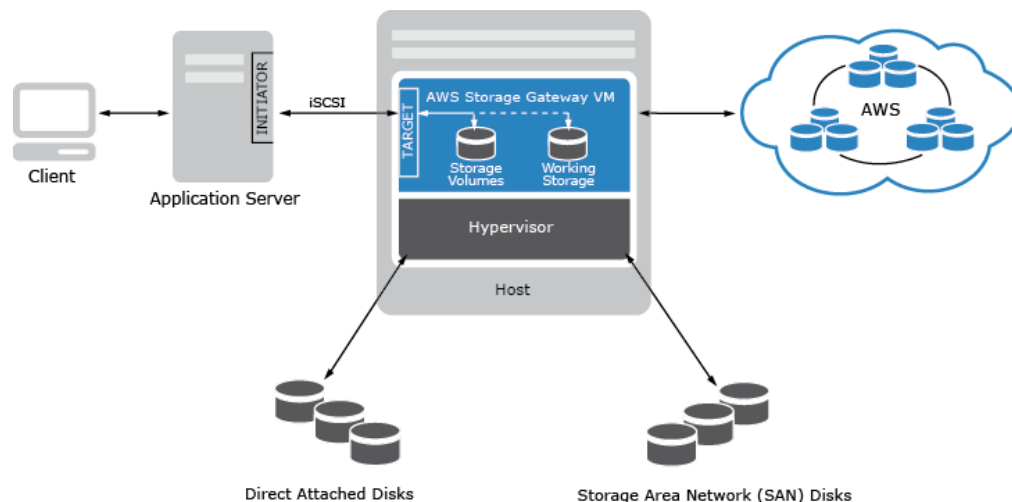
The AWS Storage Gateway supports the industry standard iSCSI interface that works with your existing applications.

AWS Storage Gateway provides the following major advantages:

- **Secure**—AWS Storage Gateway securely transfers your data to AWS over Secure Sockets Layer (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**—AWS Storage Gateway durably stores your on-premises application data by uploading it to Amazon S3. Amazon S3 is designed to sustain the concurrent loss of data in two facilities, redundantly storing your data on multiple devices across multiple facilities in a Region. Your data is stored as an Amazon EBS snapshot, providing a point-in-time backup that can be restored on-premises or used to instantiate new Amazon EBS volumes.
- **Compatible**—There is no need to re-architect your on-premises applications. AWS Storage Gateway exposes a standard iSCSI interface that works with your existing applications, while maintaining data on-premises so that applications retain the same read/write performance.
- **Cost-Effective**—By making it easy for your on-premises applications to store data on Amazon S3, the AWS Storage Gateway reduces the cost, maintenance, and scaling challenges associated with managing primary and backup storage environments. You pay only for what you use with no long-term commitments.
- **Designed for use with other Amazon Web Services**—AWS Storage Gateway allows you to easily mirror data from your on-premises applications to applications running on Amazon EC2 in the event that you require additional on-demand compute capacity for data analysis or replacement capacity for disaster recovery purposes. AWS Storage Gateway is designed to seamlessly integrate with Amazon S3, Amazon EBS, and Amazon EC2 by storing your on-premises application data in Amazon S3 as Amazon EBS snapshots. You can easily access this data from Amazon EC2 by restoring these data snapshots to Amazon EBS volumes and attaching them to your Amazon EC2 instances.
- **Optimized for Network Efficiency**—AWS Storage Gateway efficiently uses your internet bandwidth to speed up the backup of your on-premises application data to AWS. The AWS Storage Gateway only uploads data that has changed, minimizing the amount of data that is sent over the internet. You can also use AWS Direct Connect to further increase the throughput and reduce your network costs by establishing a dedicated network connection between your on-premises gateway and AWS.

How AWS Storage Gateway Works

The following diagram provides an overview of the AWS Storage Gateway deployment.



Once you've installed AWS Storage Gateway's software appliance on a host in your datacenter, 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 and buffer data for upload to Amazon S3, your gateway also stores incoming data in a staging area, referred to as Working Storage. You can use on-premises DAS or SAN disks for Working Storage. Your gateway uploads this Working Storage data 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.

When taking a snapshot, AWS Storage Gateway uploads any changed data from your gateway's Working Storage to Amazon S3 and then forms an Amazon EBS snapshot from this data. Snapshots are incremental backups that can be initiated on a scheduled or ad-hoc basis. When taking a new snapshot, only the data that has changed since your last snapshot is stored. If you have a volume with 100 GBs of data, but only 5 GBs of data have changed since your last snapshot, only the additional 5 GBs of snapshot data will be stored back to Amazon S3. 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, or use the snapshot as a starting point for a new Amazon EBS volume which you can then attach to an Amazon EC2 instance.

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

- ports 80 and 443 are used by the vSphere client to communicate to the ESXi 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

Requirements

The AWS Storage Gateway runs as a virtual machine (VM) that you deploy on a host in your datacenter. The host must be running VMware ESXi Hypervisor (v 4.1 or v 5). A free version is available on the [VMware website](#).

Once deployed, the VM will have the following configuration:

- 4 virtual processors assigned to the VM
- 7.5 GBs of RAM assigned to the VM
- 75 GBs of disk space for .ova installation and system data

You must make sure that your host provides the required hardware for the VM footprint. You will 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

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

Note

The beta version of AWS Storage Gateway currently supports applications that have an average block write size that is greater than 4KiB. For workloads that write blocks that are less than an average size of 4KiB, there is a risk that your gateway could run out of system storage space.

Pricing

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

Getting Started Tutorial for AWS Storage Gateway

Topics

- [Getting Started Requirements for AWS Storage Gateway \(p. 7\)](#)
- [Getting Started Video for AWS Storage Gateway \(p. 7\)](#)
- [Step 1: Sign Up for AWS Storage Gateway \(p. 7\)](#)
- [Step 2: Set Up and Activate AWS Storage Gateway \(p. 7\)](#)
- [Step 3: Create Storage Volumes Using the AWS Storage Gateway Console \(p. 27\)](#)
- [Step 4: Configure Working Storage for the AWS Storage Gateway VM \(p. 29\)](#)
- [Step 5: Access Your AWS Storage Gateway Volumes \(p. 32\)](#)
- [Step 6: Test the Setup - Take a Snapshot and Restore it to Another AWS Storage Gateway Volume \(p. 39\)](#)
- [Step 7: Sizing Your Working Storage for Real-World Workloads \(p. 43\)](#)
- [Where Do I Go from Here? \(p. 44\)](#)

The Getting Started section provides you with step-by-step instructions to set 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 your local storage volume, showing how AWS Storage Gateway enables you to recover your data. Below is an outline of the steps you will follow to set up your gateway and configure a connection from your on-premises client.

At the end of the getting started tutorial, you will have a working gateway with the following sample configuration:

- AWS Storage Gateway VM is deployed on your VMware ESXi hypervisor host.
- Three virtual disks on the VM in the following configuration:
 - Two disks are configured as storage volumes to store your application data.
 - One disk is configured as working storage for use by the gateway.
- Your Windows client is 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 in this getting started section, 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. 4).

The Getting Started tutorial 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 Static IP Addresses](#) (p. 105).

Getting Started Video for AWS Storage Gateway

Before you begin with the Getting Started 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 skip this step.

To sign up for AWS Storage Gateway

1. Go to <http://aws.amazon.com>, and then click **Sign Up Now**.
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.

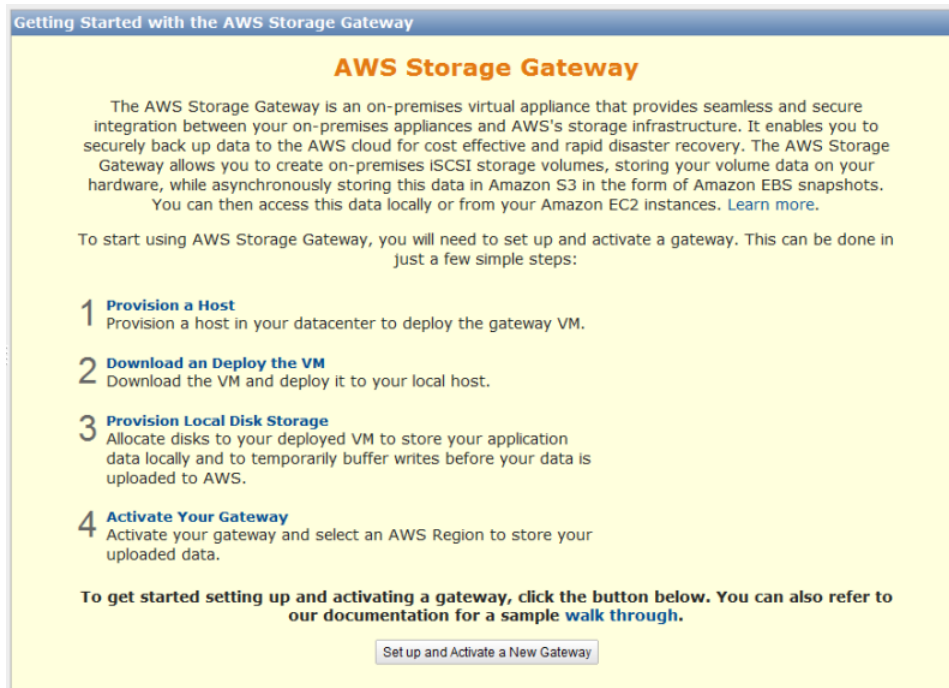
Step 2: Set Up and Activate AWS Storage Gateway

Topics

- [Provision Host to Deploy the AWS Storage Gateway VM](#) (p. 8)
- [Download and Deploy the AWS Storage Gateway VM on Your Host](#) (p. 9)

- [Provision Local Disk Storage for AWS Storage Gateway VM \(p. 18\)](#)
- [Configure the AWS Storage Gateway to Use Paravirtualized Disk Controllers \(p. 22\)](#)
- [Activate AWS Storage Gateway \(p. 23\)](#)

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 you have not activated a gateway yet, 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.

Provision Host to Deploy the AWS Storage Gateway VM

In this procedure, you create a host in your datacenter 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 \(p. 4\)](#).
2. Set up a host in your data center with the VMware ESXi hypervisor.

The 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 \(p. 234\)](#).

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 \(p. 46\)](#). In this

Getting Started tutorial, you deploy your AWS Storage Gateway virtual machine (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. In the following steps, you download this .ova file locally on your client computer and deploy it to the host.

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

Note

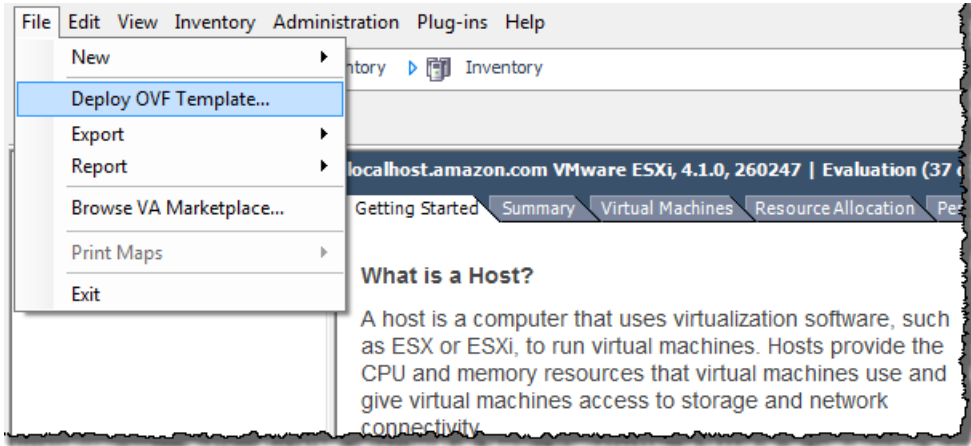
The .zip file is over 500 MB in size and it may take some time to download, depending on your network connection.

Deploy 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**.

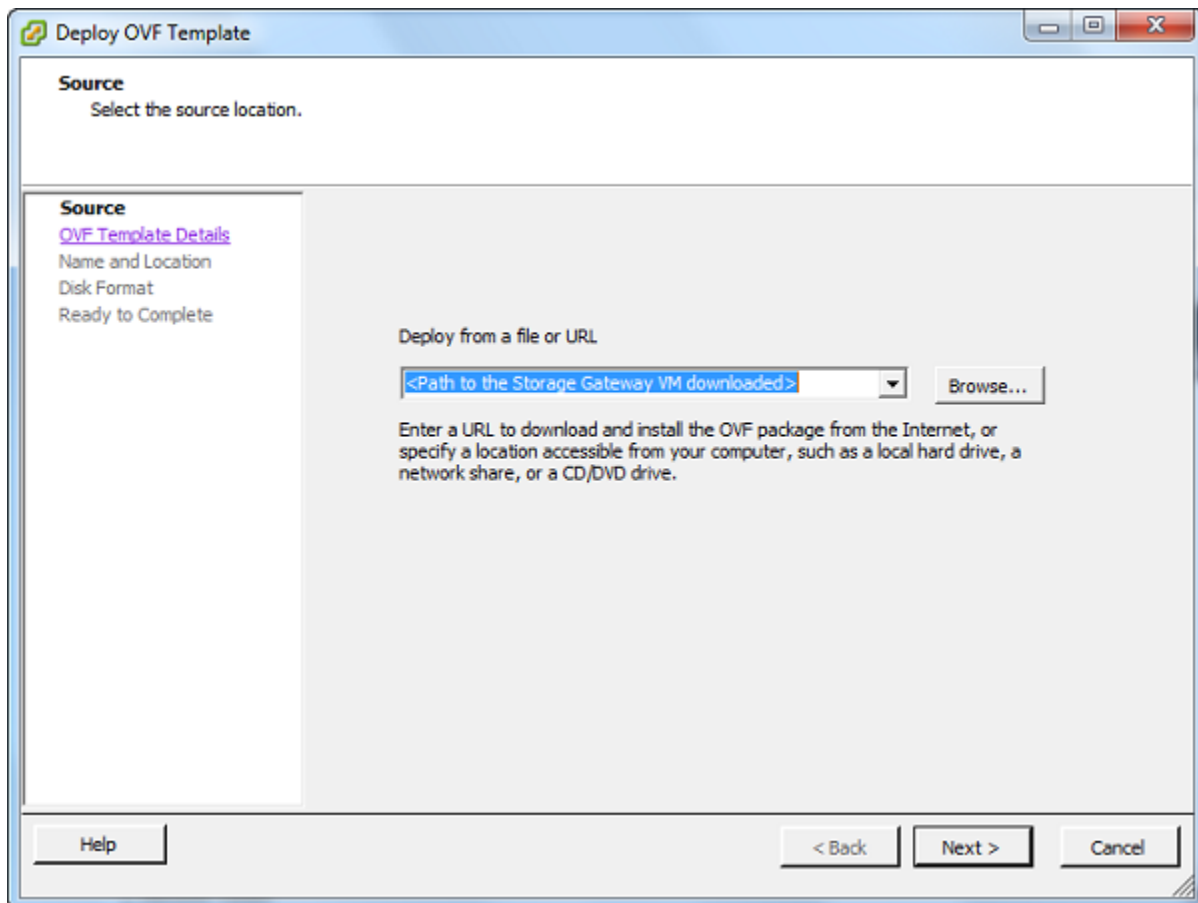
Your vSphere client is now connected to your host computer.

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



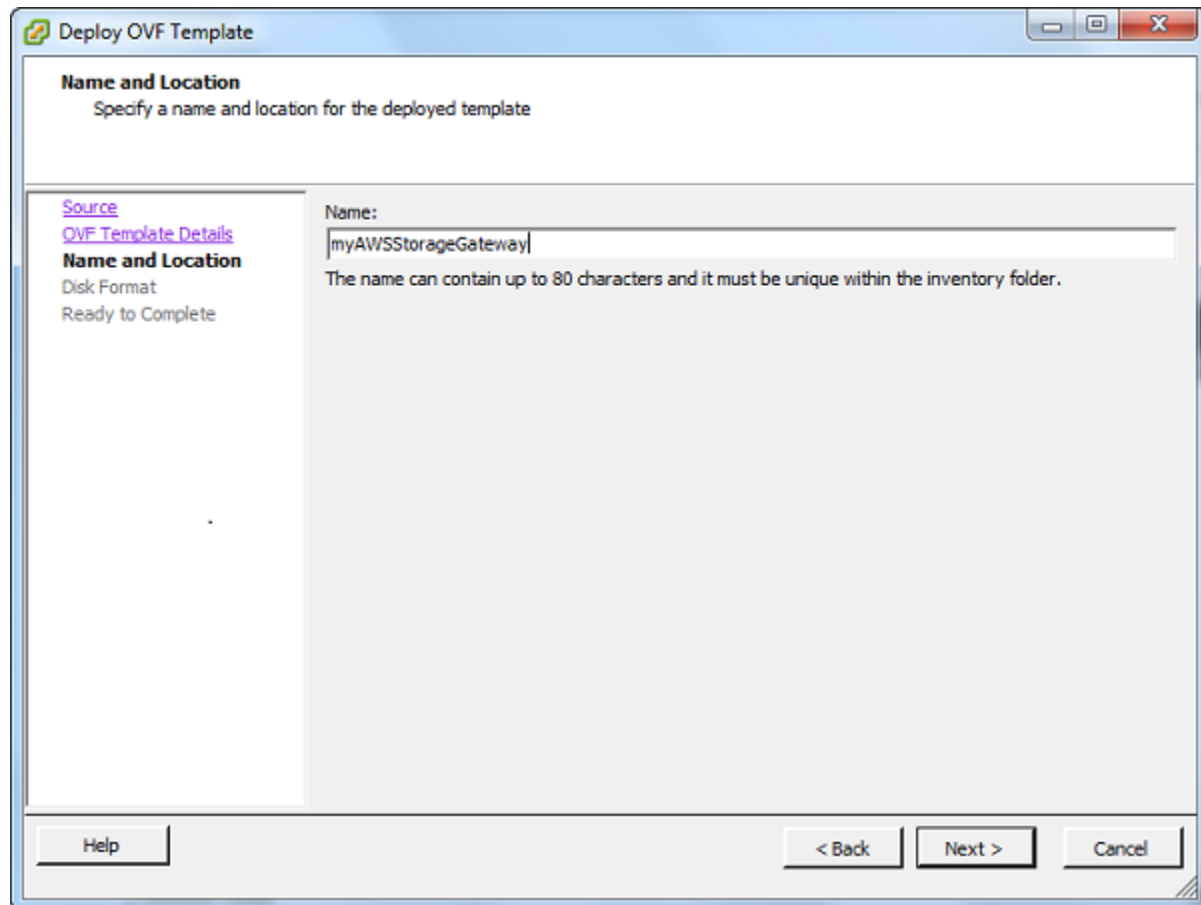
This opens the **Deploy OVF Template** wizard. The wizard is a series of steps for you to provide the required information to deploy the VM.

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



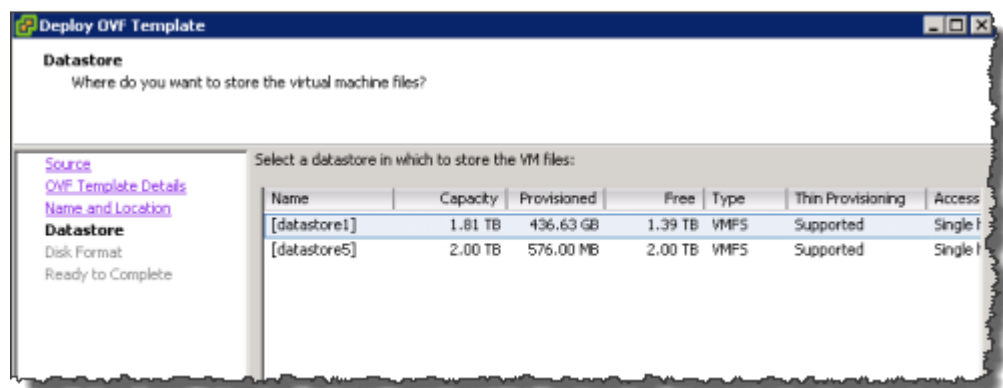
- c. In the **OVF Template Details** pane, click **Next**.
- d. In the **Name and Location** pane, enter the VM name in the **Name** field, and then click **Next**.

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



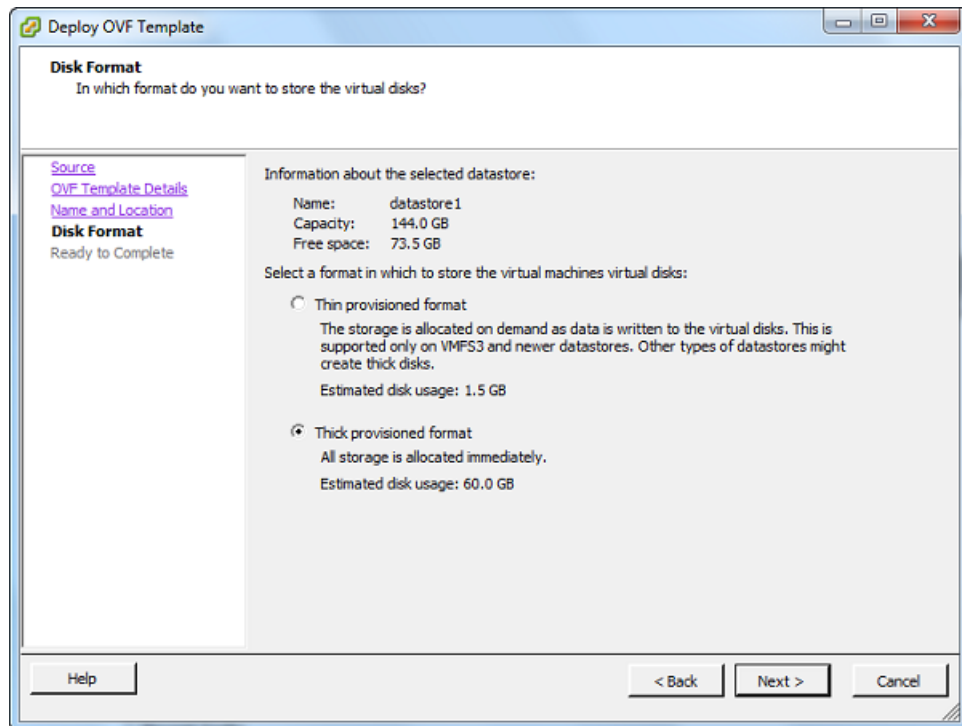
- e. The following **Datastore** pane is displayed only if your host has multiple data stores. In this pane, you select a data store where you want to deploy the VM and click **Next**. Skip to the next step if your host has only one datastore.

Datstore refers to a physical disk on your host. The following example shows a host that has two disks, datastore1 and datastore2.



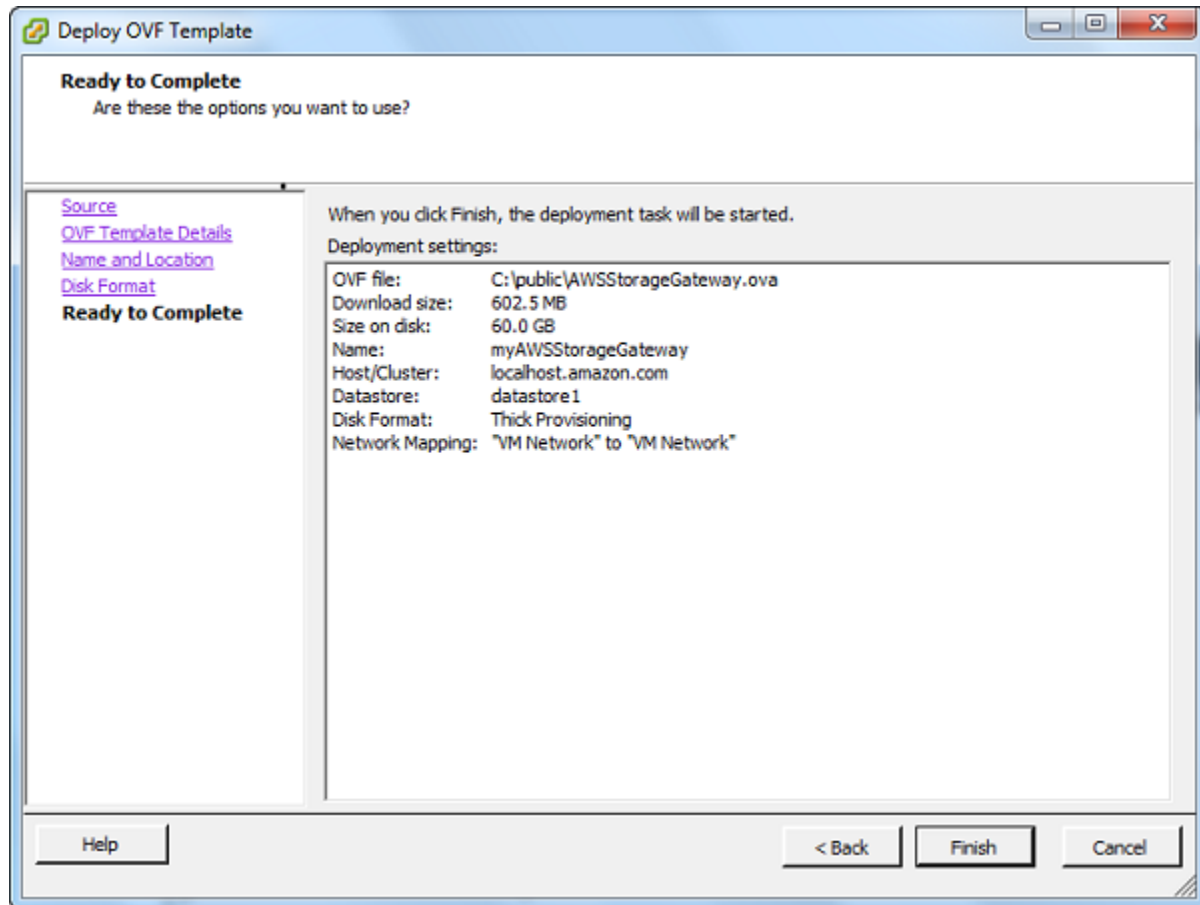
- f. In the **Disk Format** pane, select **Thick provisioned format** and click **Next**.

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

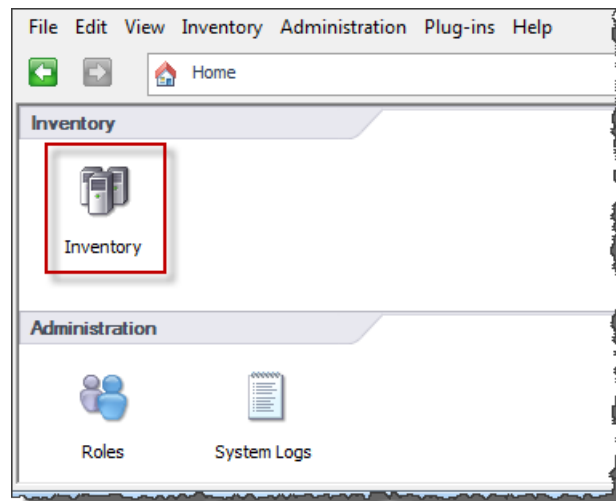


- g. In the **Ready to Complete** pane, click **Finish**.

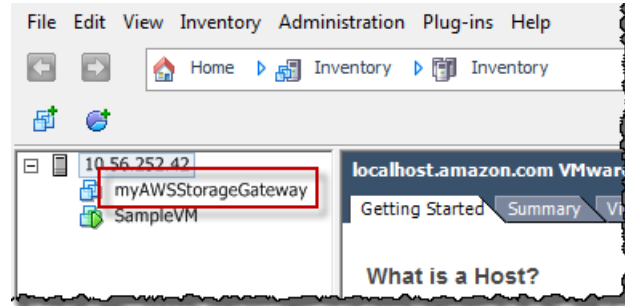
The AWS Storage Gateway VM starts deploying to your host.



- 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.



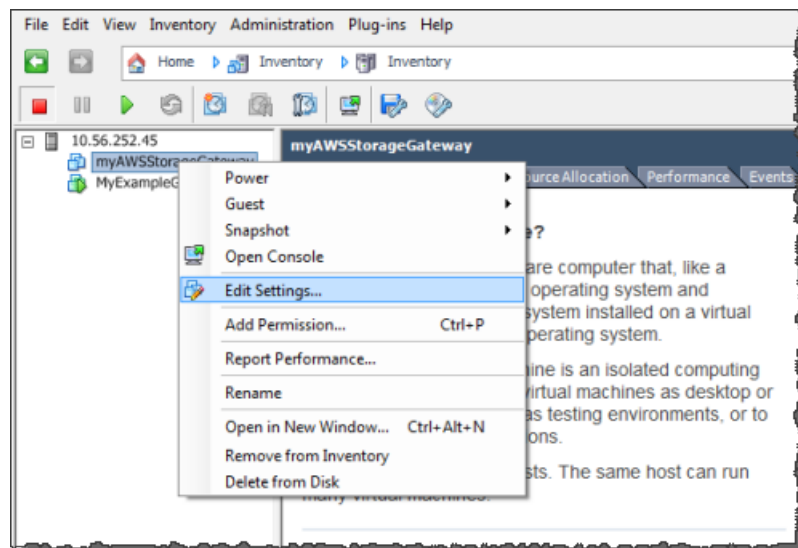
Synchronize VM Time with the Host Time

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. Then, you 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 the host time

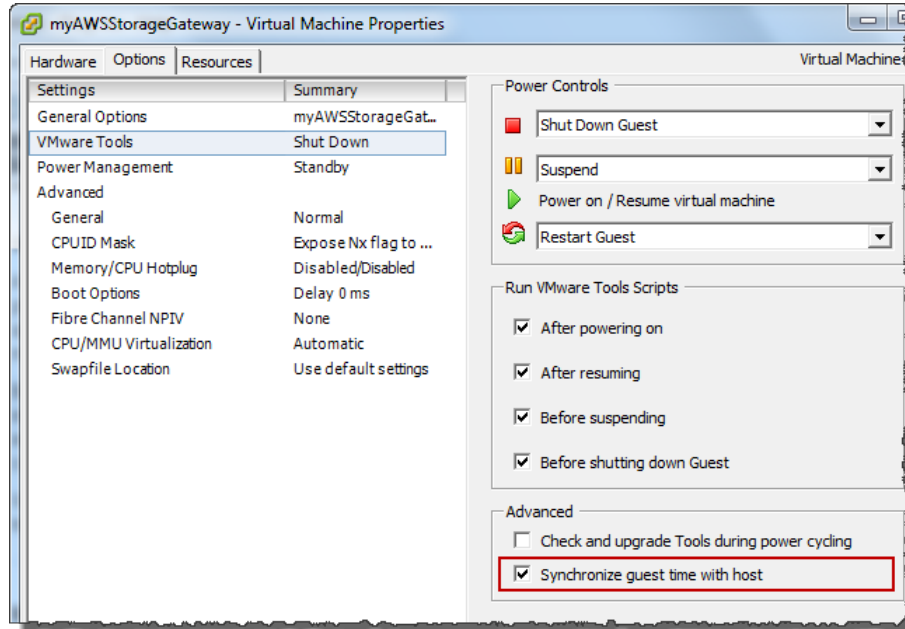
1. Configure your VM time.
 - a. In the vSphere client, right-click the name of your gateway VM and select **Edit Settings...**

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.




2. Configure the host time.

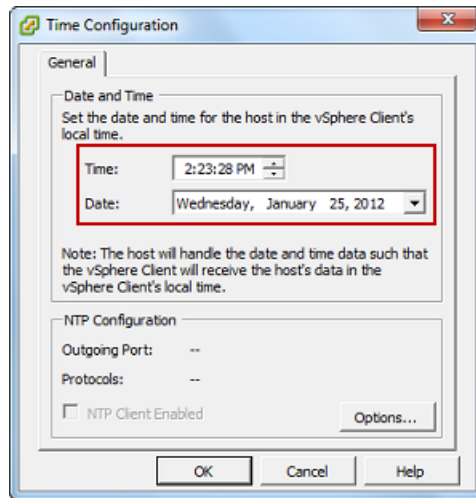
It is important to make sure that your host clock is set to the correct time. If you have not configured your host clock, use the following steps to set the host clock and synchronize it with an NTP server.

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


The **Time Configuration** dialog box displays.



- d. In the **Time Configuration** dialog box, 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:
- In the **Time Configuration** dialog box, click **Options**.
 - In the **NTP Daemon (ntpd) Options** dialog box, select **NTP Settings** in the left pane.



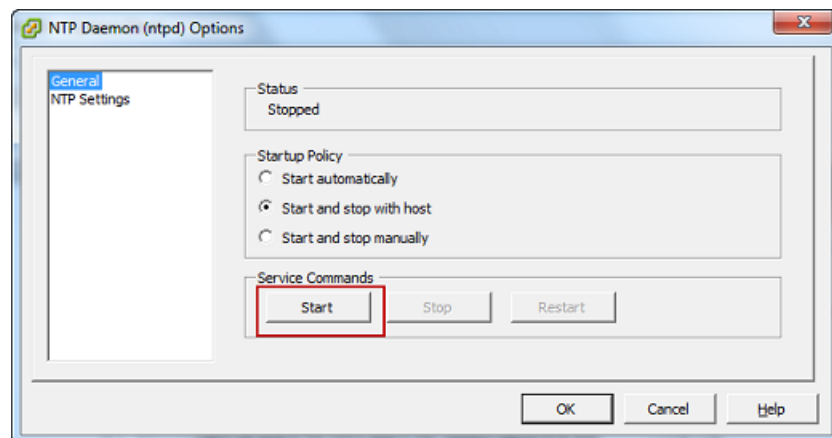
- iii. Click **Add** to add a new NTP server.
- iv. In the **Add NTP Server** dialog box, enter the IP address or the fully qualified domain name of an NTP server and click **OK**.

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.
- g. Click **OK** to close the **Time Configuration** dialog.

Provision Local Disk Storage for AWS Storage Gateway VM

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

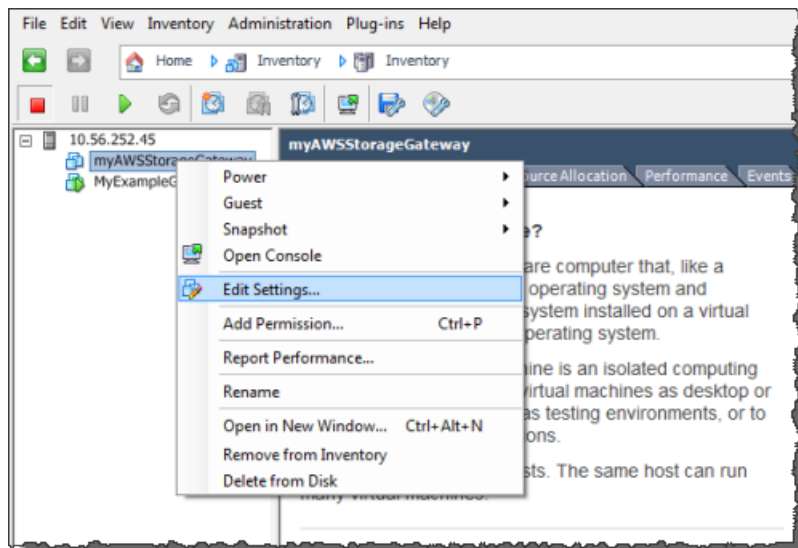
Allocate Local Disk for Application Data

All your application data is maintained locally. You must allocate a disk on the VM to store your application data. This section provides step-by-step 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.

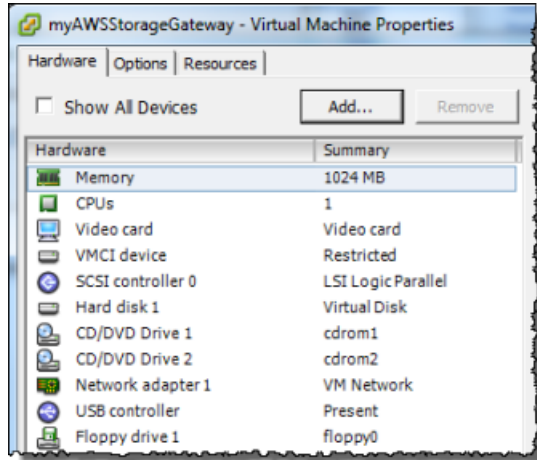
For this getting started exercise, you allocate 1 GB virtual disk to the VM for storing application data.

To allocate a local disk to store your application data

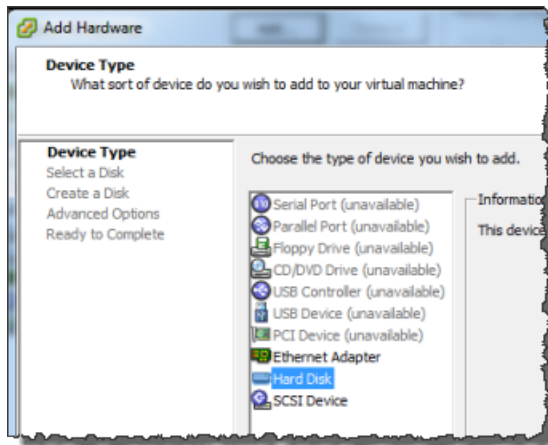
1. Start the VMware vSphere client and connect to your host.
2. In the client, right-click the name of your gateway VM and click **Edit Settings...**



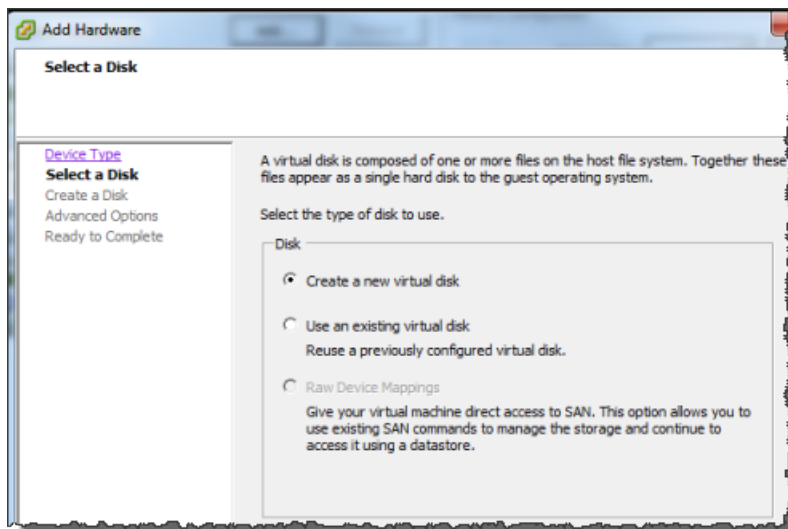
3. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, click **Add...** to add a device.



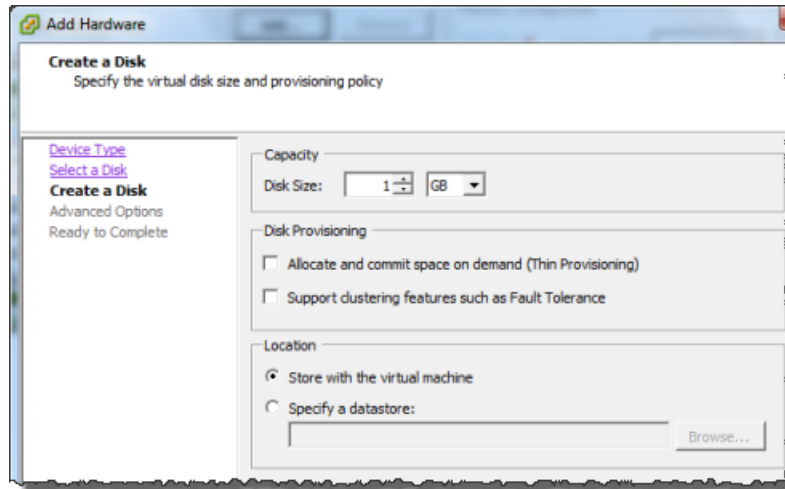
4. Follow the **Add Hardware** wizard to add a disk:
 - a. In the **Device Type** pane, click **Hard Disk** to add a disk, and click **Next**.



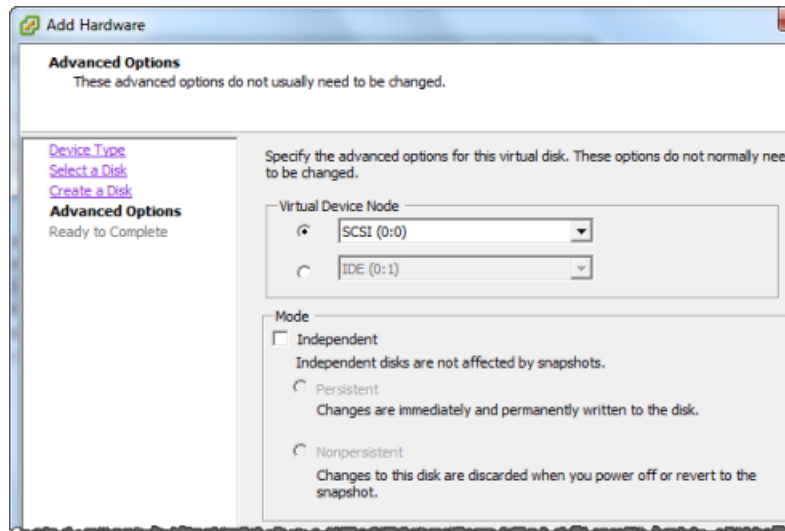
- b. In the **Select a Disk** pane, select **Create a new virtual disk** and click **Next**.



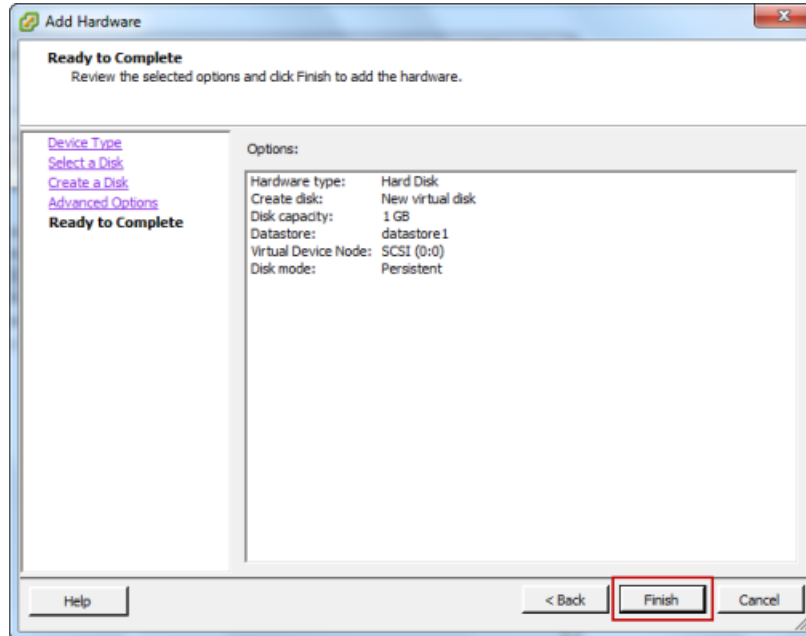
- c. In the **Create a Disk** pane, specify the size of the disk as 1 GB and click **Next**.



- 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 Local Disk for the VM to Use

The gateway needs buffer space to temporarily store data as it uploads snapshots to AWS. This is referred to as working storage. You must add virtual disks to the VM exclusively for use by the VM. The amount of working storage the gateway needs depends on the size of the disks that you allocate for storing your data. For related guidelines, see [Adding Local Disks for AWS Storage Gateway's Working Storage \(p. 51\)](#).

For this getting started tutorial, you allocate 1 GB virtual disk to the VM for exclusive use by the gateway. In the **Create a Disk** pane of the wizard enter 1 GB for the disk size.

Important

The 1 GB virtual disk you allocate for your VM to use as working storage in this tutorial is not suitable for real workloads. It is strongly recommended that you allocate at least 150 GB of working storage. In a later step in this tutorial ([Step 7: Sizing Your Working Storage for Real-World Workloads \(p. 43\)](#)) you will learn about sizing working storage appropriately for real workloads.

To allocate a local disk for working storage

- Repeat the steps in the [To allocate a local disk to store your application data \(p. 18\)](#) procedure to add another virtual disk to the gateway.

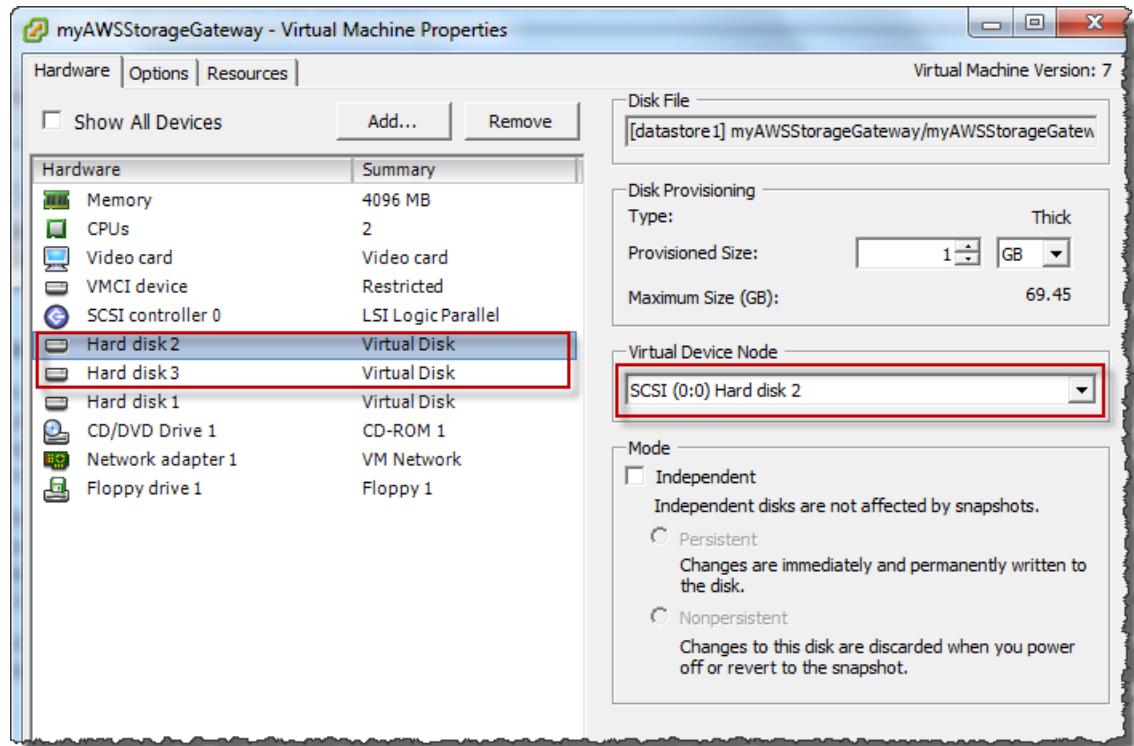
Verify the Gateway VM Has Two Disks

The remainder of the Getting Started tutorial 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 to store your application data \(p. 18\)](#) 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.



Configure the AWS Storage Gateway to Use Paravirtualized Disk Controllers

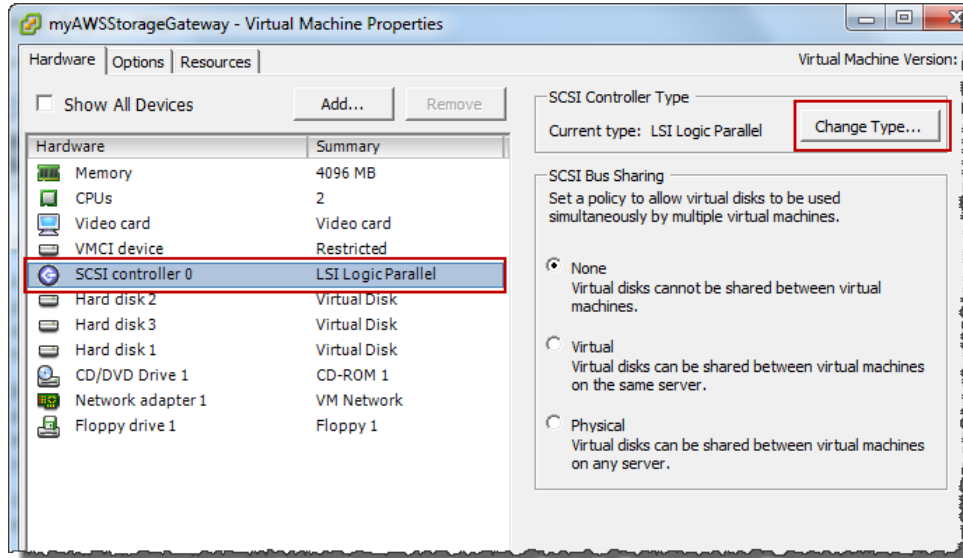
In this task, the iSCSI controller is set so that the VM uses paravirtualization, a mode where the gateway VM works with the host OS, to enable the console to properly identify the virtual disks that you add to your VM.

Note

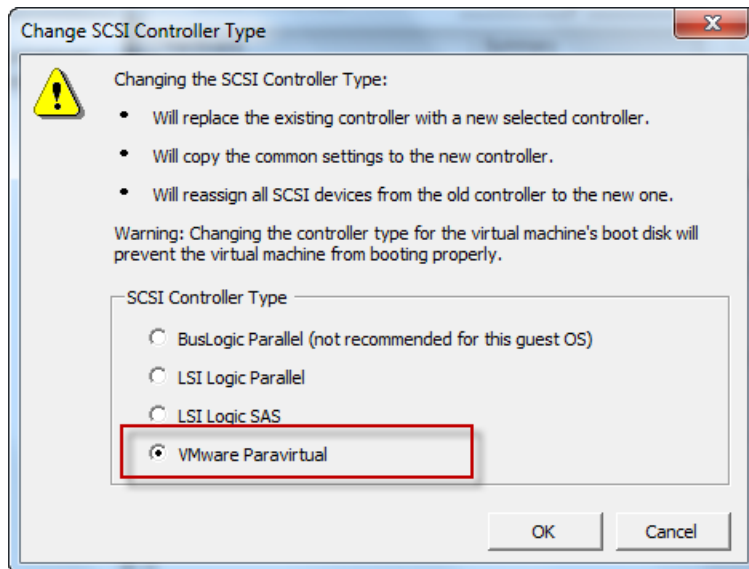
You must complete this step to avoid issues in identifying these disks in the gateway console later when you configure working storage and storage volumes.

To configure your VM to use paravirtualized controllers

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



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

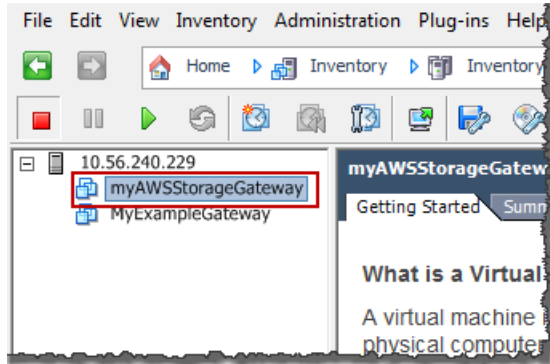


Activate AWS Storage Gateway

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.

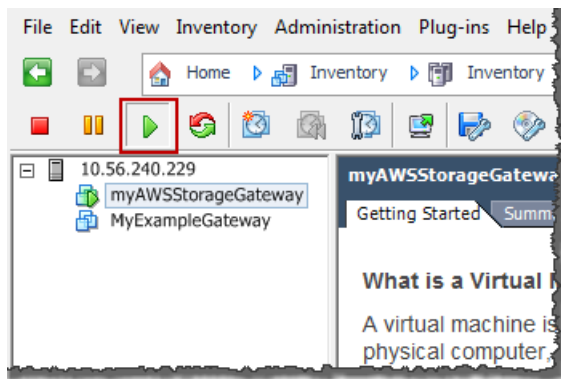
To activate your gateway

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.

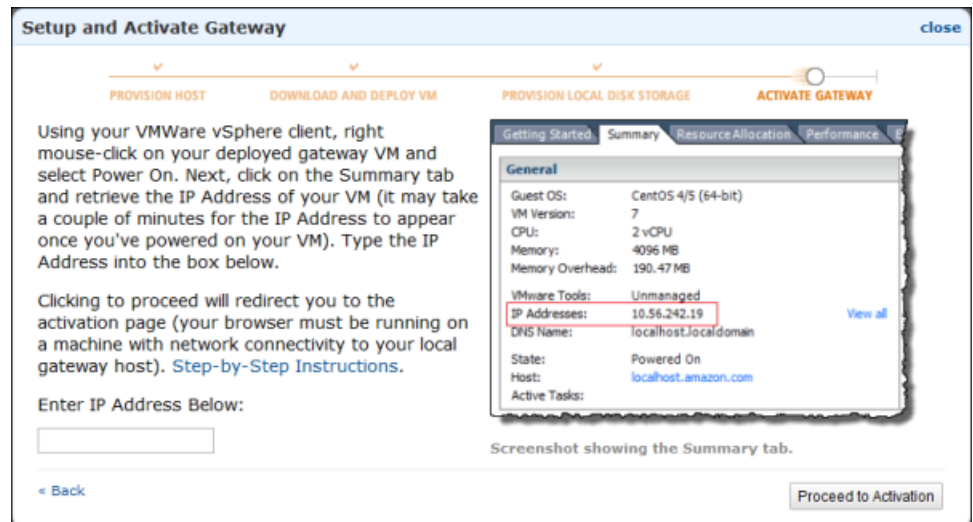
Note

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 Activate 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.

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 IP Address of the gateway you are activating.

IP Address: 10.56.250.1

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

Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway. Activated gateways are billed at \$125 per month, prorated daily. A one month free trial in the form of a \$125 credit will automatically be applied upon activation of your first gateway. Storage will be charged separately.

AWS Region: US East (Virginia) ▼

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

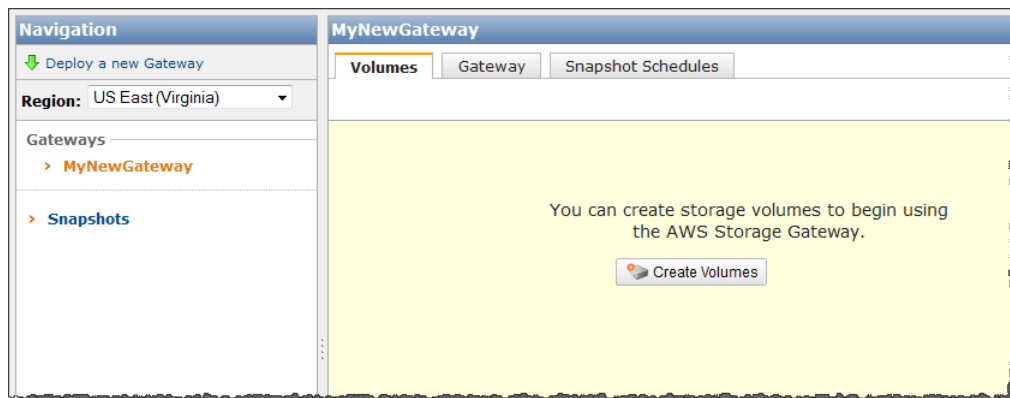
Gateway Name: MyNewGateway

Activate My Storage Gateway

- 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.

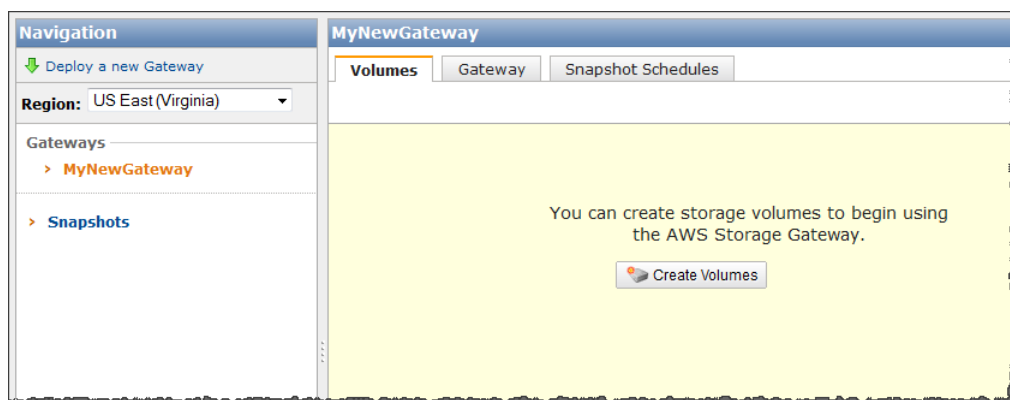


Step 3: Create Storage Volumes Using the AWS Storage Gateway Console

After activating your gateway, the next step is to create storage volumes using the AWS Storage Gateway console. In the preceding steps, you added two virtual disks to the VM you deployed. Next, you create a storage volume in the console and map it to one of these virtual disks. You create storage volume for each virtual disk on your VM on which you plan to store your application data. Your client application connects to your storage volume via the iSCSI interface. When your client application writes data to this volume, it is stored on the volume's corresponding virtual disk. AWS Storage Gateway creates a default snapshot schedule (once a day) for the storage volume that you create. Your gateway takes a daily snapshot of your volume and uploads it to AWS. This snapshot schedule is configurable.

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.

Create Storage Volume close

CREATE VOLUMES CONFIGURE LOCAL WORKING STORAGE

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: Preserve existing data

iSCSI Target Name:

Based on Snapshot ID:

Size:

Host IP:

Port:

- a. In the drop-down list of the **Disk** field, select one 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 lower case 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`. You must make sure that the target name is globally unique within your SAN network.
- 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** field is the IP address of your gateway and click **Create Volume**.

Create Storage Volume close

CREATE VOLUMES CONFIGURE LOCAL WORKING STORAGE

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: 1 GiB

Host IP: 10.56.250.1

Port: 3260

1 Volume Created

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

Step 4: Configure Working Storage for the AWS Storage Gateway VM

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

Note

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

To allocate working storage for your AWS Storage Gateway VM

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

Create Storage Volume close

CREATE VOLUMES CONFIGURE LOCAL WORKING STORAGE

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: 1 GiB

Host IP: 10.56.250.1

Port: 3260

1 Volume Created

Cancel Create Volume

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

Configure Local Working Storage

2. Select the check box next to the remaining disk to allocate the disk as working storage, 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 working storage. Working storage can be extended later without disrupting the iSCSI I/O.

Configure Local Working Storage close

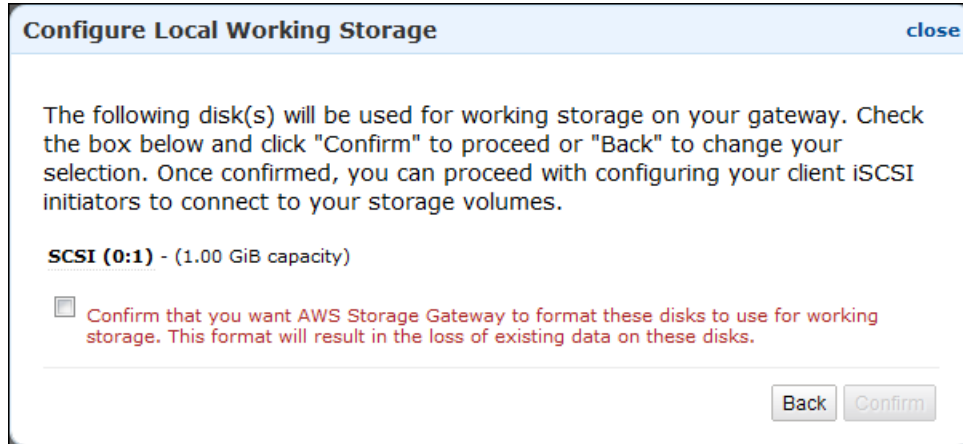
Please select which disks the VM can use for Working Storage. Please see our documentation for recommendations on how much space to provide given your workload and network connection. [Step-by-Step Instructions](#)

Local Disks

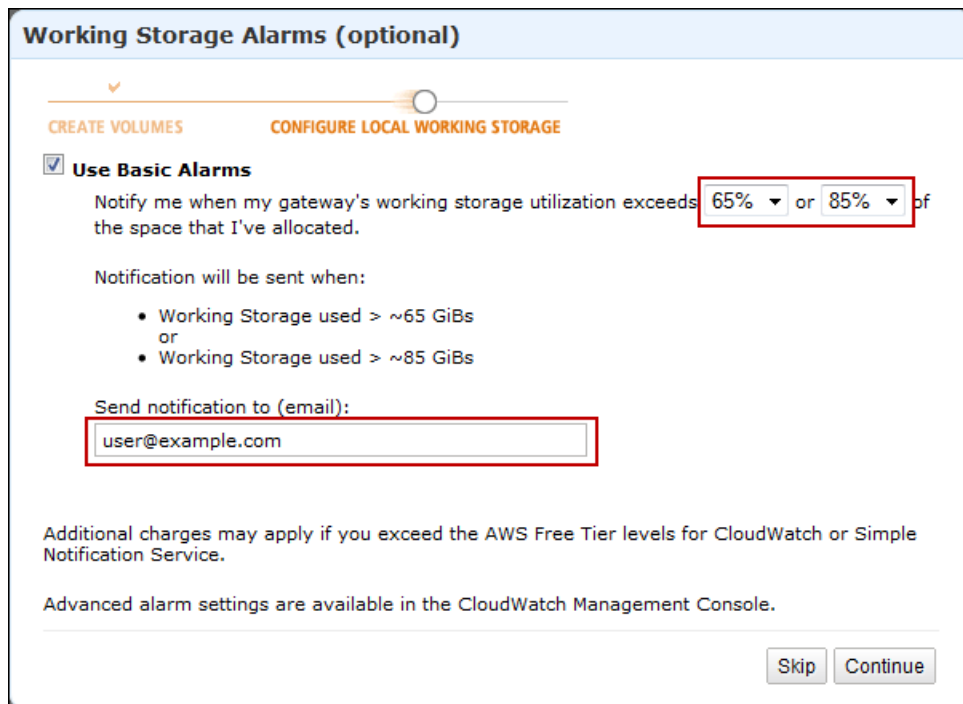
SCSI (0:1)	1.00 GiB	<input checked="" type="checkbox"/>	Use for Working Storage Space
SCSI (0:0)	1.00 GiB		In Use for Storage Volume

Cancel Next

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



4. In the **Working Storage Alarms** dialog box, configure alarms for your working storage.



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

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

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

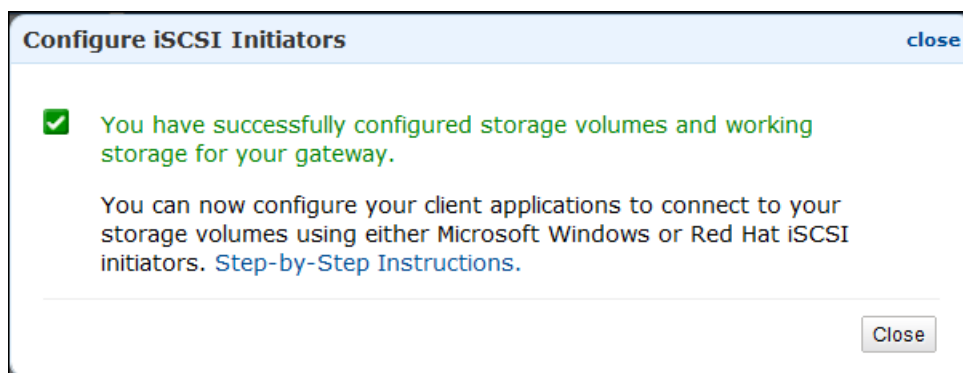
- b. Enter an email address.
c. Click **Continue**.

Two alarms are created. Using the gateway name in this tutorial as an example, the alarms would be named *MyNewGateway-WorkingStorageUtilization-Alarm1* and *MyNewGateway-WorkingStorageUtilization-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 (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 working storage alarms, see [Monitoring AWS Storage Gateway's Working Storage \(p. 124\)](#).

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



Step 5: Access Your AWS Storage Gateway Volumes

You are now ready to connect your Windows client to your storage volume. You make this connection by using the Microsoft iSCSI Initiator on your client.

Note

You must have administrator rights to run the iSCSI Initiator.

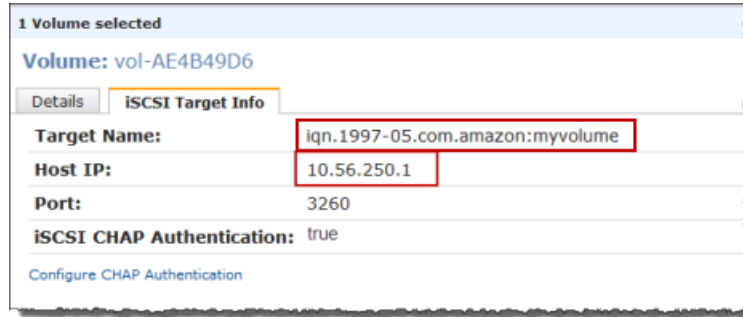
To connect your Windows client to the storage volume

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.

AWS Storage Gateway User Guide

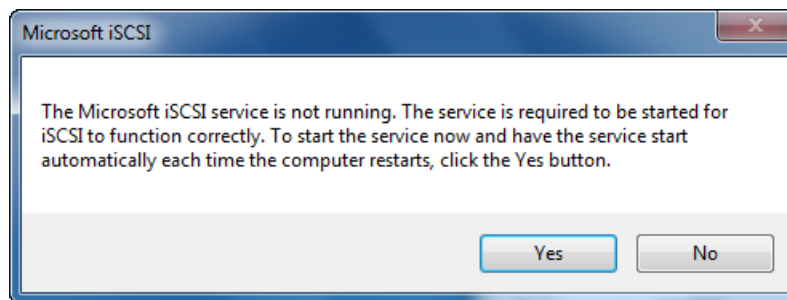
Step 5: Access Your Volumes



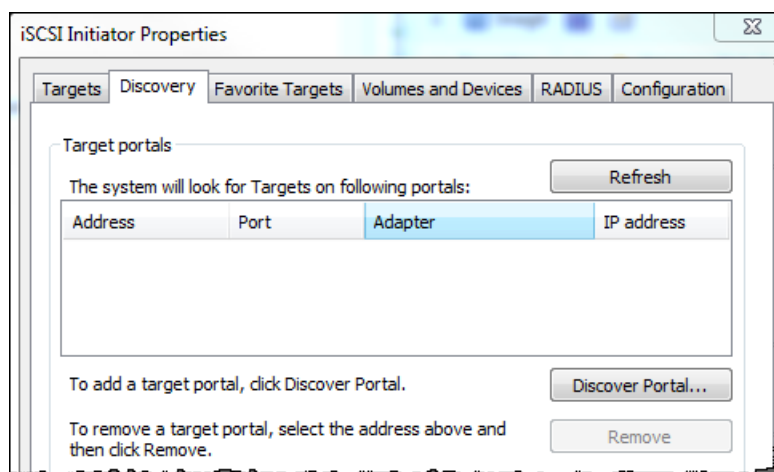
2. Start the iSCSI Initiator.
 - a. In the **Start** menu of your Windows client computer, type `iscsicpl.exe` and run the program.

The **iSCSI Initiator Properties** dialog box appears if the iSCSI Initiator Service is running.

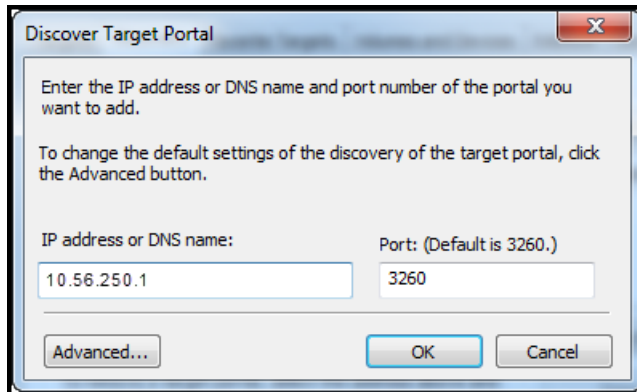
- b. If the Microsoft iSCSI Initiator Service is not running, you are prompted to start the service and have the service start automatically each time the computer restarts. Click **Yes** in the **Microsoft iSCSI** dialog box to start the service.



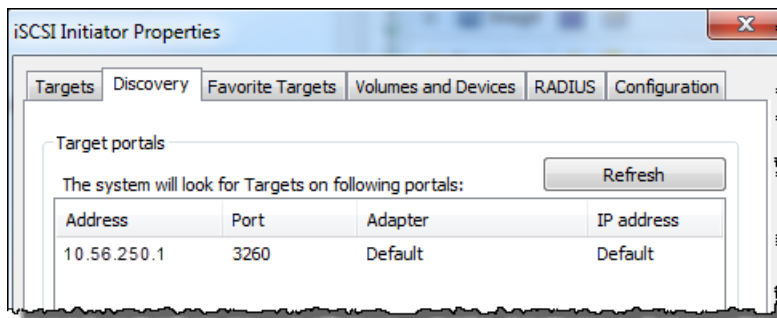
3. Discover the gateway:
 - a. In the **iSCSI Initiator Properties** dialog box, click the **Discovery** tab, and click the **Discovery Portal** button.



- 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**.



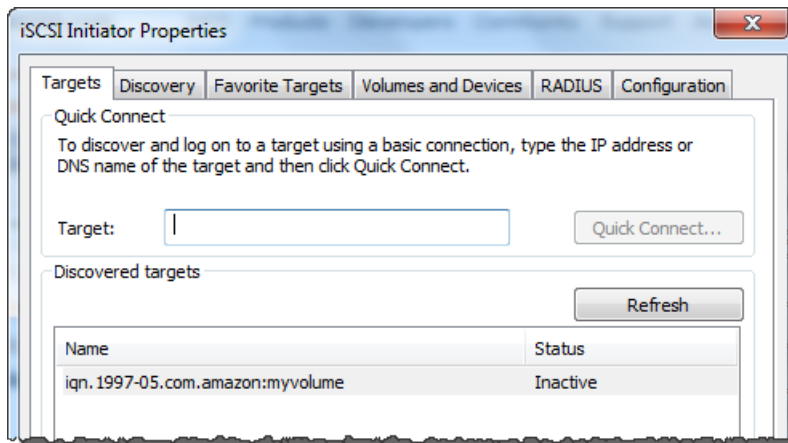
- c. The IP address is now displayed in the lists of **Target portals** in the **Discovery** tab.



- 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 what you noted for your storage volume in step 1.




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

In the **Connect to Target** dialog box, enter the name of the target that you specified when you created the storage volume (see [Step 3: Create Storage Volumes Using the AWS Storage Gateway Console \(p. 27\)](#)), 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

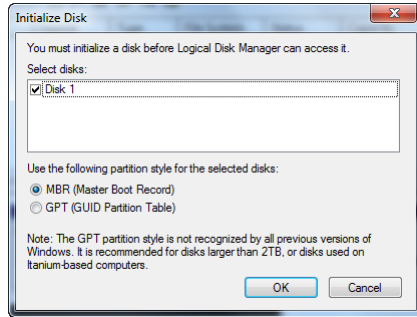
While not required for this Getting Started tutorial, we highly recommend that you customize your iSCSI settings for a real application as discussed in the topic [Customizing Your Windows iSCSI Settings \(p. 78\)](#).

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**.


AWS Storage Gateway User Guide

Step 5: Access Your Volumes

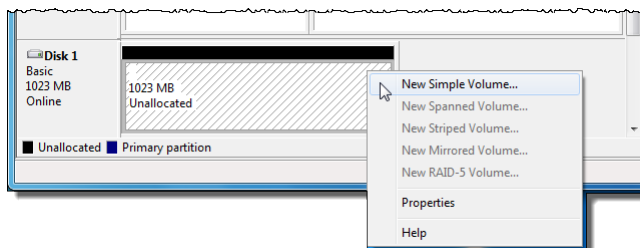


3. Create a simple volume:

- 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 example below, **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 on the disk and select **New Simple Volume ...**




Important


You should exercise caution so as 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 has a status of **Unallocated**.



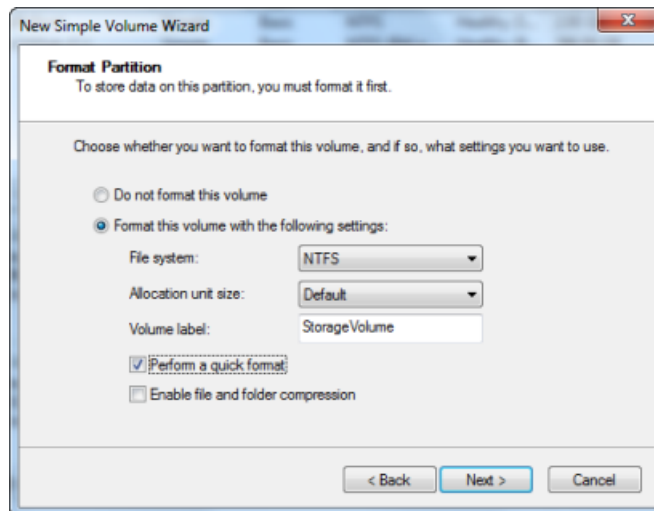
- c. In the **New Simple Volume Wizard**, click **Next**.
- d. In the **Specify Volume Size** dialog box, leave the default values and click **Next**.



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



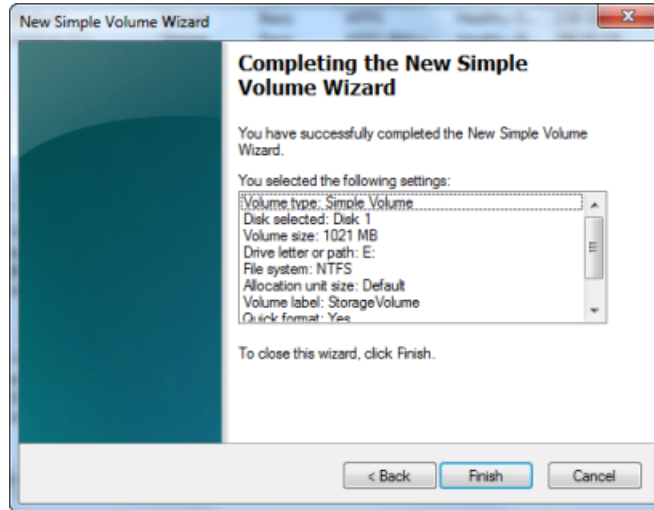
- f. In the **Format Partition** dialog box, specify a **Volume label** field and ensure that **Perform a quick format** is selected. Click **Next**.



- 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 or more to complete.



Step 6: Test the Setup - Take a Snapshot and Restore it to Another AWS Storage Gateway Volume

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 on 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.

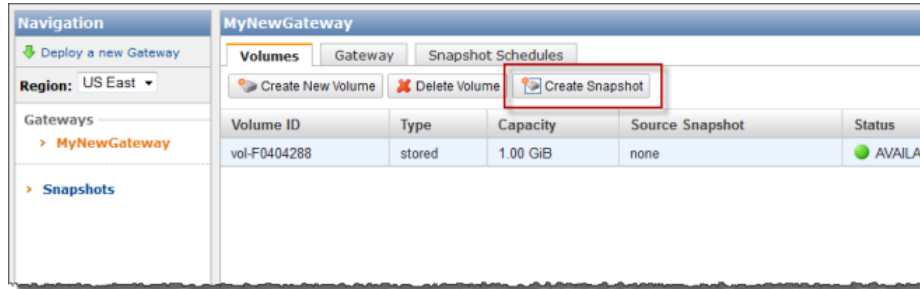
There should only be one storage volume for this gateway. Selecting the volume displays the volume's properties.

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.

AWS Storage Gateway User Guide

Step 6: Test the Setup



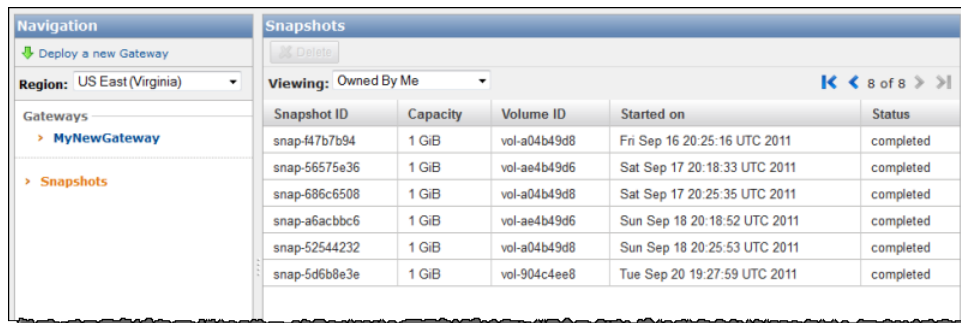
5. In the Snapshot Taken confirmation window, click **Close**.



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 the storage volume on which you created a snapshot was 1 GB in size, create a new virtual disk of the same size. For step-by-step instructions, see [To allocate a local disk to store your application data \(p. 18\)](#).
2. In the AWS Storage Gateway console, click the name of gateway in the **Navigation** pane.
3. Click the **Volumes** tab, and 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** field, enter a name for your iSCSI target, for example **myvolumerestored**.

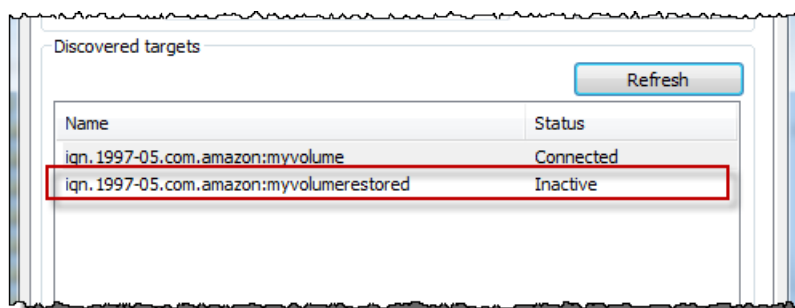
The target name can contain lower case letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name in the **Targets** tab of the **iSCSI Microsoft Initiator** GUI after discovery. For example, a name `target1` would appear as `iqn.1007-05.com.amazon:target1`. You must make sure 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 Volume**.


This creates a storage volume based on your snapshot. The storage volume details appear in the 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**.

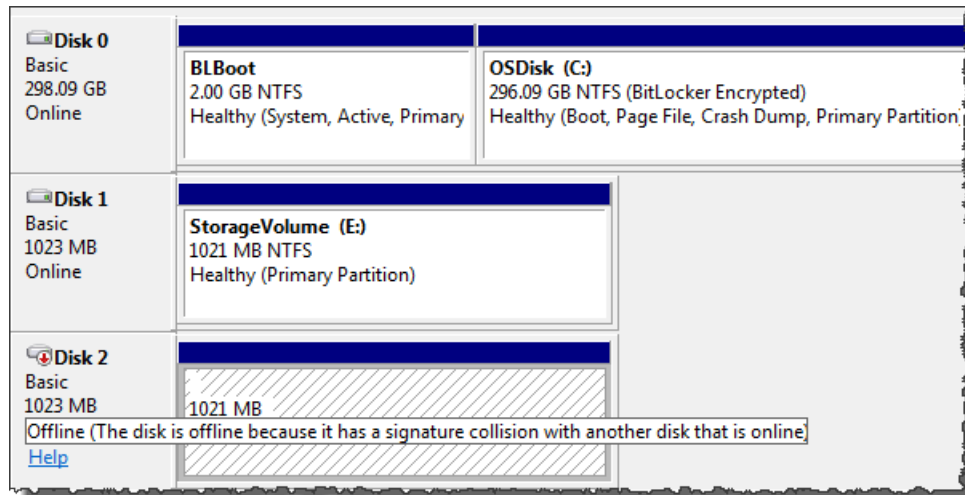


- c. Select the new target and click **Connect**.
- d. In the **Connect to Target** dialog box, Click **OK**.




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.

Step 7: Sizing Your Working Storage for Real-World Workloads

At this point in the Getting Started tutorial you have a simple, working gateway. However, the assumptions used to create this gateway are not appropriate for real-world workloads. In particular, you need to do two things to ensure that your gateway is appropriate for real-world workloads. First, you need to size your working storage appropriately. Second, you need to set up monitoring for your working storage. Both of these tasks are shown in this step.

To size your working storage

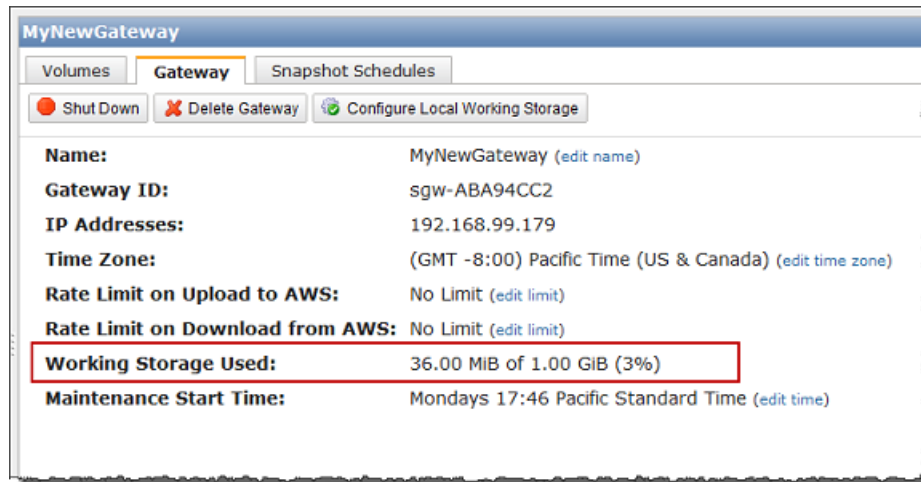
- Use the formula discussed in [Sizing Working Storage Capacity for AWS Storage Gateway \(p. 55\)](#). However, we strongly recommend that you allocate at least 150 GB of working storage. Therefore, if the working storage formula yields a value less than 150 GB, use 150 GB as your allocated working storage.

The working storage 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 GB of working storage space.

To monitor your working storage

1. View your gateway's current working storage.
 - In the **Gateway** tab in the AWS Storage Gateway console, find the **Working Storage Used** field.

The following example shows that working storage at three percent.



2. Set an alarm on working storage.

It is highly recommended that you create a working storage alarm in the Amazon CloudWatch console. For more information, see [To set an upper threshold alarm for a gateway's working storage \(p. 125\)](#).

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 the [Getting Started Tutorial for AWS Storage Gateway \(p. 6\)](#), you created a gateway, provisioned the gateway, and connected your Windows host to the gateway's storage volume. You added data to the gateway's storage volume, took a snapshot of the volume, restored the volume to a new disk, and verified that the data shows on the new disk.

After you finish the Getting Started tutorial, we recommend that you focus on the following areas:

- To understand what storage volumes are and how to create them, see [Managing Storage Volumes in AWS Storage Gateway \(p. 61\)](#).
- To understand what working storage is and how to size it for a gateway, see [Configuring Working Storage in AWS Storage Gateway \(p. 70\)](#).
- To troubleshoot gateway problems, see [Troubleshooting in AWS Storage Gateway \(p. 113\)](#).
- To understand how to optimize your gateway, see [Optimizing AWS Storage Gateway Performance \(p. 117\)](#).
- To understand Storage Gateway metrics and how you can monitor how your gateway performs, see [Monitoring Your AWS Storage Gateway \(p. 118\)](#).
- To connect to the gateway's iSCSI targets to store data, see [Managing Your Application Access to Storage Volumes \(p. 76\)](#).

Setting Up AWS Storage Gateway

Topics

- [Downloading and Deploying AWS Storage Gateway VM \(p. 46\)](#)
- [Provisioning Local Disk Storage for AWS Storage Gateway VM \(p. 47\)](#)
- [Activating AWS Storage Gateway \(p. 57\)](#)

In this section, we discuss the fundamentals of setting up your AWS Storage Gateway which includes downloading and deploying the gateway VM, adding disk storage, and activating the gateway. We recommend that you review the getting started section (see [Getting Started Tutorial for AWS Storage Gateway \(p. 6\)](#)), before continuing in this section.

To begin setup


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

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

2. If you have not already activated a gateway, your experience starts with the following page in the AWS Storage Gateway console. Clicking the **Setup and Activate a New Gateway** button starts the **Setup and Activate Gateway Wizard**.



If you already have one or more gateways activated, you must click the **Deploy a New Gateway** button in the **Navigation** pane to start **Setup and Activate New Gateway** wizard.



Downloading and Deploying AWS Storage Gateway VM

After you provision a local host, the next step in the **Set up and Activate a New Gateway** wizard is to download the AWS Storage Gateway VM software locally, and then deploy the VM to your host.

For step-by-step instructions, see [Download and Deploy the AWS Storage Gateway VM on Your Host](#) (p. 9).

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.

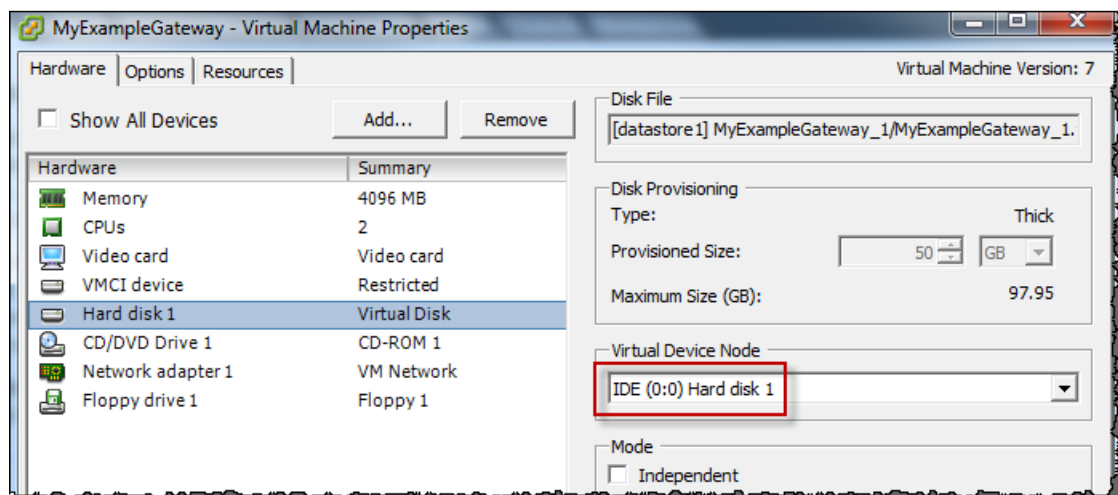
- Deploy the gateway Open Virtualization Application (OVA) on only one host in a cluster.

- 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 \(p. 78\)](#)) and Linux clients (see [Customizing Your Linux iSCSI Settings \(p. 79\)](#)) 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 AWS Storage Gateway VM

After you deploy your AWS Storage Gateway VM, you must add local disks to it as described in the overview (see [How AWS Storage Gateway Works \(p. 3\)](#)). You can use these disks to store your application data locally. You must also allocate one or more of these disks as exclusive working storage for the gateway. Your gateway requires working storage to temporarily buffer your data prior to uploading it to AWS.

Note that after you deploy the 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.




Adding Local Disks to AWS Storage Gateway for Storing Your Application Data

Your application data is stored locally. You must allocate local virtual disks to the VM to store your data. Each disk that you allocate to your VM is used by one of your AWS Storage Gateway iSCSI storage volumes.

AWS Storage Gateway User Guide

Adding Local Disks for Your Application Data

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

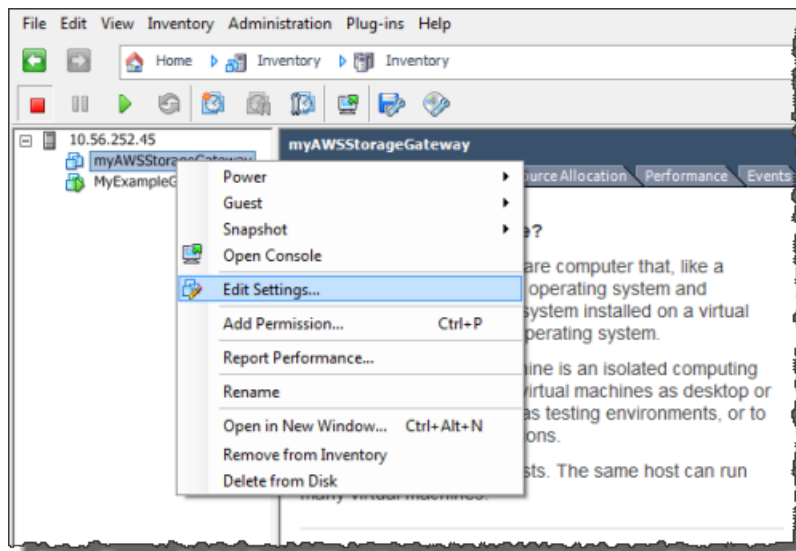


Each disk can be up to 1 TiB in size and must be rounded to the nearest GiB, where GiB is calculated using Base 2 (i.e. GiB = 1024³ bytes).


You can 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. For your application storage, you can allocate disks with existing data. We preserve this data when creating your iSCSI storage volumes. The following procedure provides step-by-step instructions to add a virtual disk from a DAS disk that is available on the host.

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

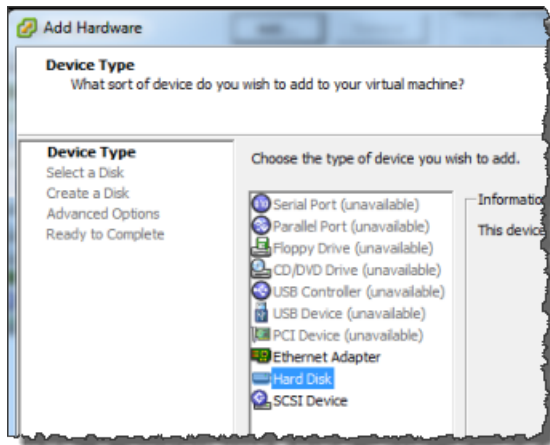
1. Start the VMware vSphere client and connect to your host.
2. In the client, right-click the name of your gateway VM and click **Edit Settings....**



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

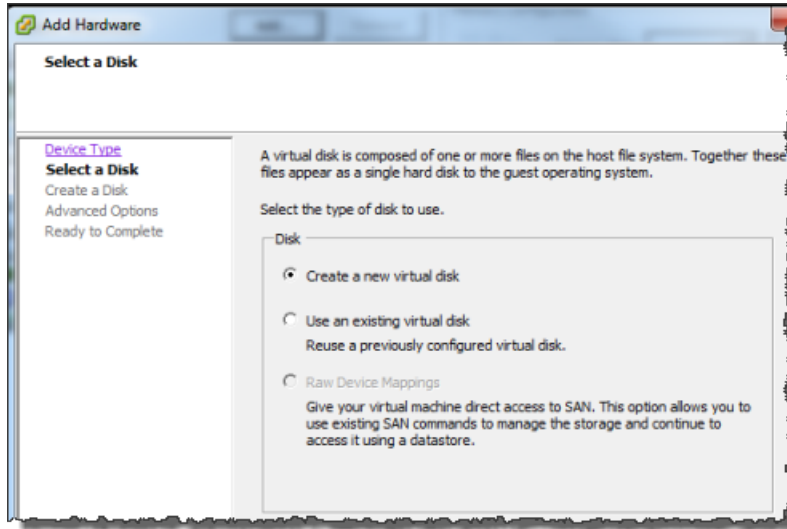


4. Follow the **Add Hardware** wizard to add a disk:
 - a. In the **Device Type** pane, click **Hard Disk** to add a disk, and click **Next**.

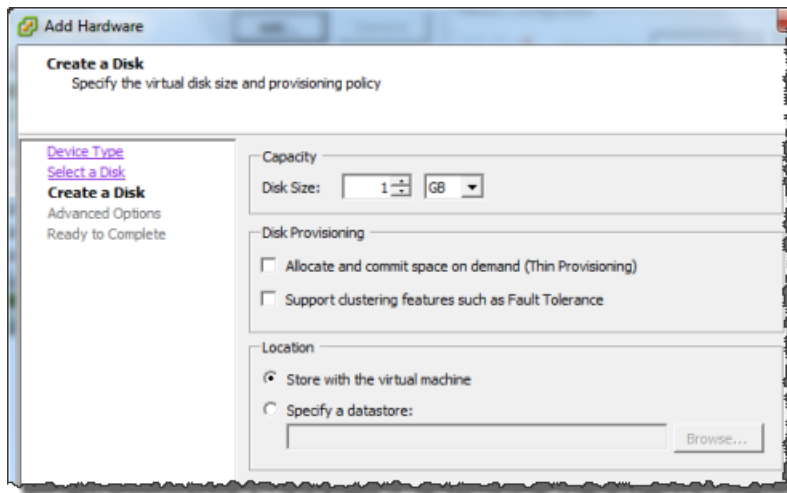


- b. In the **Select a Disk** pane, select **Create a new virtual disk** and click **Next**.

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**..



- 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 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.

To configure your VM to use paravirtualized controllers

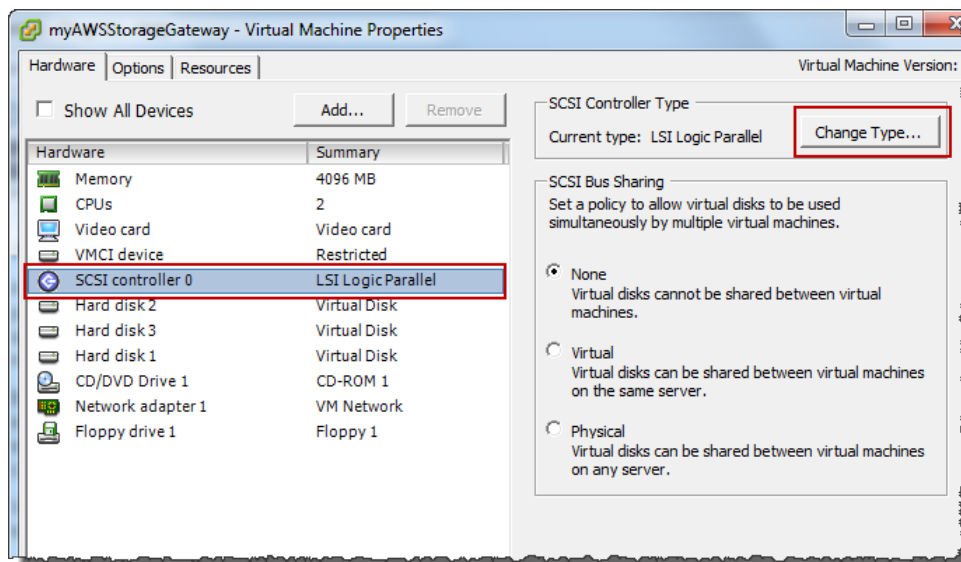
- 1. In the VMware vSphere client, right-click the name of your gateway virtual machine.

Verify that the VM is powered off. If not, power it off. For more information, see [Steps for Activating a Gateway \(p. 58\)](#). Before powering off the VM, make sure that the gateway is not in use.

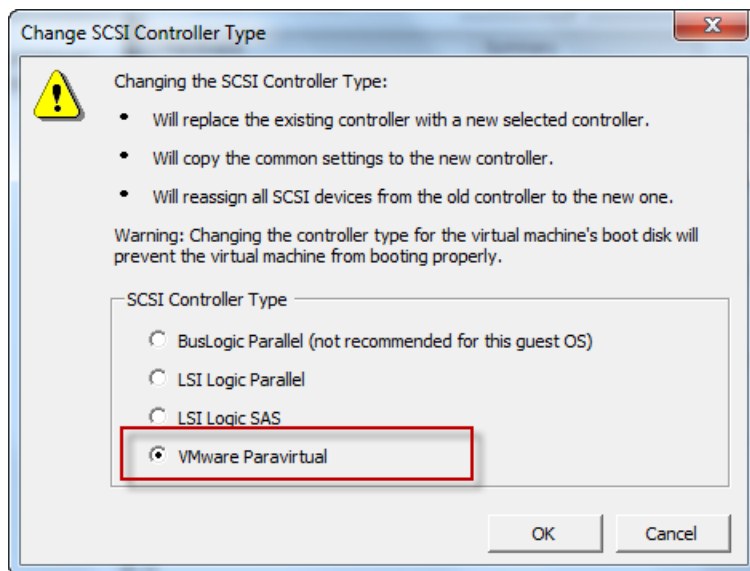
2. Select **Edit Settings....**

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**.



Adding Local Disks for AWS Storage Gateway's Working Storage

Topics


- [Adding a Virtual Disk for Working Storage in AWS Storage Gateway \(p. 52\)](#)
- [Sizing Working Storage Capacity for AWS Storage Gateway \(p. 55\)](#)

You must add virtual disks to your deployed VM for the gateway to use exclusively for working storage. Your gateway requires working storage to temporarily buffer your snapshot data prior to uploading it to AWS. We strongly recommend that you allocate at least 150 GB of working storage to avoid exceeding your gateway's working storage. When working storage is exceeded, 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 working storage and not uploading any of this data to AWS. To size your working storage appropriately, you can use the methods discussed in [Sizing Working Storage Capacity for AWS Storage Gateway \(p. 55\)](#).

Important

Your gateway must have at least one local disk configured as working storage. Working storage disks are required in addition to the local disks that you've allocated to store your application data (storage volume data).

The following diagram highlights the working storage in the larger picture of the AWS Storage Gateway architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)). It illustrates that you must allocate working storage for the AWS Storage Gateway.



Adding a Virtual Disk for Working Storage in AWS Storage Gateway

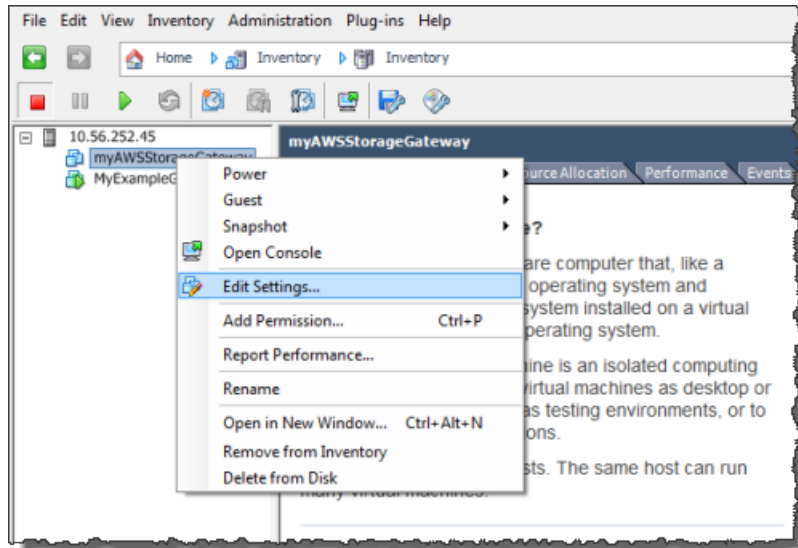
In this section, you allocate a virtual disk to your VM that will be used as working storage for your gateway. To roughly estimate the amount of working storage your gateway requires, simply add up the capacity of the application storage disks that you provisioned, and calculate 20% of this total capacity. To estimate the working storage size more precisely, see [Sizing Working Storage Capacity for AWS Storage Gateway \(p. 55\)](#).

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.

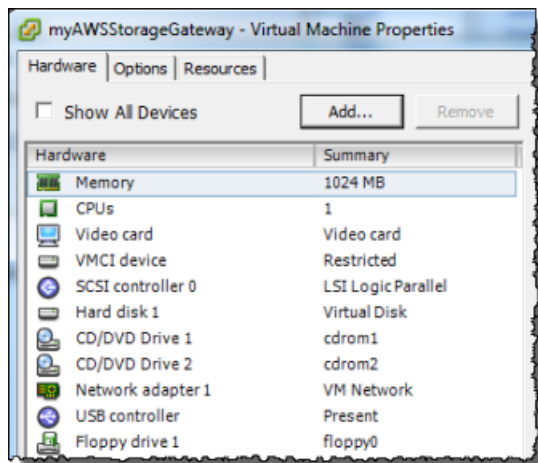
To allocate a new virtual disk to the VM for working storage

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....**



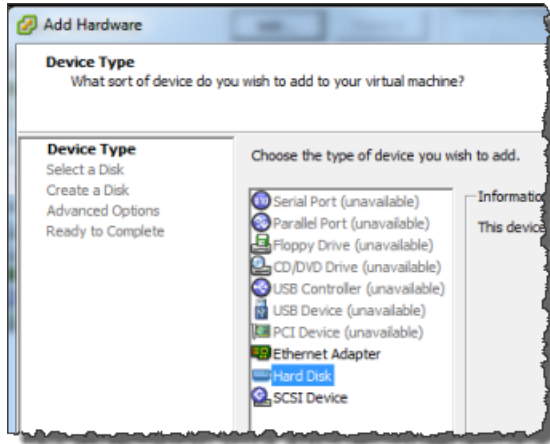
3. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, click **Add...** to add a device.



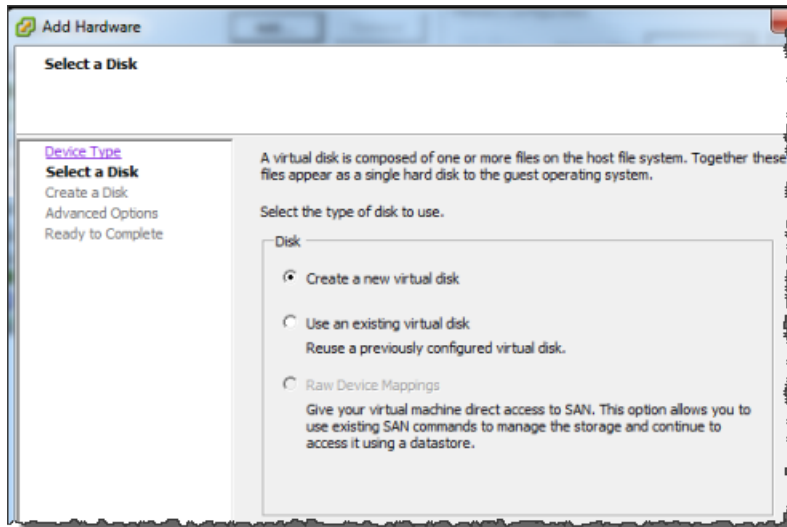
4. Follow the **Add Hardware** wizard to add a disk:
 - a. In the **Device Type** pane click **Hard Disk** to add a disk, and click **Next**.

AWS Storage Gateway User Guide

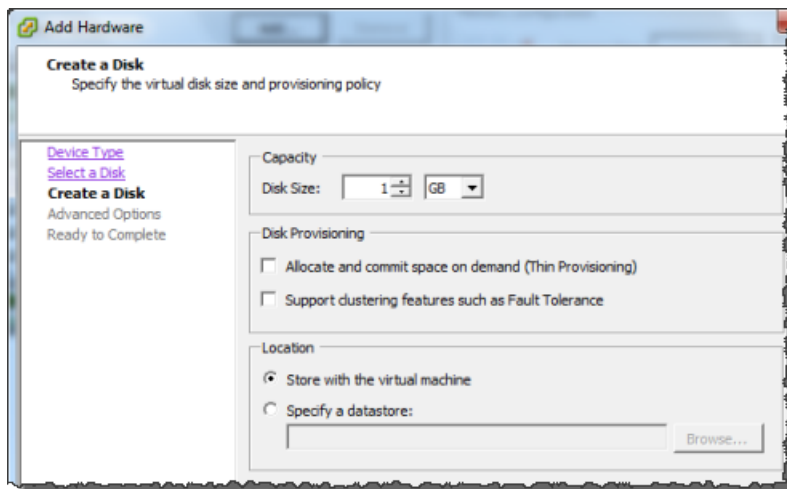
Adding Local Disks for the Gateway Working Storage



- 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 and click **Next**.



- d. In the **Advanced Options** pane, click **Next**.
- e. In the **Ready to Complete** pane, click **Finish**.
- f. If you have not already done so, you must configure your gateway VM to use paravirtualization. For more information, see [Configure AWS Storage Gateway VM to Use Paravirtualization \(p. 56\)](#).

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.

Sizing Working Storage Capacity for AWS Storage Gateway

You can determine the size of your working storage by using a simple, rough estimate or a more precise working storage formula. Regardless of the method you use, it is strongly recommended that you allocate at least 150 GB of working storage. Therefore, if the estimate or formula returns a value less than 150 GB, use 150 GB as the amount you allocate to working storage.

To roughly estimate the amount of working storage your gateway requires, simply add up the capacity of the application storage disks that you provisioned for your storage volumes, and calculate 20% of this total capacity. You can allocate up to 2 TiB of working storage per gateway.

To estimate the amount of working storage 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 working storage 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 working storage your gateway needs.

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

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 working storage.

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

Note that you can initially use this approximation to determine the disk size that you want to allocate to the gateway as working storage. Add more working storage as needed using the AWS Storage Gateway console. Also, you can use the Amazon CloudWatch operational metrics to monitor the working storage

usage and determine additional storage requirements. You can use the `WorkingStorageUsed`, `WorkingStorageFree`, `WorkingStoragePercentUsed` metrics and you can also set up alarms. For more information, see [Monitoring AWS Storage Gateway's Working Storage](#) (p. 124).

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 AWS Storage Gateway VM](#) (p. 47).

To configure your VM to use paravirtualized controllers

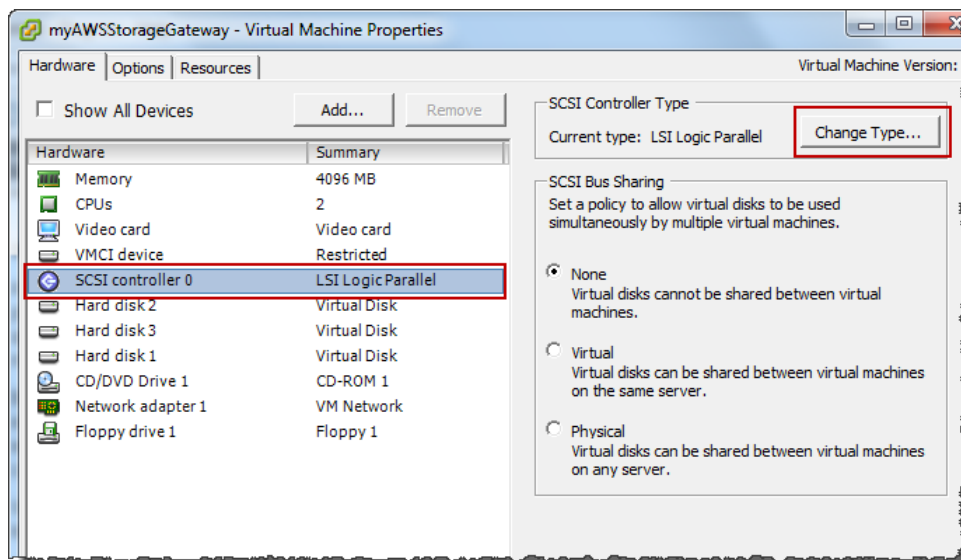
1. In the VMware vSphere client, right-click the name of your gateway virtual machine.

Verify that the VM is powered off. If not, power it off. For more information, see [Steps for Activating a Gateway](#) (p. 58). Before powering off the VM, make sure that the gateway is not in use.

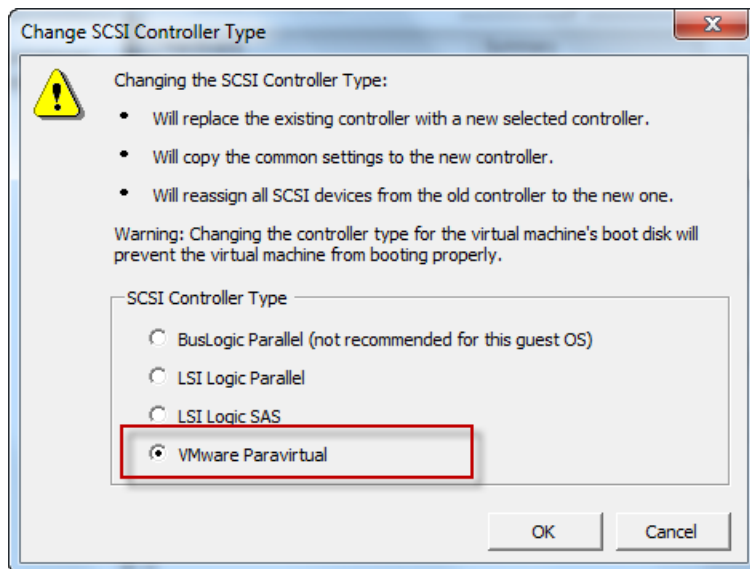
2. Select **Edit Settings...**

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.

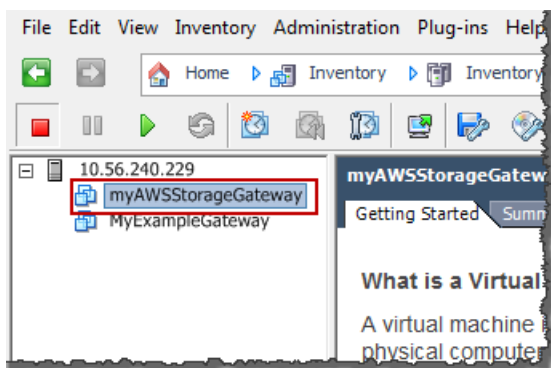
Step	Description
Download and Deploy the VM	In the AWS Storage Gateway console, download the latest virtual machine (VM) that is distributed as an .ova file and deploy this VM on your VMware host. For more information, see Downloading and Deploying AWS Storage Gateway VM (p. 46) .
Provision local disks to the VM	The provisioned VM has no disks. You must add local disks to store your application data and also provide disks to the gateway as working storage. For more information, see Provisioning Local Disk Storage for AWS Storage Gateway VM (p. 47) .
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. 56) .

Steps for Activating a Gateway

To activate a gateway, you need to know the IP address of the gateway VM. You use the IP address in the AWS Management Console to connect to the gateway and initiate the activation process. Before starting the activation process, you should ensure that you have network access to the gateway from the computer that you will use to perform the activation.

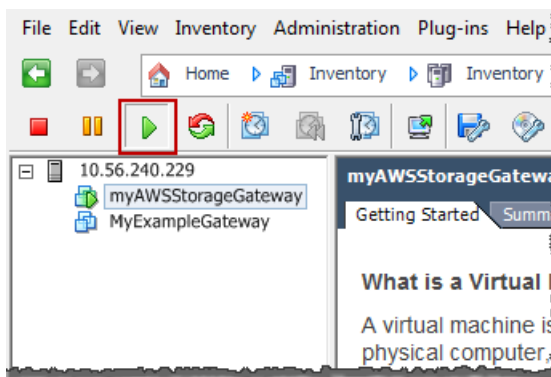
To activate your gateway

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



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

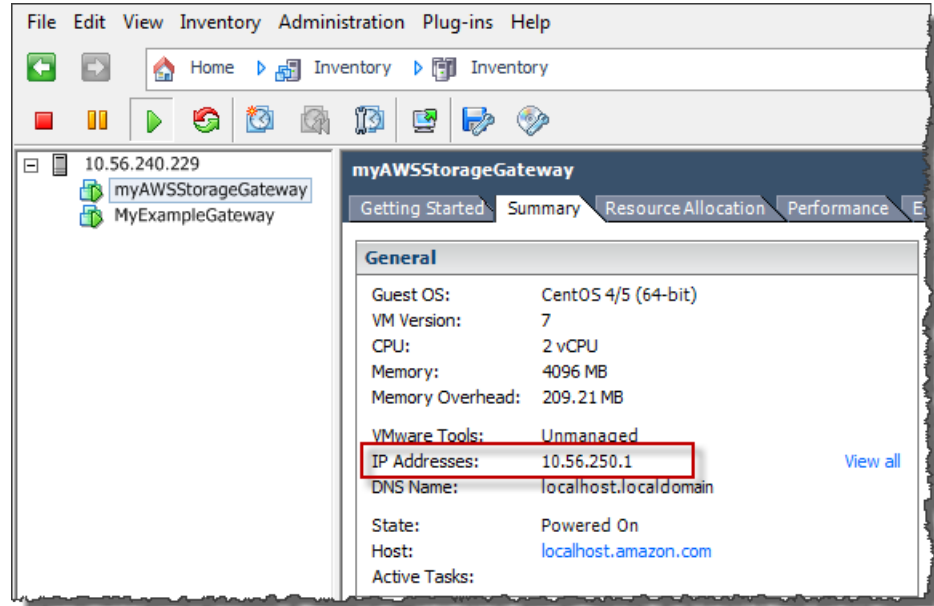
Your gateway VM icon now includes a green arrow icon indicating you have powered on the VM.



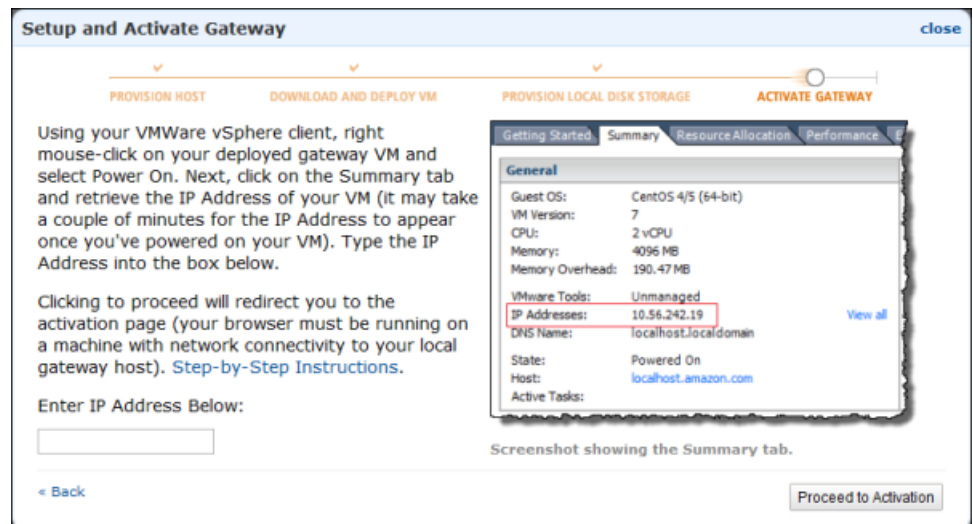
2. Activate the gateway.
 - a. Obtain the IP address of your gateway. Note that, after powering on the VM, it might take a few minutes for the IP address to appear.
 - i. Using the vSphere client, log in to your host.
 - ii. Select the deployed gateway VM.
 - iii. Click the **Summary** tab for the IP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

AWS Storage Gateway User Guide

Steps for Activating a Gateway



- 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

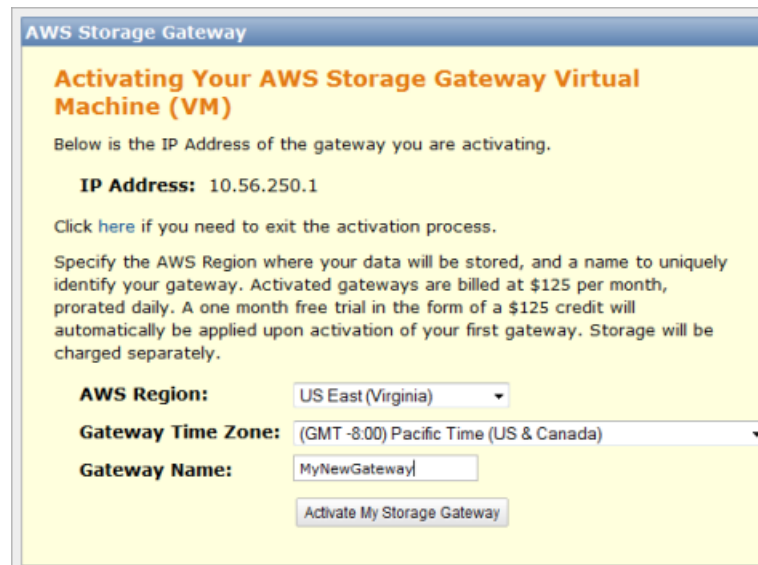
If activation fails, check that the IP address you entered is correct and try to activate again. If the IP address is correct, then confirm that the gateway can access the Internet and if needed set up a proxy (see [Routing AWS Storage Gateway Through a Proxy](#) (p. 104)).

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

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.



AWS Storage Gateway

Activating Your AWS Storage Gateway Virtual Machine (VM)

Below is the IP Address of the gateway you are activating.

IP Address: 10.56.250.1

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

Specify the AWS Region where your data will be stored, and a name to uniquely identify your gateway. Activated gateways are billed at \$125 per month, prorated daily. A one month free trial in the form of a \$125 credit will automatically be applied upon activation of your first gateway. Storage will be charged separately.

AWS Region: US East (Virginia)

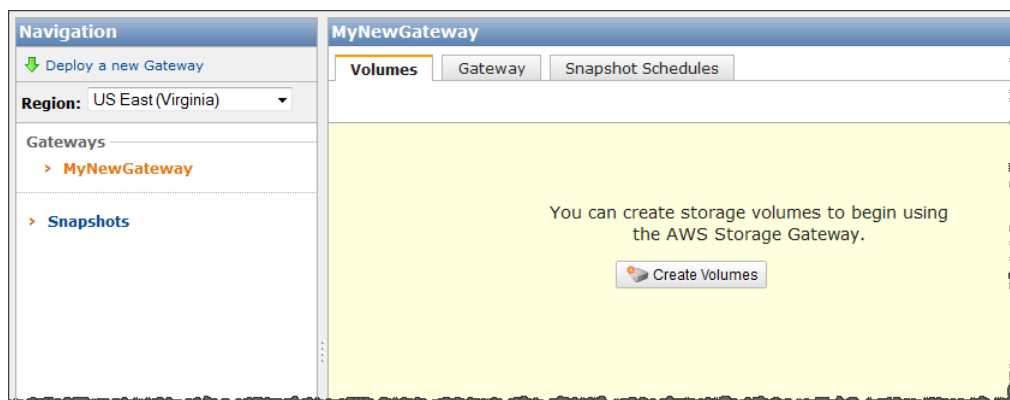
Gateway Time Zone: (GMT -8:00) Pacific Time (US & Canada)

Gateway Name: MyNewGateway

Activate My Storage Gateway

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

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



Related Section

- [API Reference for AWS Storage Gateway \(p. 137\)](#)

Managing Your Activated Gateway

Topics

- [Related Section \(p. 61\)](#)
- [Managing Storage Volumes in AWS Storage Gateway \(p. 61\)](#)
- [Configuring Working Storage in AWS Storage Gateway \(p. 70\)](#)
- [Managing Your Application Access to Storage Volumes \(p. 76\)](#)
- [Working With Snapshots in the AWS Storage Gateway Console \(p. 88\)](#)
- [Performing Maintenance Tasks in AWS Storage Gateway \(p. 97\)](#)
- [Troubleshooting in AWS Storage Gateway \(p. 113\)](#)
- [Optimizing AWS Storage Gateway Performance \(p. 117\)](#)
- [Monitoring Your AWS Storage Gateway \(p. 118\)](#)

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 working storage, working with snapshots, general maintenance, troubleshooting, and monitoring your gateway. If you have not set up a gateway, see [Setting Up AWS Storage Gateway \(p. 45\)](#).

Related Section

- [API Reference for AWS Storage Gateway \(p. 137\)](#)

Managing Storage Volumes in AWS Storage Gateway

Topics

- [Storage Volume Status in AWS Storage Gateway \(p. 67\)](#)

To enable your applications to save data on the local virtual disks that you added to your gateway VM (see [Provisioning Local Disk Storage for AWS Storage Gateway VM \(p. 47\)](#)), the gateway exposes the disks as iSCSI targets. For a gateway to expose a disk to your applications, you must create iSCSI storage volumes in the gateway, one for each local disk on which you want to store application data. When creating

a storage volume, you must identify the disk and provide additional information such as the iSCSI target name.

The AWS Storage Gateway console is the management tool you use to manage your gateway and storage volumes. The console provides the following UI to create a storage volume.

Create Storage Volume close

CREATE VOLUMES CONFIGURE LOCAL WORKING STORAGE

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: Preserve existing data

iSCSI Target Name:

Based on Snapshot ID:

Size:

Host IP:

Port:

In the **Create Storage Volume** wizard, the **Disk** drop-down shows a list of local disks on your gateway. You select a disk and provide information including the **iSCSI Target Name**. Your gateway exposes this volume as an iSCSI target with the name you specify, prepended by `iqn.1997-05.com.amazon:.` For example, based on the information in the preceding screen shot, the iSCSI target will appear as `iqn.1997-05.com.amazon:myvolume`. You then configure your applications to mount this volume over iSCSI. Your gateway stores your application data locally on your storage volume's disk, while asynchronously uploading your data to AWS.

Note

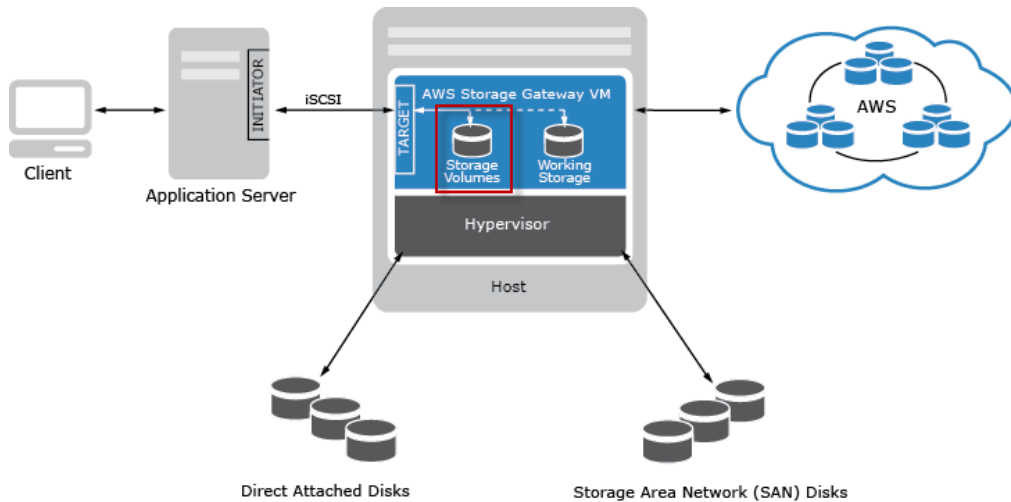
Storage volumes can range from 1 GiB to 1 TiB in size and must be rounded to the nearest GiB. Each gateway can support up to 12 volumes and up to 12 TiB of local storage. If you need to store more data locally on your gateway, complete this [request form](#) and your use case and storage increase will be considered.

Note

Resizing the underlying disk size 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 a storage volume \(p. 65\)](#). For steps describing how to add a storage volume and preserve existing data, see [To create a storage volume \(p. 63\)](#).

When creating a storage volume, you can optionally provide an existing snapshot ID in which case the gateway downloads the snapshot data from AWS to the storage volume's local disk. This is a useful scenario if you want to restore a backup to another storage volume.

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

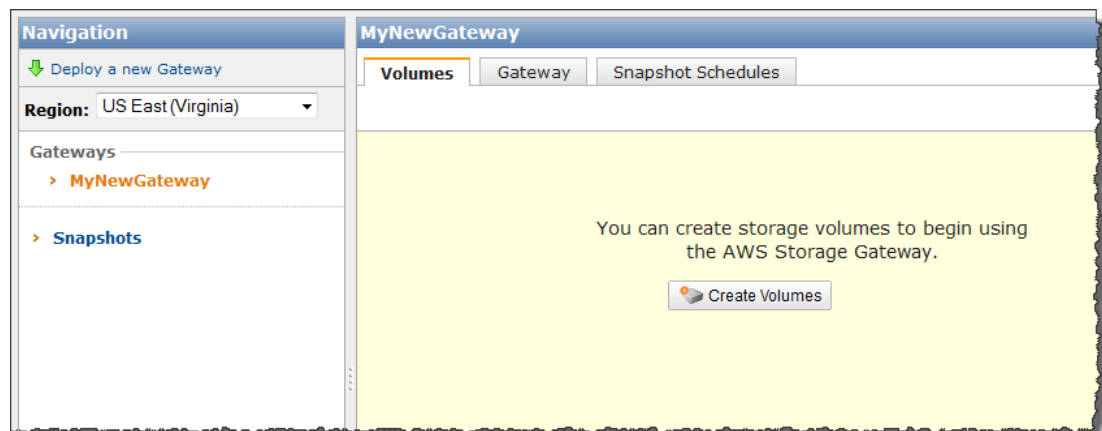


The following tasks assume that you already have a deployed and activated gateway. Furthermore, it is assumed that there is at least one locally provisioned disk of the gateway that is not used. You create a storage volume exposing this disk as the iSCSI target.

To create a storage volume

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**.

Create Storage Volume close

CREATE VOLUMES CONFIGURE LOCAL WORKING STORAGE

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: Preserve existing data

iSCSI Target Name:

Based on Snapshot ID:

Size: --

Host IP:

Port:

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 AWS Storage Gateway VM \(p. 47\)](#).
 - 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 lower case 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.1997-05.com.amazon:target1`. You must make sure 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 your existing snapshot data to the volume. To learn about how to find a snapshot you want to use, see [Finding a Snapshot Using the AWS Storage Gateway Console \(p. 90\)](#). 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 (NIC), 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 only supports 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 Using the AWS Storage Gateway Console \(p. 91\)](#).

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.

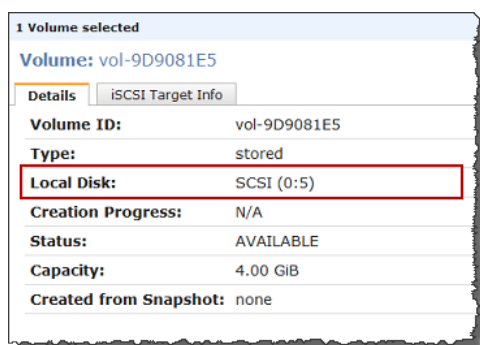
As you add more storage volumes, you must consider the size of your working storage you allocated to the gateway. The gateway must have sufficient buffer space. For more information, see [Configuring Working Storage in AWS Storage Gateway \(p. 70\)](#).

To remove 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 Using the AWS Storage Gateway Console \(p. 91\)](#). Perform the following task to remove a storage volume 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. 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.

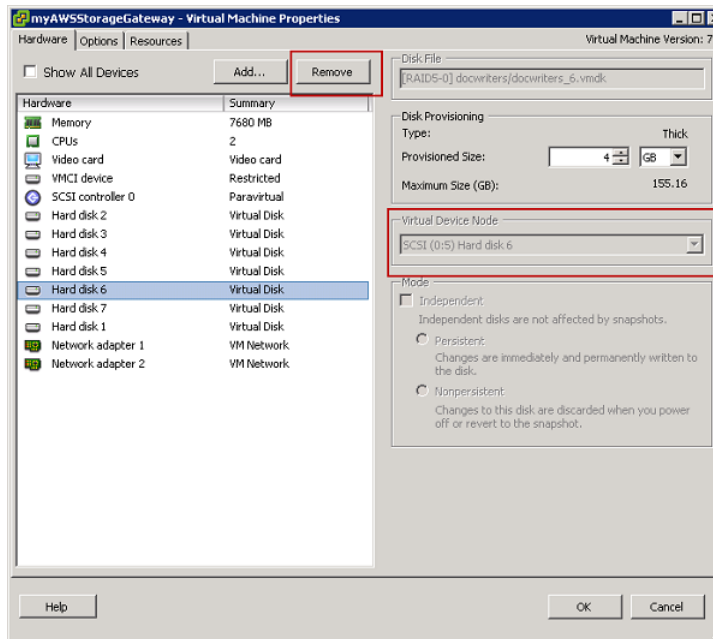
This value is the disk's **Virtual Device Node** value that you use in the vSphere client to ensure that you remove the correct disk.



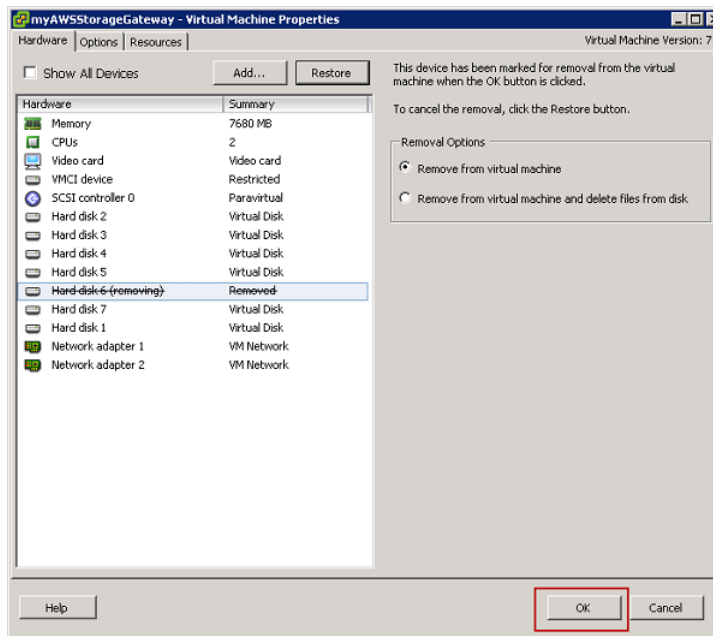
3. Click **Delete Volume**.
4. If you want to remove the disk from the VM, use the vSphere client.
 - a. In the vSphere client, right-click the name of your gateway VM and click **Edit Settings...**
 - b. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, select the disk to remove and click **Remove**.

AWS Storage Gateway User Guide Managing Storage Volumes

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.

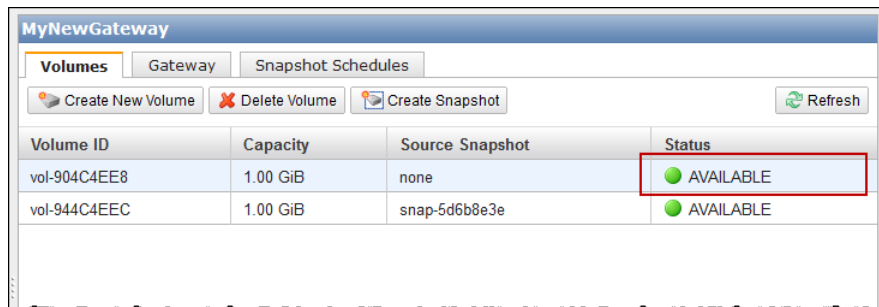


- c. Choose an option in the **Removal Options** panel and click **OK** to complete the process of removing the disk.



Storage Volume Status in AWS Storage Gateway

The AWS Storage Gateway console shows a **Status** field for each of the storage volume on your gateway as shown in the following screen shot.



The following table describes the storage volume status values and their meaning.

Status	Description
AVAILABLE	The normal running status for a volume.
BOOTSTRAPPING	Bootstrapping occurs when you choose to create a volume using a disk that has existing data that you want to preserve. In this case, your gateway starts uploading all of the data up to AWS. Bootstrapping also occurs when your volume is in PASS THROUGH and the amount of free working storage increases sufficiently. You can provide additional working storage as one way to increase the percent of free working storage. In the bootstrapping process, 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.
CREATING	The volume is currently being created and is not ready to be used.
DELETING	The volume is currently being deleted.
IRRECOVERABLE	An error occurred from which the volume cannot recover.

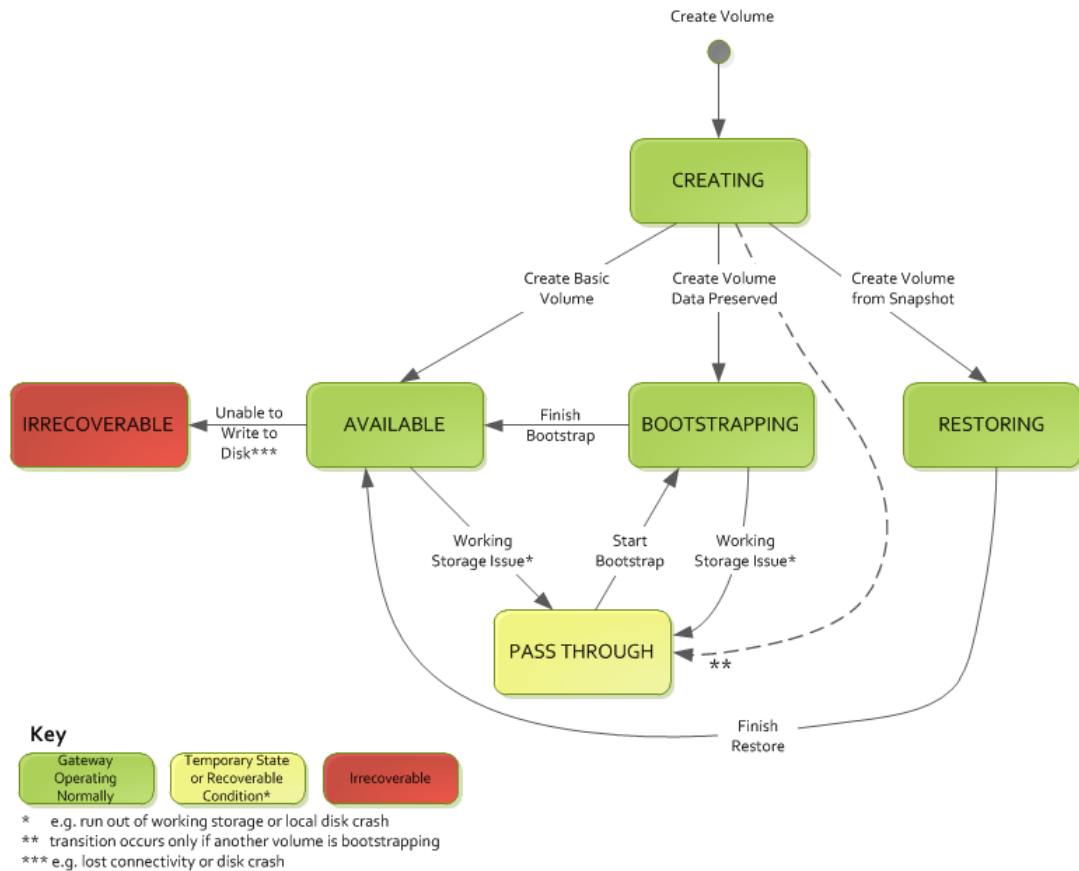
Status	Description
PASS THROUGH	<p>This volume status can occur for several reasons. One reason that can cause PASS THROUGH status is if your gateway has run out of working storage. Your applications can continue to read from and write data to your storage volumes while they are in PASS THROUGH; however, the gateway is not writing any of your volume data to its working storage 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 working storage is exceeded, see Troubleshooting Storage Volume Issues (p. 115).</p> <p>Another reason for a volume to indicate the PASS THROUGH status is because there is more than one storage volume boot strapping at once. Only one gateway storage volume can boot strap 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 boot strapping. 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.</p> <p>Infrequently, the PASS THROUGH status can indicate that a working storage disk has failed. For information about what action to take in this scenario, see Troubleshooting Storage Volume Issues (p. 115).</p>
RESTORING	<p>The volume is being restored from an existing snapshot.</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>
RESTORING_PASS_THROUGH	<p>The volume is being restored from an existing snapshot and encountered a working storage issue.</p> <p>One reason that can cause the RESTORING_PASS_THROUGH status is if your gateway has run out of working storage. 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 working storage is exceeded, see Troubleshooting Storage Volume Issues (p. 115).</p> <p>Infrequently, the RESTORING_PASS_THROUGH status can indicate that a working storage disk has failed. For information about what action to take in this scenario, see Troubleshooting Storage Volume Issues (p. 115).</p>
WORKING STORAGE NOT CONFIGURED	<p>The volume cannot be created or used because the gateway does not have working storage configured. To add working storage, see Configuring Working Storage in AWS Storage Gateway (p. 70).</p>

Storage Volume Status Transitions in AWS Storage Gateway

The following state diagram describes the most common transitions between volume status states. The diagram shows neither the WORKING STORAGE NOT CONFIGURED state nor the DELETING state. 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 state 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. If this state is because working storage is exceeded, then in some cases working storage may become available again. At that point, the storage volume self-corrects and becomes AVAILABLE. In other cases, you may have to add more working storage to your gateway to allow the storage volume state to become AVAILABLE. To troubleshoot when working storage is exceeded, see [Troubleshooting Storage Volume Issues \(p. 115\)](#). To add working storage, see [Configuring Working Storage in AWS Storage Gateway \(p. 70\)](#).
- **Red**—The storage volume has become IRRECOVERABLE. In this case, you should delete the volume (see [To remove a storage volume \(p. 65\)](#)).

In the diagram, a transition between two states is depicted with a labeled line. For example, the transition from the CREATING state to the AVAILABLE state 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 \(p. 63\)](#).



Note

The volume state of PASS THROUGH is depicted as yellow in this diagram and does not match the color of this state's icon in the **Status** field of the AWS Storage Gateway console.

Configuring Working Storage in AWS Storage Gateway

Topics


- [Ongoing Management of Working Storage for a Gateway \(p. 73\)](#)

Your gateway requires working storage to temporarily buffer your snapshot data prior to uploading it to AWS.

Important

Your gateway must have at least one local disk configured as working storage. Working storage disks are required in addition to the local disks that you've allocated to store your application data (storage volume data).

The following diagram highlights the working storage in the larger picture of the AWS Storage Gateway architecture (see [How AWS Storage Gateway Works \(p. 3\)](#)). It illustrates that you must allocate working storage disks for use by the gateway.



As part of setting up and deploying your VM, you add local virtual disks for your gateway (see [Adding Local Disks for AWS Storage Gateway's Working Storage \(p. 51\)](#)) to use as working storage. You can allocate up to 2 TiB of working storage per gateway. This section describes how to configure your gateway to use these disks. When you allocate disks as working storage, you lose any of your pre-existing data on these disks.

You can provide additional working storage to your gateway without interrupting any gateway functions. The AWS Storage Gateway console is your management tool to configure additional working storage. The console shows you available virtual disks on your VM that you can configure as additional working storage.

If there are no virtual disks available on your VM, then you must first add more disks to your VM. For step-by-step instructions about adding more local disks to your VM, see [Adding Local Disks for AWS Storage Gateway's Working Storage \(p. 51\)](#). Note that you can add more working storage with the gateway VM powered on; however, when you reduce the amount of working storage, you must first power off the VM.


The following tasks assume that you already have a deployed and activated gateway and you have at least one local disk available on your VM that you can allocate as working storage to the gateway.

To configure working storage for your gateway

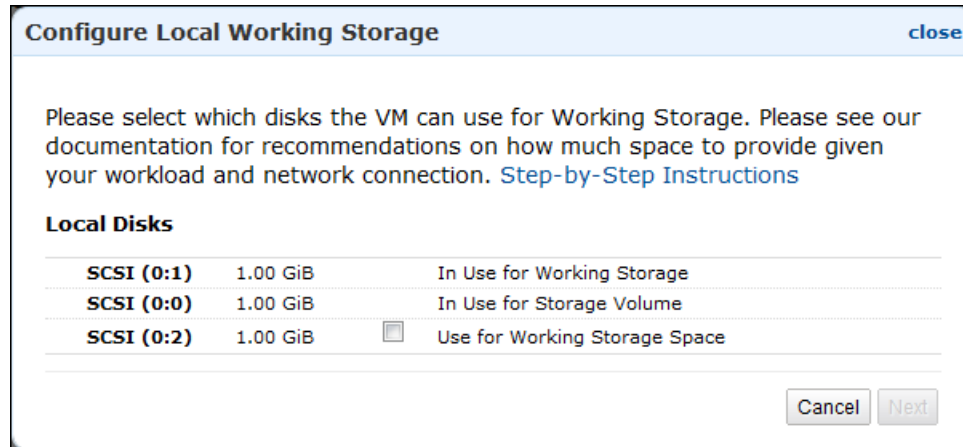
1. In the AWS Storage Gateway console, start the **Configure Local Working Storage** wizard.

You may also be directed to this wizard if you are creating a storage volume for the first time.

- a. Click the gateway in the **Navigation** pane.
- b. Select the **Gateway** tab.
- c. Click **Configure Local Working Storage**.



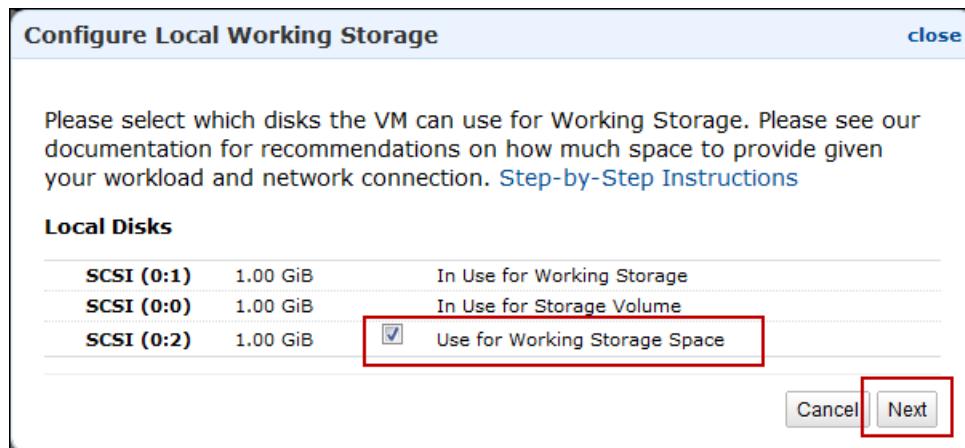
The **Configure Local Working Storage** wizard shows a list of available disks on your local VM.



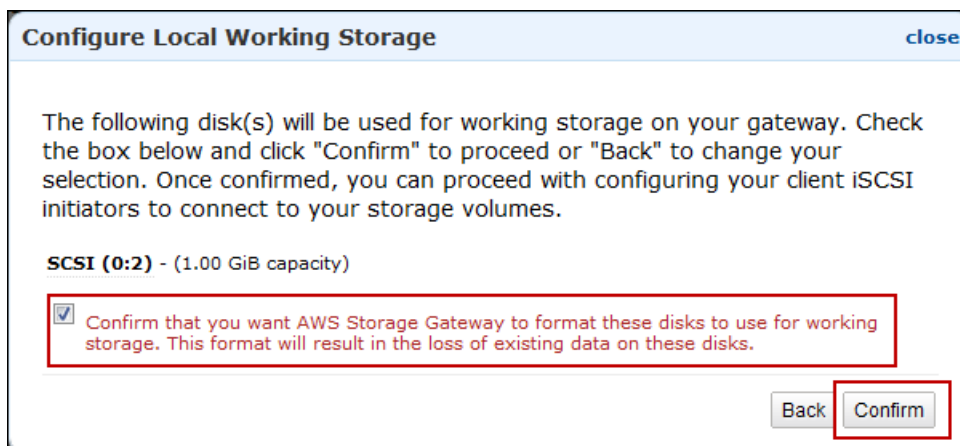
2. If there are no local disks available to configure as working storage, then you must first add a local disk to your gateway VM. For more information, see [Adding Local Disks for AWS Storage Gateway's Working Storage](#) (p. 51).
3. If there are disks available to configure as working storage, then configure the gateway to use them.
 - a. Select the check box next to the disks that you want to allocate to the gateway as working storage and click **Next**. The **Next** button is only enabled if you select at least one disk.

Important

After configuring a disk as working 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 working storage for the gateway.



Ongoing Management of Working Storage for a Gateway

Topics

- [Adding More Working Storage \(p. 74\)](#)
- [Removing Working Storage \(p. 74\)](#)

As part of initial deployment, you configure working storage for the gateway based on an estimation (see [Adding Local Disks for AWS Storage Gateway's Working Storage \(p. 51\)](#)). However, as you add storage volumes to provide more storage for your application data, you might need to add more working storage that provides additional buffer space for the gateway. Conversely, you might need to remove a working storage disk, for example, because you want to reduce the amount of working storage for a gateway or replace a working storage disk that has failed. This section reviews how you would add or remove a disk that is used as working storage.

The amount of working storage 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 local storage volumes, and network throughput is not sufficient for the gateway to upload data to AWS, then eventually your working storage will be filled with data waiting to be uploaded to AWS. You must ensure that working storage space does not fill up. You can do this by monitoring the percentage of working storage that is being used by your gateway. Amazon CloudWatch provides usage metrics such as the `WorkingStoragePercentUsed` metric for monitoring gateway working storage (see [Monitoring AWS Storage Gateway's Working Storage \(p. 124\)](#)). You can set a threshold to trigger a notification to you when working storage usage exceeds a certain capacity.

If your working storage is getting filled close to capacity, you should consider adding more working storage to the gateway. If the working storage capacity is exceeded, the impacted storage volumes go into a [PASS THROUGH \(p. 68\)](#) mode. That is, your applications can continue to operate, writing and reading data to and from your storage volume's local disk, but, data upload to AWS and the taking snapshots come to a halt. If this happens, you must add more working storage to the gateway. However, 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 working storage you provision. In this case, you should then considering optimizing your gateway for better performance (see [Optimizing AWS Storage Gateway Performance \(p. 117\)](#)).

Adding More Working Storage


To add more working storage to a gateway

- Follow the steps in [To configure working storage for your gateway \(p. 71\)](#).

Removing Working Storage

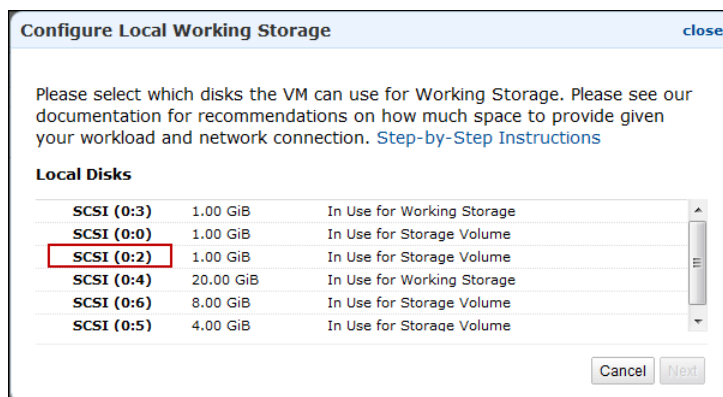
To remove a disk as working storage from a gateway

1. In the AWS Storage Gateway console, in the **Gateway** tab, click the **Configure Local Working Storage** button.



2. In the **Configure Local Working Storage** 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.



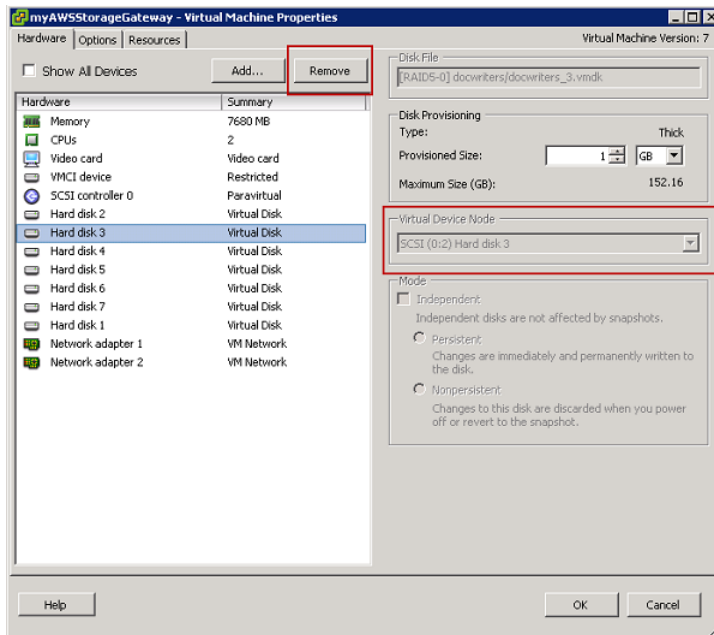
3. Shut down the gateway by following the steps in the [Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console \(p. 99\)](#) procedure.

Note

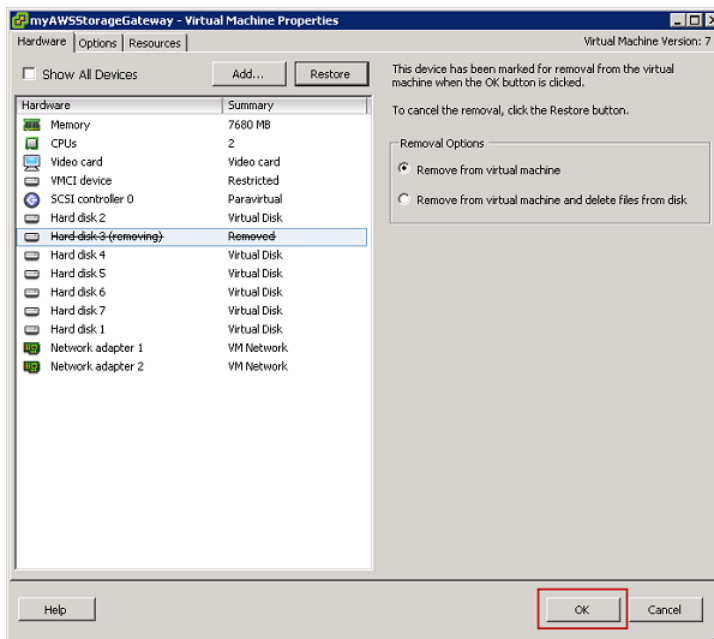
Before shutting down the gateway, ensure that it is not in use by an application that is writing data to it and that no snapshots are progress. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the console. For more information, see [Editing a Snapshot Schedule Using the AWS Storage Gateway Console \(p. 91\)](#).

4. In the vSphere client, right-click the name of your gateway VM and click **Edit Settings...**
5. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, select the working storage disk, 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.



6. Choose an option in the **Removal Options** panel and click **OK** to complete the process of removing the disk.



7. In the AWS Storage Gateway console, turn on the gateway.

Important

When removing a disk used as working storage, you must only add new disks to the VM after you have turned the gateway back on.

8. In the AWS Storage Gateway console, in the **Volumes** tab, check that all storage volumes have a status of [AVAILABLE](#) (p. 68).

After a gateway restart, a storage volume may go through the [PASS THROUGH](#) (p. 68) and [BOOTSTRAPPING](#) (p. 68) states as the gateway adjusts to the working storage disk that you removed. A storage volume that passes through these two states will eventually come to the [AVAILABLE](#) (p. 68) state. You can use a storage volume during the pass through and bootstrapping states; however, you cannot take snapshots.

Managing Your Application Access to Storage Volumes

Topics


- [Connecting from a Windows Client to Your Storage Volume](#) (p. 77)
- [Connecting from a Red Hat Client to Your Storage Volume](#) (p. 78)
- [Configuring CHAP Authentication for Your Storage Volume](#) (p. 80)

After you add local disks to your VM and create storage volumes, the gateway exposes these disks as iSCSI targets. Your client applications must connect to these iSCSI targets to store data. 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 AWS Storage Gateway VM](#) (p. 47). To learn more about creating storage volumes, see [Managing Storage Volumes in AWS Storage Gateway](#) (p. 61).

Note

To connect to your storage volume target, your gateway must have working storage configured. If working storage is not configured for your gateway, then the status of your storage volumes is displayed as **WORKING STORAGE NOT CONFIGURED**. To configure working storage, see [To configure working storage for your gateway](#) (p. 71).

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)).



AWS Storage Gateway User Guide
Connecting from a Windows Client to Your Storage Volume

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.
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 iscsi-initiator-utils 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 5: Access Your AWS Storage Gateway Volumes (p. 32) in the <i>Getting Started</i> tutorial
Connect to your storage volume from Red Hat Linux.	Connecting from a Red Hat Client to Your Storage Volume (p. 78)
Configure CHAP Authentication for Windows and Red Hat Linux.	Configuring CHAP Authentication for Your Storage Volume (p. 80)

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 provides step-by-step instructions about how to connect to your storage volumes. For more information, see [Step 5: Access Your AWS Storage Gateway Volumes \(p. 32\)](#).

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 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.

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. In the registry, navigate to **HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters**.
 - b. Set the **MaxRequestHoldTime** key to 600.

This value represents a hold time of 600 seconds.
2. Increase the disk timeout value.
 - a. In the registry, navigate to **HKLM\System\CurrentControlSet\Services\Disk**.
 - b. Set the **TimeoutValue** to 180.

This value represents a timeout value of 180 seconds.
3. 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-initiatorutils` RPM package to connect to your gateway storage volume.

To connect a Linux client to the storage volume

1. 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
```

2. Ensure that the iSCSI daemon is running.
 - a. Verify that the iSCSI daemon is running using the following command.

AWS Storage Gateway User Guide

Connecting from a Red Hat Client to Your Storage Volume

```
sudo /etc/init.d/iscsi status
```

- b. 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
```

3. 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 \(p. 79\)](#).

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.

AWS Storage Gateway User Guide

Configuring CHAP Authentication for Your Storage Volume

- a. Open the `/etc/iscsi/iscsid.conf` file and find the following line.

```
node.session.timeo.replacement_timeout = [timeout_value]
```

- b. Set the `timeout_value` value to 600.
This value represents a timeout of 600 seconds.

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 180.
This value represents a timeout value of 180 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:

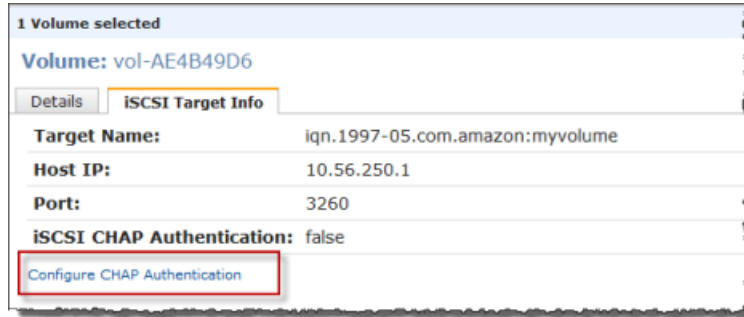
- First, configure CHAP in the AWS Storage Gateway console.
 - [To configure CHAP in the AWS Storage Gateway console \(p. 80\)](#)
- Second, in your client initiator software, complete the CHAP configuration.
 - [To configure mutual CHAP on a Windows client \(p. 82\)](#)
 - [To configure mutual CHAP on a Red Hat Linux client \(p. 87\)](#)

To configure CHAP in the AWS Storage Gateway console

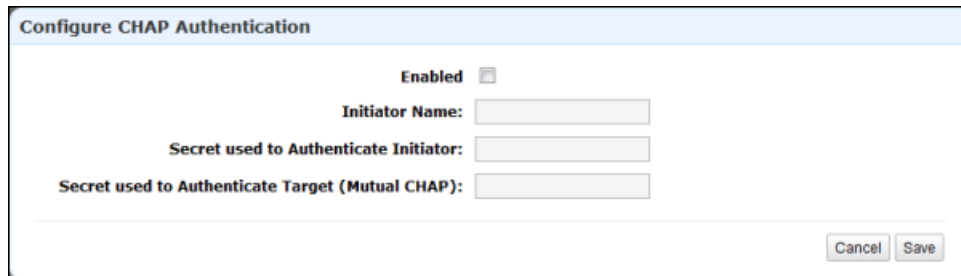
In this procedure, you specify two secret keys that are used to read and write to a storage volume. These same keys are used in the procedure to configure the client initiator.

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** checkbox.
- 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 \(p. 82\)](#).

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 twelve 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 twelve 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.

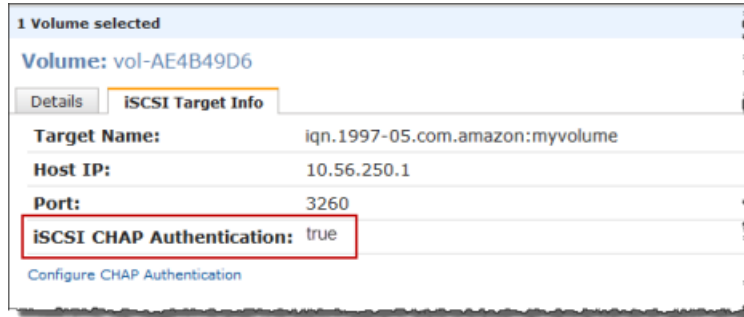
- e. Click **Save**.
- f. Click **Close** in the confirmation dialog box.

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 iSCSI Initiator is not already started, then in the **Start** menu of your Windows client computer, type `iscsiicpl.exe` and run the program.
2. Configure the initiator's (the Windows client) mutual CHAP configuration.
 - a. Click the **Configuration** tab.

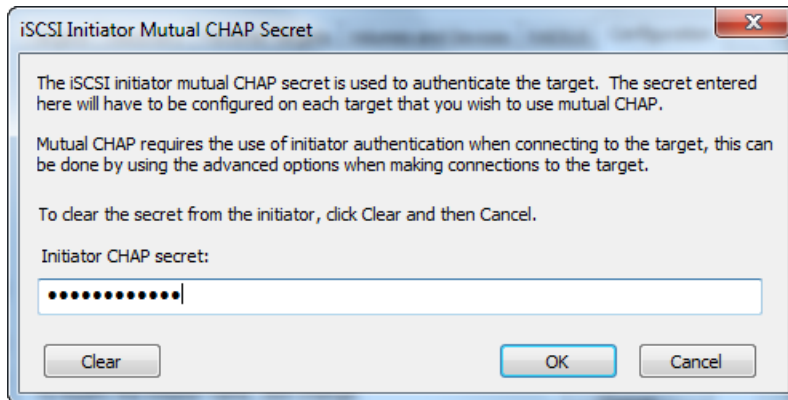


- b. Note the **Initiator Name** field will be unique to your initiator and company. The name shown here is the value that you used in the **Configure CHAP Authentication** dialog box of the AWS Storage Gateway console.

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, 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 \(p. 80\)](#).

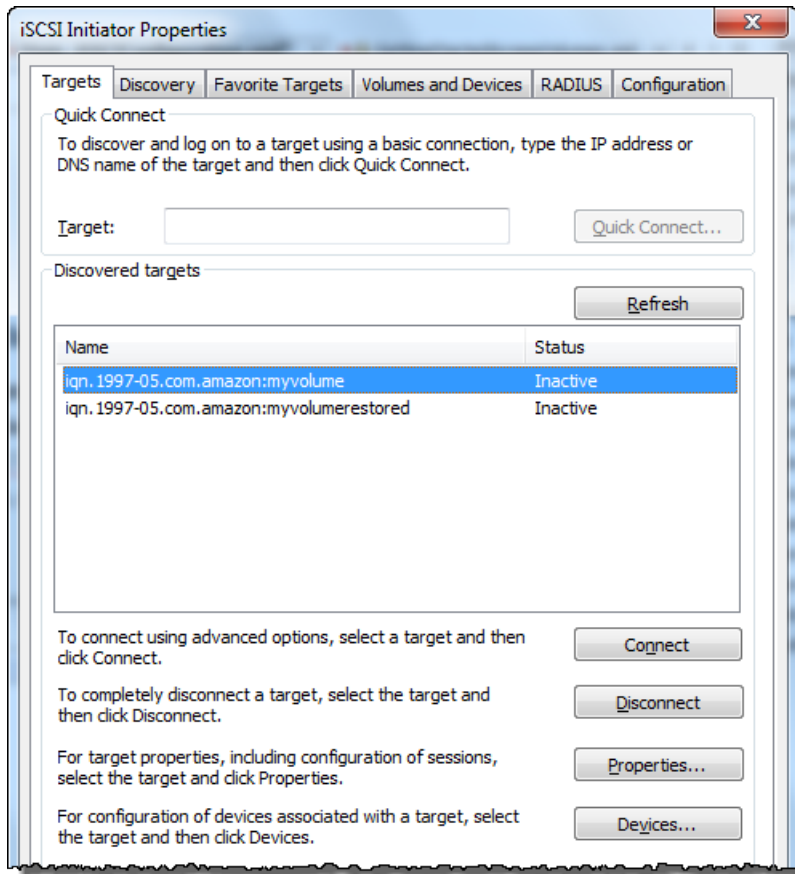
- e. If the key that you enter is less than twelve characters, you get an **Initiator CHAP secret** error dialog box.

Click **OK** and try entering the key again.



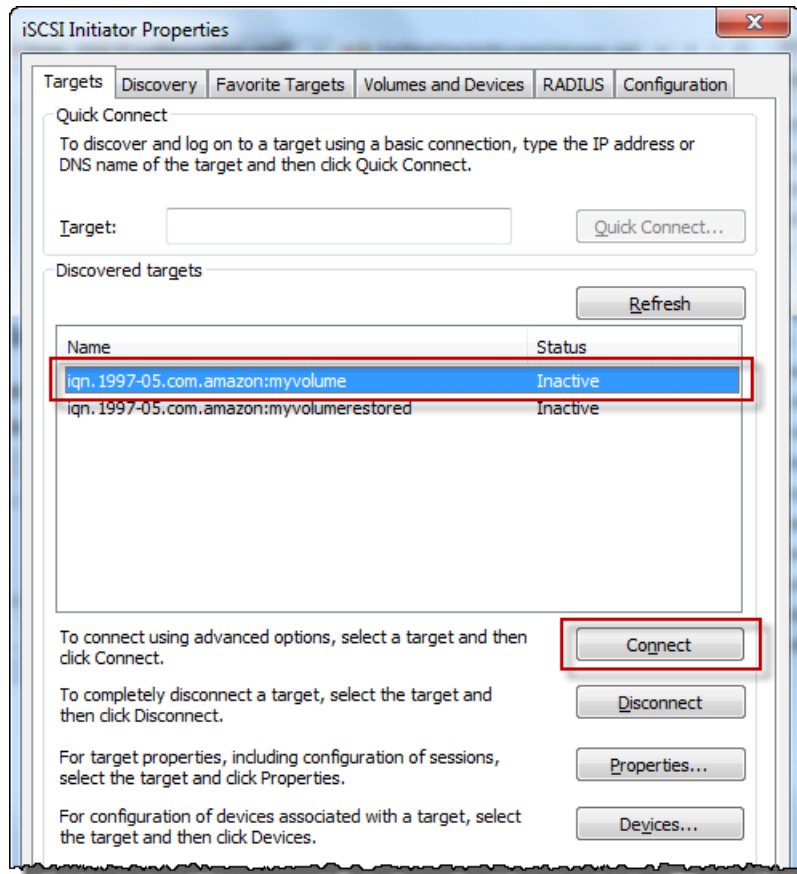
3. Configure the target with the initiator's secret to complete the mutual CHAP configuration.
 - a. Click the **Targets** tab.

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.

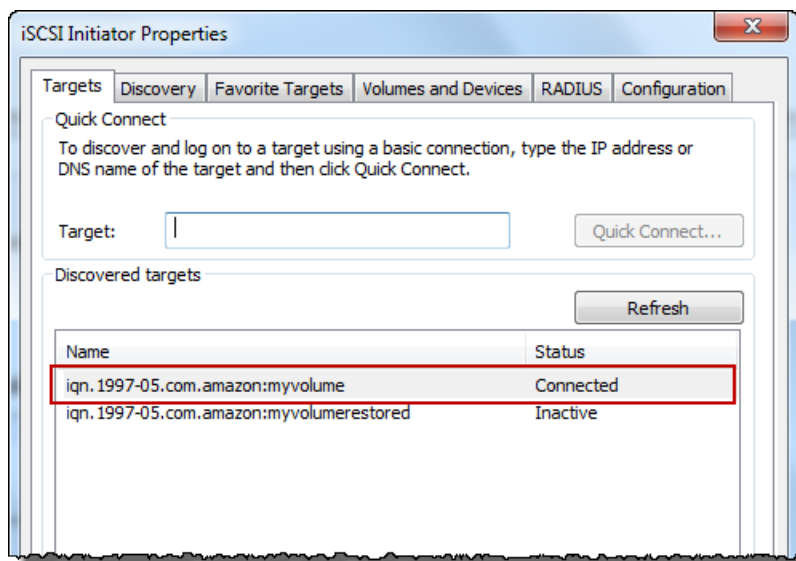
AWS Storage Gateway User Guide
Configuring CHAP Authentication for Your Storage Volume



- 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 \(p. 80\)](#).
 - 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 \(p. 78\)](#).

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.

- 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 \(p. 78\)](#) to discover and log into the target.

Working With Snapshots in the AWS Storage Gateway Console

Topics

- [Overview \(p. 89\)](#)

- [Summary of Snapshot Tasks \(p. 89\)](#)
- [Finding a Snapshot Using the AWS Storage Gateway Console \(p. 90\)](#)
- [Editing a Snapshot Schedule Using the AWS Storage Gateway Console \(p. 91\)](#)
- [Creating an Adhoc Snapshot Using the AWS Storage Gateway Console \(p. 92\)](#)
- [Deleting a Snapshot Using in the AWS Storage Gateway Console \(p. 93\)](#)
- [Restoring a Snapshot Using the AWS Storage Gateway Console \(p. 94\)](#)

Overview

AWS Storage Gateway provides the ability to back up point-in-time snapshots of your data to Amazon S3 for durable recovery. You can take snapshots on a scheduled or ad-hoc basis. You can restore these snapshots to your local gateway storage volumes, or to Amazon EBS volumes enabling access from your Amazon EC2-based applications. 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 gateway storage volume, and restoring a snapshot to an Amazon EBS volume, which can then be attached to an Amazon EC2 instance.

Snapshots are incremental backups, that is, the gateway uploads only the blocks on your volume that have changed since your last snapshot. For example, if you have 100 GB of data and only 5 GB data changed since the last snapshot, then the gateway uploads only the 5 GB 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. Each snapshot has a unique identifier.

When you restore a snapshot to a new storage volume, you can then mount the storage volume as an iSCSI device to your on-premises application server and access the contents of this snapshot. New volumes that are created from existing snapshots load in the background. This means that once a volume is created 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 which hasn't yet been loaded, the gateway will immediately download the requested data from Amazon S3, and will then continue loading the rest of the volume's data in the background.

Snapshot Consistency

Snapshots provide a point-in-time backup 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 like [Diskpart.exe](#). In Linux, use the `mount` and `umount` commands.

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. The tasks summarized here are provided as a reference. Each task is covered in detail in the linked section.

The starting point for all of these tasks is the AWS Storage Gateway console. To work with the tasks, you should have one or more gateways that have been running for enough time so that they have generated snapshots.

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 Using the AWS Storage Gateway Console (p. 90) .
Scheduling	When you first set up a storage 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 Using the AWS Storage Gateway Console (p. 91) .
Creating	Snapshots are automatically created by default and you can change the schedule of the snapshot. However, you can take an instantaneous snapshot at any time. For more information, see Creating an Adhoc Snapshot Using the AWS Storage Gateway Console (p. 92) .
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. 95) and Restoring a Snapshot to an Amazon EBS Volume (p. 97) .
Deleting	If you don't need a snapshot anymore, you can delete the snapshot. 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 Using in the AWS Storage Gateway Console (p. 93) .

Finding a Snapshot Using the AWS Storage Gateway Console

You can list all your snapshots in the **AWS Storage Gateway** console. The list includes all your snapshots generated from your gateway and snapshots that you might have generated from Amazon EBS. For each snapshot, the console shows snapshot 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

1. On the **AWS Storage Gateway** console, in the **Navigation** pane, click **Snapshots**.

The **Snapshots** window shows a list of your snapshots.

The screenshot shows the AWS Storage Gateway console interface. On the left is a navigation pane with 'Snapshots' selected. The main area displays a table of snapshots. The selected snapshot, 'snap-5d6b8e3e', is highlighted in blue. Below the table, a details panel for this snapshot is shown, including its ID, status, volume, and start time.

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-a6acb6c6	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. Click the snapshot row to display the snapshot details.

Editing a Snapshot Schedule Using the AWS Storage Gateway Console

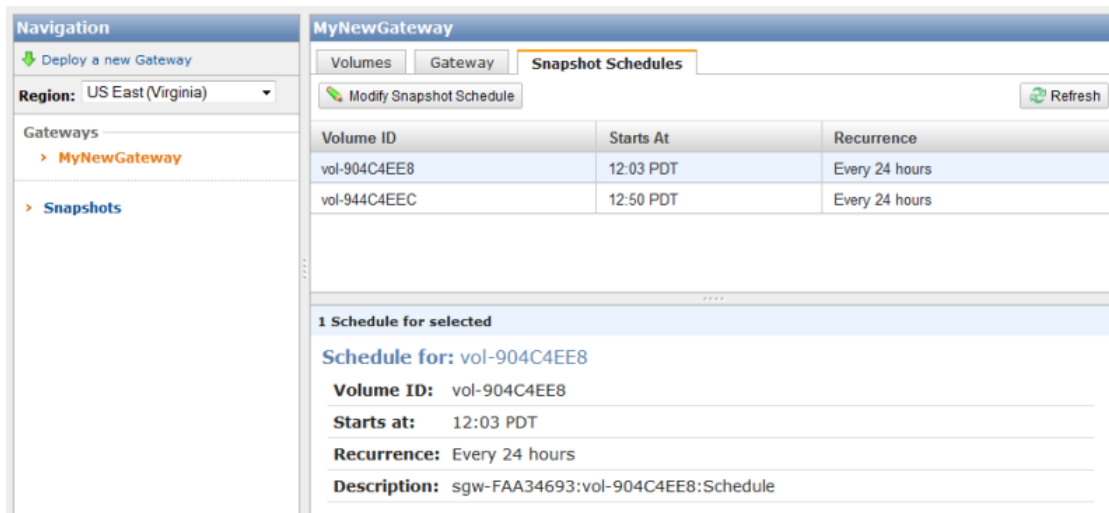
For each storage volume you create on your gateway to store your application data, AWS Storage Gateway creates a default snapshot schedule of once a day. To ensure that your gateways can keep up with the rate of incoming writes on your local storage volumes, you must maintain a snapshot schedule, at least once a day set by default.

You can optionally modify the default 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). In the following steps, we show you how to edit the snapshot schedule of a storage volume.

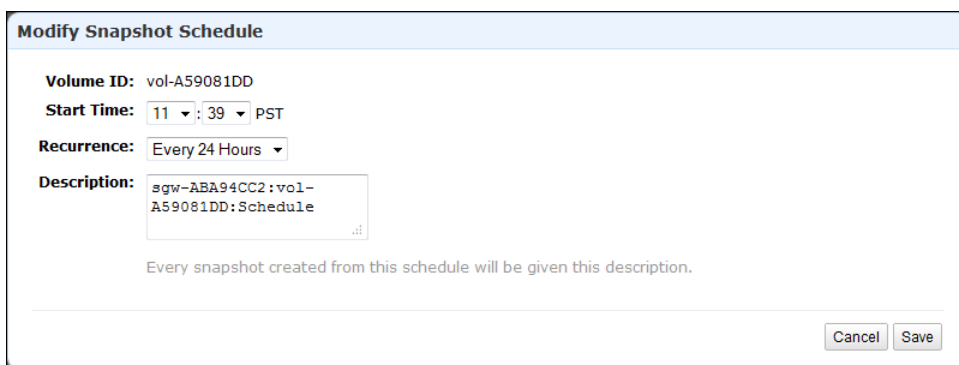
To edit snapshot schedule

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.

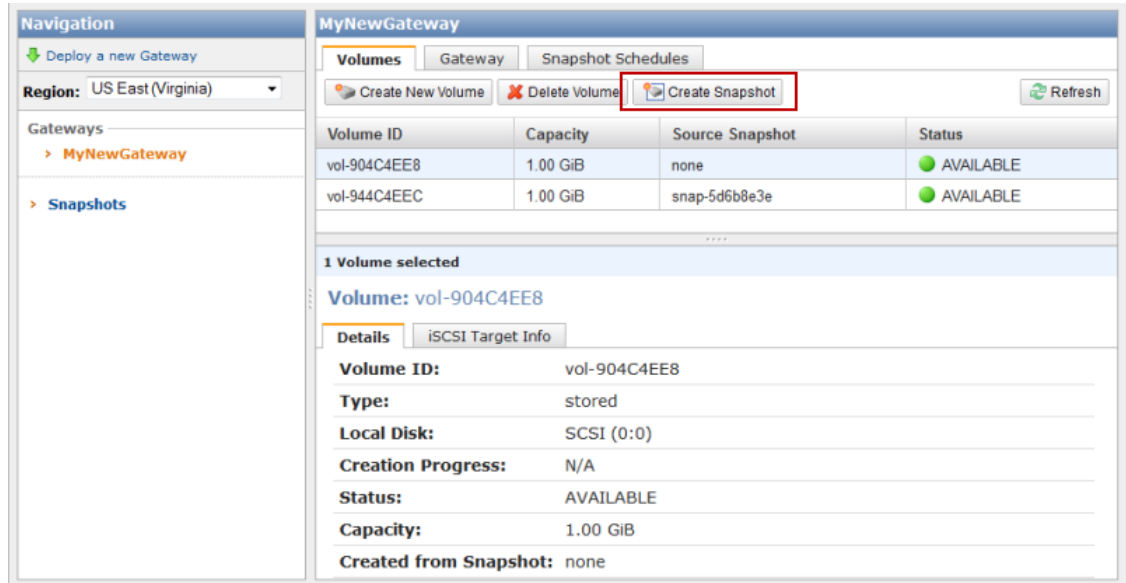
Creating an Adhoc Snapshot Using the AWS Storage Gateway Console

In addition to scheduled snapshots, **AWS Storage Gateway** allows you to take adhoc snapshots, enabling you to back up your storage volume immediately without waiting for the next scheduled snapshot.

To take an adhoc 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 Using the AWS Storage Gateway Console](#) (p. 90).

Deleting a Snapshot Using in the AWS Storage Gateway Console

You can delete a snapshot using the **AWS Storage Gateway** console. For example, there might be several snapshots taken of your storage volume over a period of time and you might want to delete old snapshots.

To delete a snapshot

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 AWS Storage Gateway console interface. On the left is a navigation pane with 'Snapshots' selected. The main area displays a table of snapshots. The snapshot 'snap-5d6b8e3e' is selected, and its details are shown below the table.

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-a6acb6c6	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. Select the snapshot that you want to delete and click **Delete**.

This screenshot is identical to the previous one, but the 'Delete' button (marked with a red 'X') is highlighted with a red box, indicating the next step in the process.

3. Clicking **OK** to confirm that you want to delete the snapshot.

Restoring a Snapshot Using the AWS Storage Gateway Console

Topics

- [Restoring a Snapshot to an AWS Storage Gateway Volume \(p. 95\)](#)
- [Restoring a Snapshot to an Amazon EBS Volume \(p. 97\)](#)

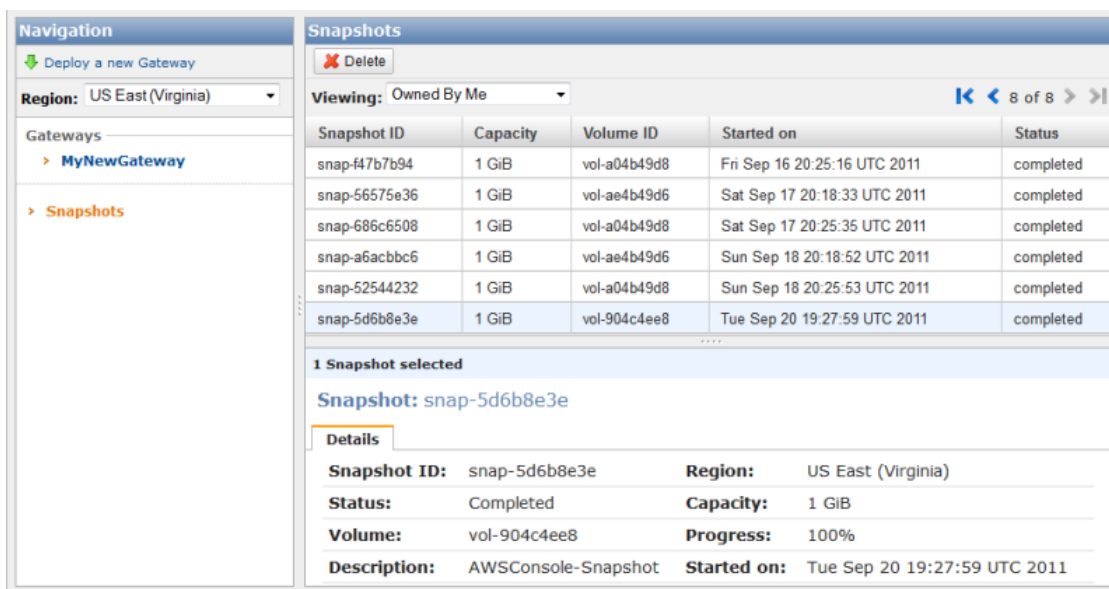
You can restore the snapshot locally to a new AWS Storage Gateway volume, or you can use the snapshot to create an Amazon Elastic Block Store (EBS) volume and attach this volume to an Amazon EC2 instance.

Restoring a Snapshot to an AWS Storage Gateway Volume

In the event that you lose your storage volume, for example when a disk crash occurs, you can restore a saved snapshot to a new storage volume using the **AWS Storage Gateway** console. When you restore a snapshot to a new storage volume, you can then mount the storage volume as an iSCSI device to your on-premises application server and access the contents of this snapshot. New volumes created from existing snapshots load in the background. This means that once a volume is created 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 which hasn't yet been loaded, the volume will immediately download the requested data from Amazon S3, and then will continue loading the rest of the volume's data in the background.

To create a storage volume from an existing snapshot

1. On the **AWS Storage Gateway** console, click **Snapshots**.



The screenshot shows the AWS Storage Gateway console interface. On the left is a navigation pane with 'Snapshots' selected. The main area displays a table of snapshots. The selected snapshot, 'snap-5d6b8e3e', is highlighted. Below the table, the details for this snapshot are shown, including its ID, region, status, capacity, volume ID, progress, and description.

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-a6acb9c6	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.

AWS Storage Gateway User Guide

Restoring a Snapshot

Navigation

Deploy a new Gateway

Region: US East (Virginia)

Gateways

- MyNewGateway
- 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-a6acb6c6	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

- In the **Navigation** pane, select the gateway that to which you want to restore the snapshot.
- Click **Create Volume**.

Navigation

Deploy a new Gateway

Region: US East (Virginia)

Gateways

- MyNewGateway
- Snapshots

MyNewGateway

Volumes Gateway Snapshot Schedules

Create New Volume Delete Volume Create Snapshot Refresh

Volume ID	Capacity	Source Snapshot	Status
vol-904C4EE8	1.00 GiB	none	AVAILABLE
vol-944C4EEC	1.00 GiB	snap-5d6b8e3e	AVAILABLE

1 Volume selected

Volume: vol-904C4EE8

Details iSCSI Target Info

Volume ID: vol-904C4EE8

Type: stored

Local Disk: SCSI (0:0)

Creation Progress: N/A

Status: AVAILABLE

Capacity: 1.00 GiB

Created from Snapshot: none

- In the dialog box, paste the Snapshot ID you copied above into the **Based on Snapshot ID** field.

Create Storage Volume close

Disk: SCSI (0:2) Preserve existing data

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

Based on Snapshot ID: snap-5d6b8e3e

Size: 1 GiB

Host IP: 10.56.250.1

Port: 3260

Cancel Create Volume

- 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 storage volume, see [Adding Local Disks to AWS Storage Gateway for Storing Your Application Data](#) (p. 47). You can now access the contents of this volume from your on-premises applications (see [Managing Your Application Access to Storage Volumes](#) (p. 76)).

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. Therefore, 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*.

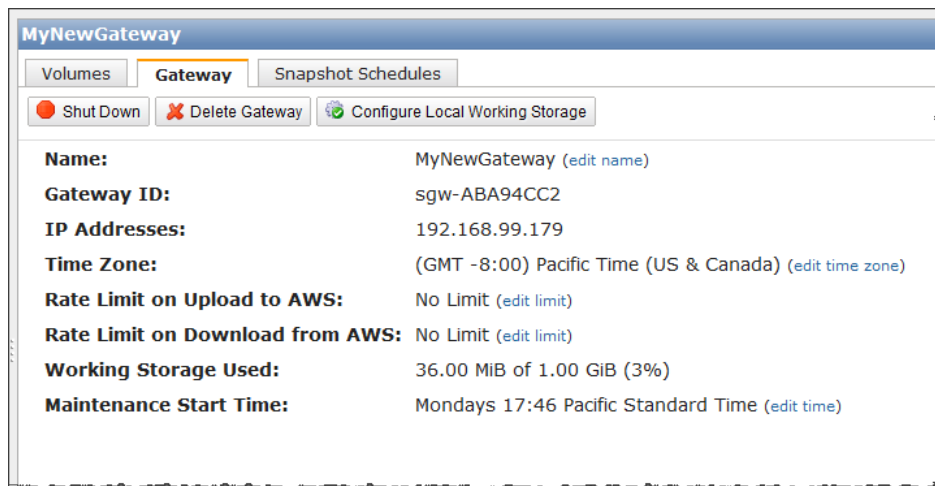
Performing Maintenance Tasks in AWS Storage Gateway

Topics

- [Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console](#) (p. 99)
- [Managing Gateway Updates Using the AWS Storage Gateway Console](#) (p. 100)
- [Deleting a Gateway Using the AWS Storage Gateway Console](#) (p. 101)
- [Logging into Your AWS Storage Gateway Local Console](#) (p. 102)
- [Routing AWS Storage Gateway Through a Proxy](#) (p. 104)
- [Configuring Your AWS Storage Gateway to Use Static IP Addresses](#) (p. 105)
- [Test Your AWS Storage Gateway Connection to the Internet](#) (p. 108)
- [Configuring AWS Storage Gateway for Multiple Network Adapters \(NICs\)](#) (p. 109)
- [Creating a Storage Volume in AWS Storage Gateway with Multiple Network Adapters](#) (p. 112)

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.

AWS Storage Gateway User Guide Performing Maintenance Tasks



The table below summarizes the updatable fields on the **Gateway** tab. Click the **edit** link at the end of 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. 119)), you might want to be aware of the previous name and the new name to avoid confusion or 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.
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.
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.
Working Storage Used	Displays the working storage used. For information about how to monitor working storage and how it changes in time, see Monitoring AWS Storage Gateway's Working Storage (p. 124) .
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 timezone of the gateway.

Maintenance Item	Comments
Apply Update Now	If there is an update for your AWS Storage Gateway, a message appears in the console. Click the Apply Update Now button to apply the update immediately. If you do not apply the update, AWS Storage Gateway applies the update based on your Maintenance Start Time . For more information, see Managing Gateway Updates Using the AWS Storage Gateway Console (p. 100).

Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console

This section discusses shutting down a gateway. You might need to shut down your gateway, for example, to apply a patch to your hypervisor host. While a gateway is shutdown, 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 Using the AWS Storage Gateway Console](#) (p. 91).

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 vSphere client 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 **AWS Storage Gateway** console **Navigation** pane, select the gateway.
2. Click the **Shut Down** button.



3. In the confirmation dialog, click **Okay**.

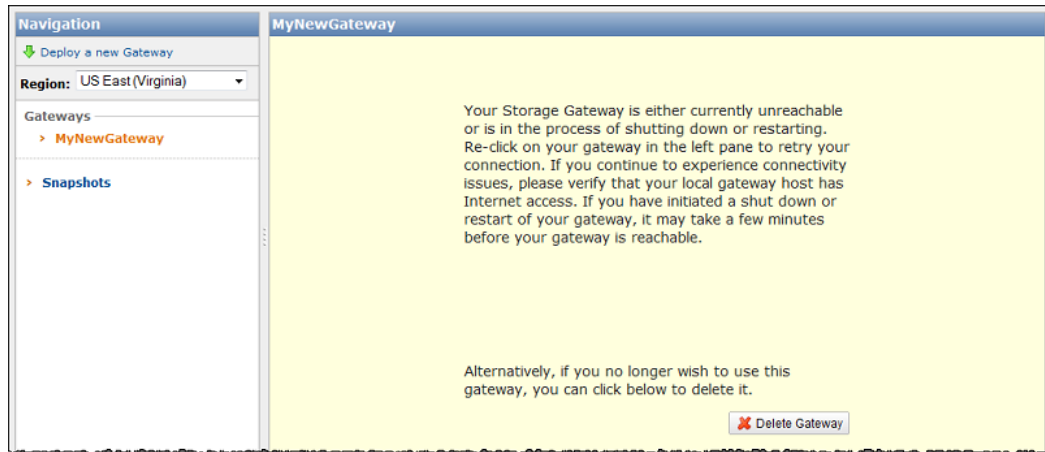


4. In the **Gateway Shut Down** dialog box, click **Close**.



5. While the gateway is shutting down you may see a message that your gateway is in the process of shutting down.

You have the option of deleting the gateway at this point. Do not delete the gateway if you plan to restart the gateway and continue working with it.



6. Select the gateway in the left navigation pane.

A **Restart** button is displayed.



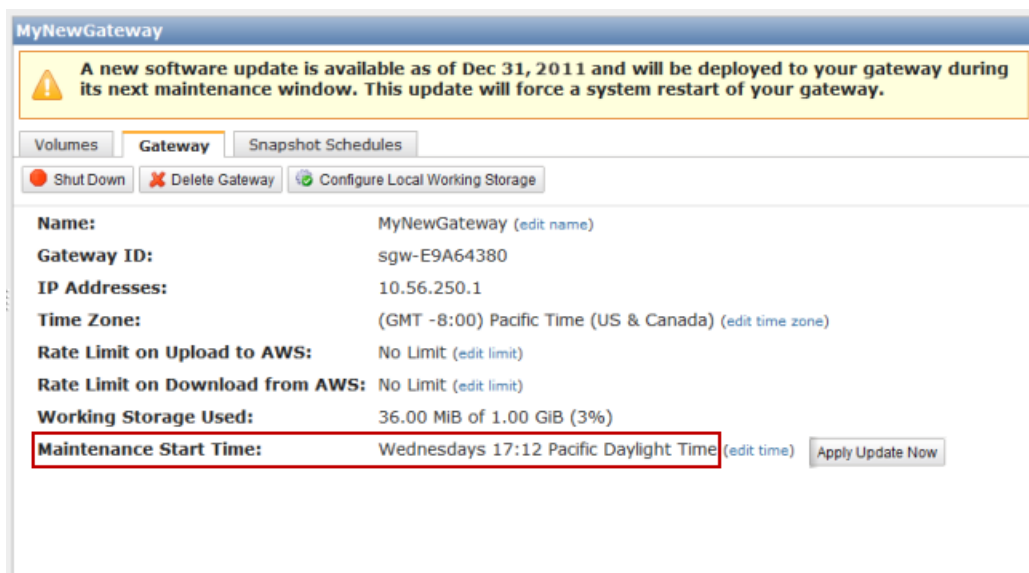
Managing Gateway Updates Using the AWS Storage Gateway Console

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. Depending on the update, AWS Storage Gateway may need to force a system restart of your gateway; however, this should generally take only a few minutes to complete. Updates which require a restart occur infrequently.

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 the **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 \(p. 78\)](#) and [Customizing Your Linux iSCSI Settings \(p. 79\)](#), respectively.



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.

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 gateway (see [Activating AWS Storage Gateway \(p. 57\)](#)) and configure it to match the storage volumes and working storage of the deleted gateway. When you create storage volumes on the new gateway, you can use the same underlying disks as the deleted gateway and preserve the data on the disks. For more information, see [To create a storage volume \(p. 63\)](#).

To delete a gateway

1. In the **AWS Storage Gateway** console **Navigation** pane, 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 the gateway 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 the VMware ESXi host environment. To remove the VM, use the VMware vSphere client to connect to the host and remove the VM.

Logging into Your AWS Storage Gateway Local Console

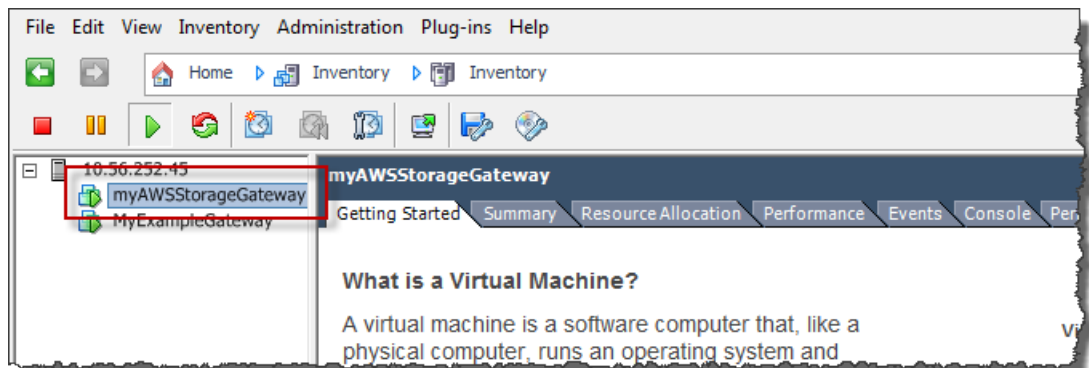
Some gateway maintenance tasks require that you log in to your gateway's local console. The local console is accessible through your VMware client. The user is *sguser* and the password is *sgpassword*. These login credentials give you access to configuration menus, where you can configure gateway network settings.

To route your gateway Internet traffic through a local proxy server

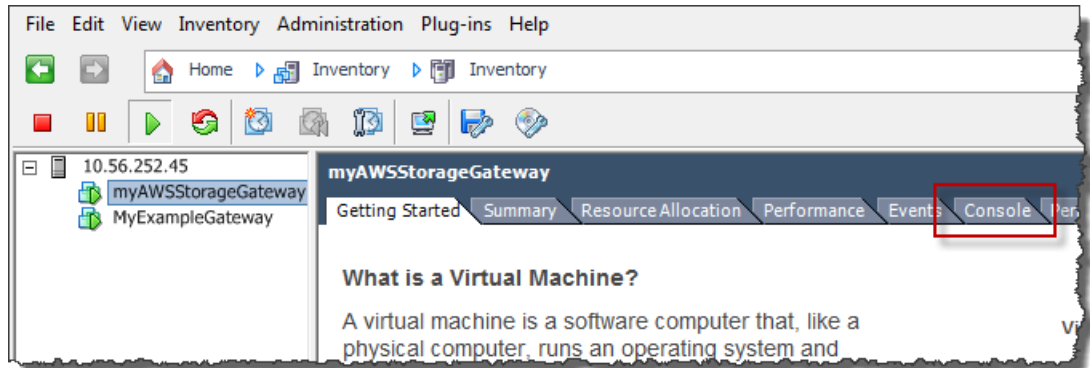
1. In the VMware Vsphere client, select your gateway VM.
2. Ensure that the gateway is powered on.

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 out by clicking the green **Power On** icon in the **Toolbar** menu.



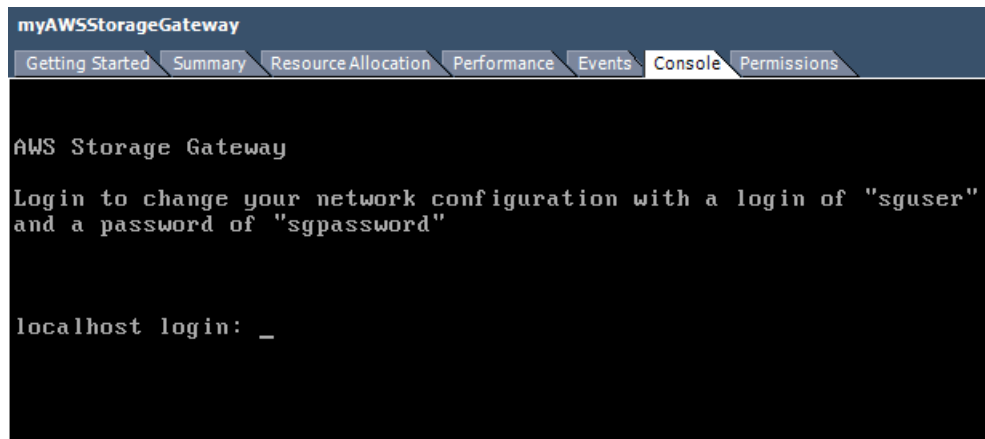
- Click the **Console** tab.



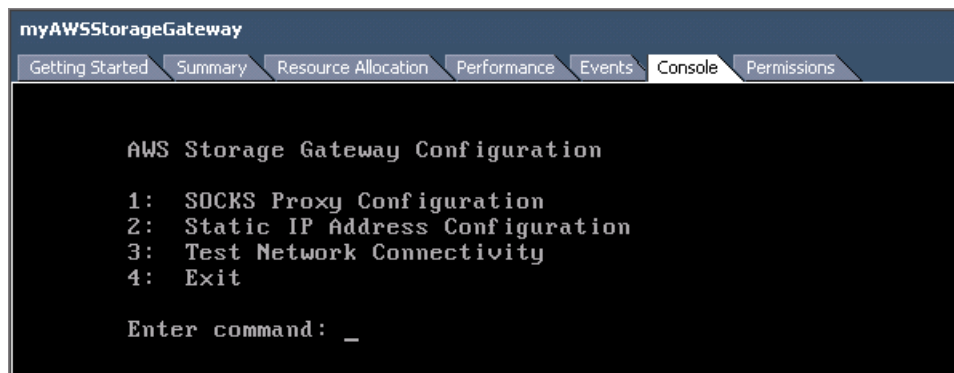
- In the login screen, log in to the VM with the user name and password provided in the **Console** tab.

Note

To release the cursor from the console window, press CTRL + ALT.



- After you log in, you will see the **AWS Storage Gateway Configuration** main menu.



To...	See...
Configure a SOCKS proxy for your gateway	Routing AWS Storage Gateway Through a Proxy (p. 104)

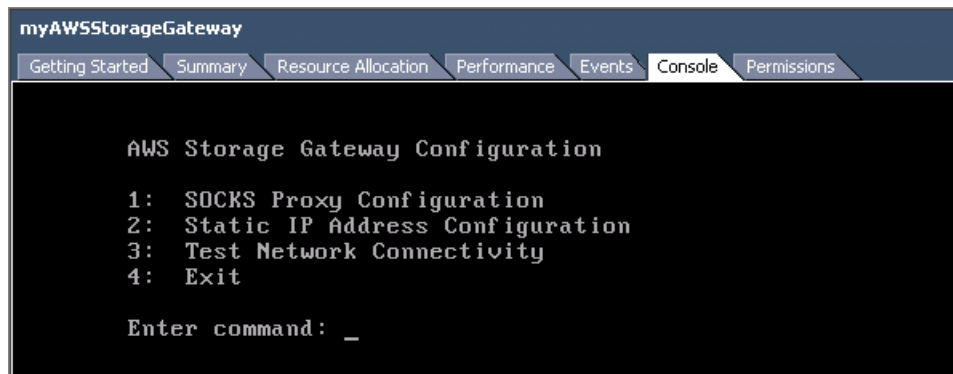
To...	See...
Configure static IP addresses for your gateway's interfaces	Configuring Your AWS Storage Gateway to Use Static IP Addresses (p. 105)
Test network connectivity	Test Your AWS Storage Gateway Connection to the Internet (p. 108)

Routing AWS Storage Gateway Through a Proxy

The AWS Storage Gateway supports the configuration of a SOCKS proxy between your gateway and AWS. If your gateway must use a proxy server to communicate to the internet, then you need to configure 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.

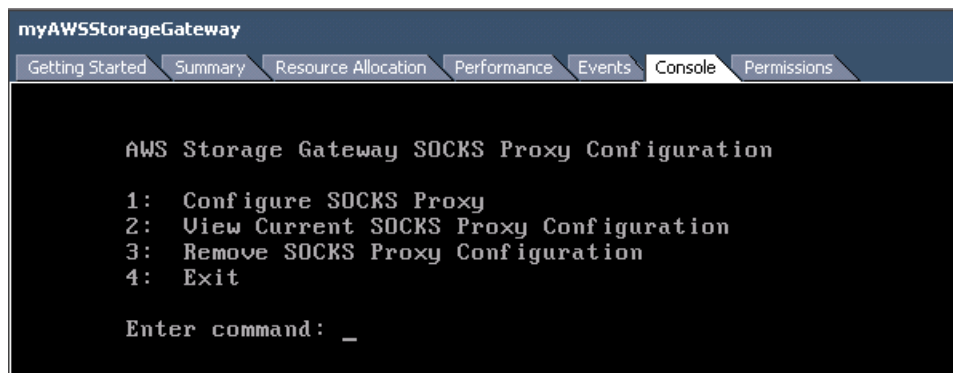
To route your gateway internet traffic through a local proxy server

1. Log into your gateway's local console. For instructions, see [Logging into Your AWS Storage Gateway Local Console \(p. 102\)](#).
2. In the **AWS Storage Gateway Configuration** main menu enter option 1.



```
myAWSStorageGateway
Getting Started Summary Resource Allocation Performance Events Console Permissions
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:



```
myAWSStorageGateway
Getting Started Summary Resource Allocation Performance Events Console Permissions
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: _
```

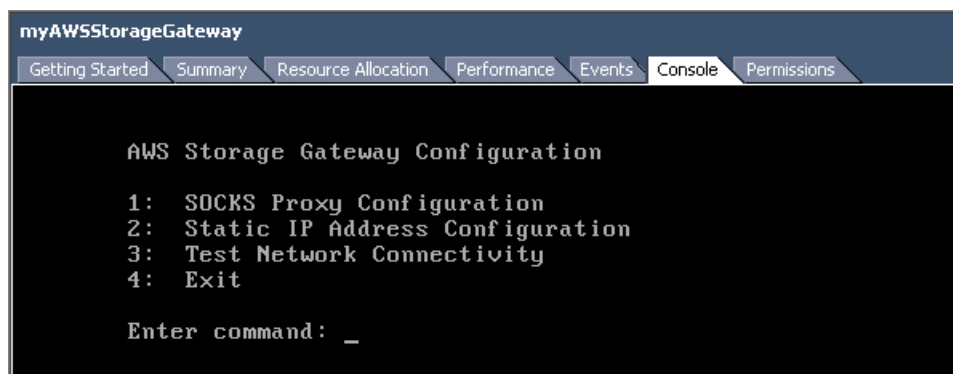
To...	Do this...
Configure a SOCKS proxy	Enter option 1. You will need to supply a host name and port to complete configuration.
View the current SOCKS proxy configuration	Enter option 2. If a SOCKS proxy is not configured the message "SOCKS Proxy not configured" is displayed. If a SOCKS proxy is configured 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 Static IP Addresses

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 \(p. 102\)](#).
2. In the **AWS Storage Gateway Configuration** main menu, select option 2.



```
myAWSStorageGateway
Getting Started Summary Resource Allocation Performance Events Console Permissions
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:

```
myAWSStorageGateway
Getting Started Summary Resource Allocation Performance Events Console Permissions

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</p> <p>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. 99).</p> <p>If your gateway uses more than one network interface, then all enabled interfaces must be set 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 then you 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</p> <p>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. 99).</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 .

Test 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.

To test your gateway's connection to the Internet

1. Log in to your gateway's local console. For instructions, see [Logging into Your AWS Storage Gateway Local Console \(p. 102\)](#).
2. In the **AWS Storage Gateway Configuration** main menu, select option **3**.

```

myAWSStorageGateway
Getting Started Summary Resource Allocation Performance Events Console Permissions
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

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. Scenarios where you would want to configure a gateway to use multiple network adapters include maximizing throughput to a gateway when network adapters are the bottleneck or separating your applications and how they write to a gateway's storage volumes. In the latter scenario, you might choose, for example, to have a critical storage application exclusively use one of one adapters defined for a gateway.

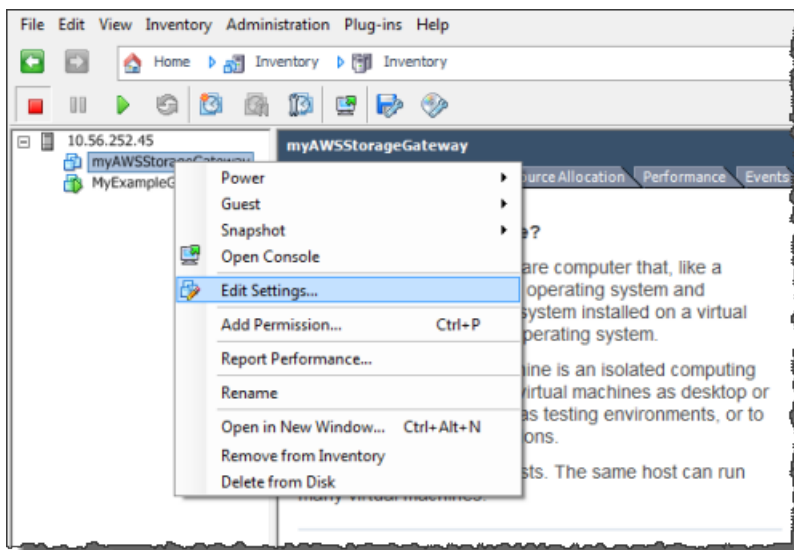
The following procedure assumes that your gateway VM already has one network adapter defined and you will add a second adapter.

To configure your gateway to use an additional network adapter

1. In the VMware Vsphere client, select your gateway VM.

The VM can remain powered on for this procedure.


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



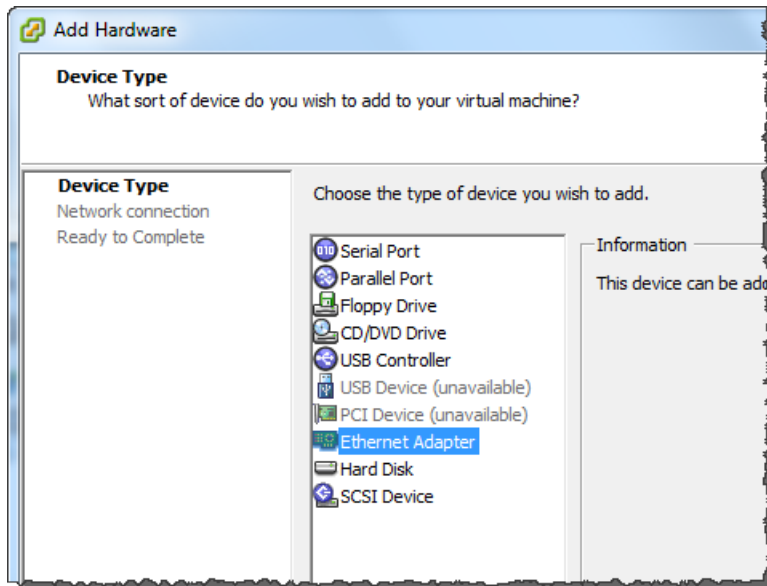
3. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, click **Add...** to add a device.

AWS Storage Gateway User Guide

Configuring Your Gateway for Multiple Network Adapters (NICs)



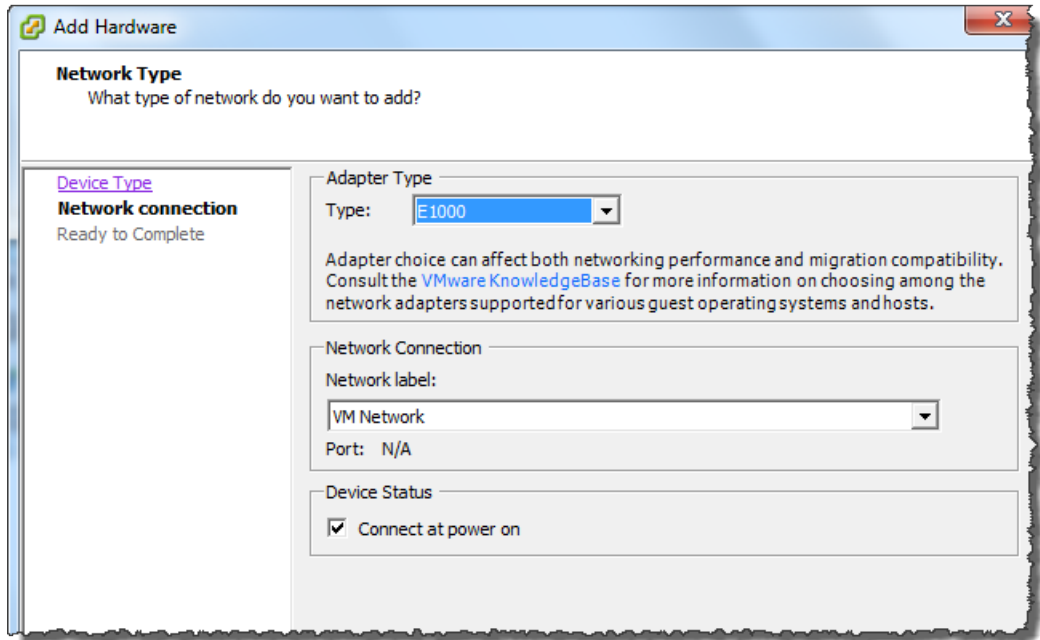
4. Follow the Add Hardware wizard to add a network adapter:
 - a. In the **Device Type** pane, click **Ethernet Adapter** to add an adapter, and click **Next**.



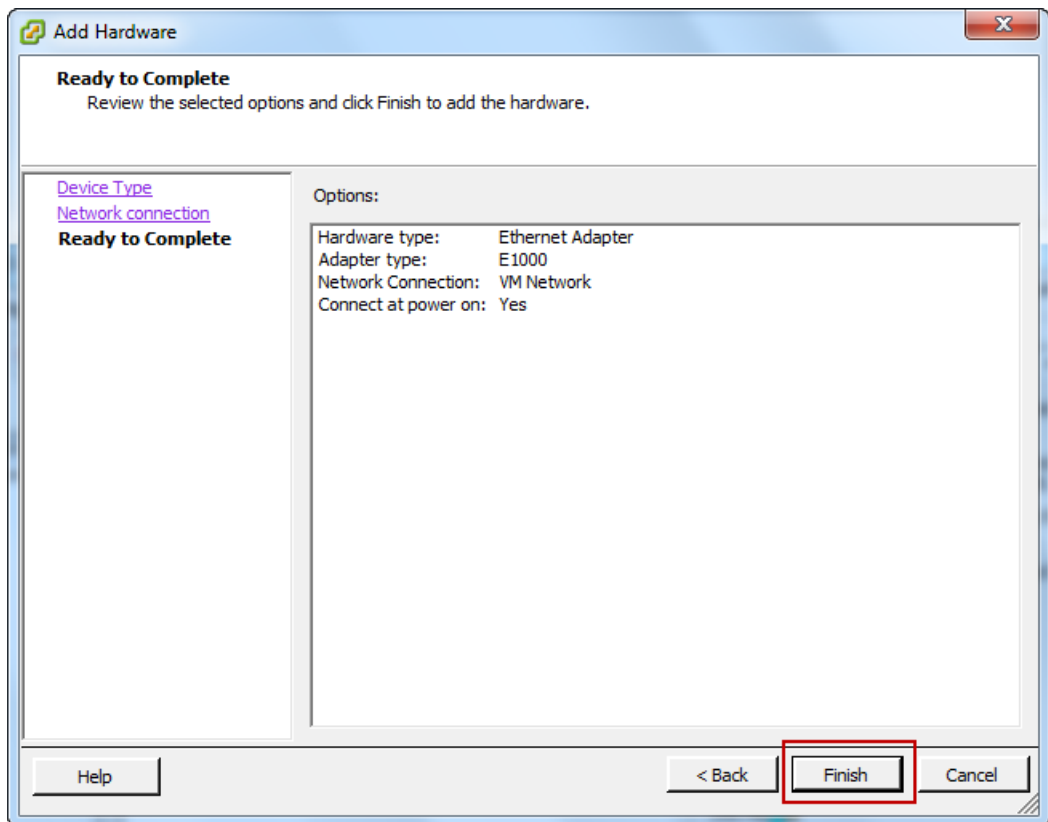
- b. In the **Network Type** pane, in the **Type** drop-down list, select an adapter type, ensure that **Connect at power on** is selected, and click **Next**.

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](#).

AWS Storage Gateway User Guide
Configuring Your Gateway for Multiple Network Adapters (NICs)



- c. In the **Ready to Complete** pane, review the information and click **Finish**.



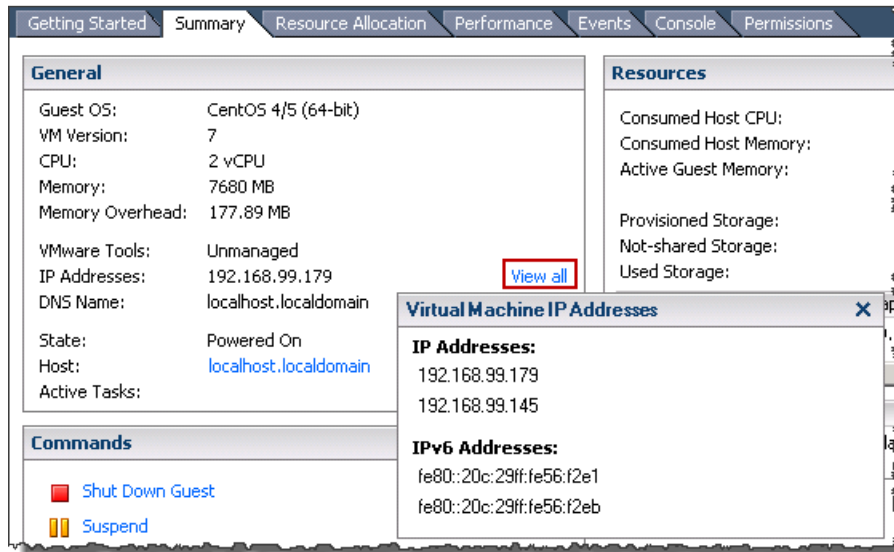
5. Click the **Summary** tab of the VM, and click **View All** next to the **IP Address** field. A Virtual Machine IP Addresses window displays all the IP addresses you can use to access the gateway. Confirm that a second IP address is listed for the gateway.

AWS Storage Gateway User Guide

Creating a Storage Volume on a Gateway with Multiple Network Adapters

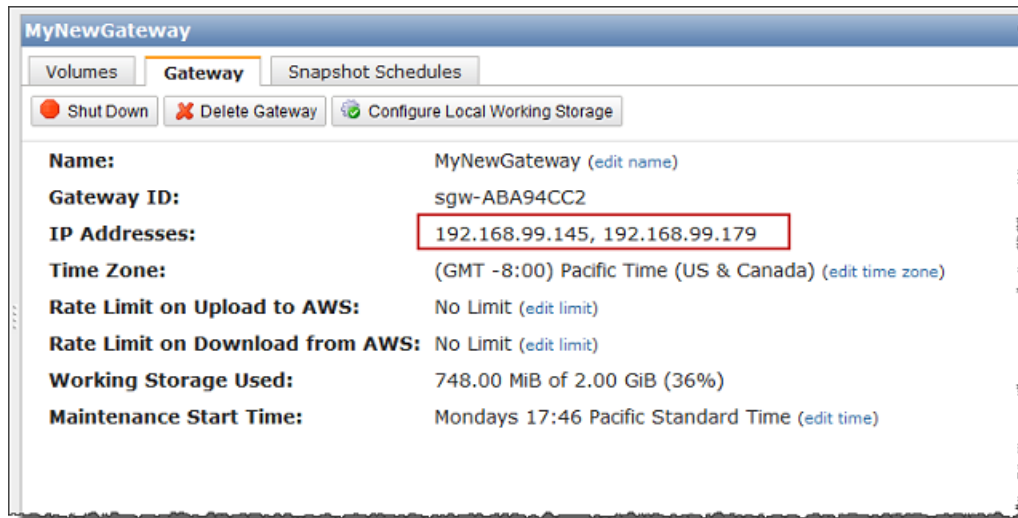
Note

It might take several moments for the adapter changes to take effect and the VM summary information to refresh.



6. 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 what 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. 109).

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, [To create a storage volume](#) (p. 63).
4. Select an IP address to use to access the volume.

Note that the **Create Storage Volume** dialog displays a drop-down list for **Host IP**, one IP address per adapter configured for the gateway VM. If the gateway VM is only configured for 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. It has a title bar with 'Create Storage Volume' and a 'close' button. The form contains the following fields and options:

- Disk:** A dropdown menu set to 'SCSI (0:4)' and a checkbox for 'Preserve existing data'.
- iSCSI Target Name:** A text input field containing 'iqn.1997-05.com.amazon:' and a sub-field containing 'myvolume'.
- Based on Snapshot ID:** An empty text input field.
- Size:** A text input field containing '20 GiB'.
- Host IP:** A dropdown menu with '192.168.99.145' selected. This field is highlighted with a red rectangular box.
- Port:** A text input field containing '3260'.

At the bottom right, there are two buttons: 'Cancel' and 'Create Volume'.

5. Click **Create Volume**.

To create a connection to the storage volume, see [Managing Your Application Access to Storage Volumes](#) (p. 76).

Troubleshooting in AWS Storage Gateway

This section discusses troubleshooting gateway-related and storage volume-related issues.

Troubleshooting Gateway Issues

The following table lists typical issues that you might encounter working with gateways.

Scenario	Action to Take
Your gateway VM's IP address doesn't appear in the vSphere client, Summary tab.	<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 IP address to appear in the Summary tab of your vSphere client. For more information, see Activating AWS Storage Gateway (p. 57)

Scenario	Action to Take
<p>Your gateway's activation fails when you click the Proceed to Activation button in the AWS Storage Gateway console.</p>	<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. 104). • Check that the host has the correct time and is configured to synchronize its time automatically to a Network Time Protocol (NTP) server. For more information, see Synchronize VM Time with the Host Time (p. 14). • Check that the VM is time-synchronized. For more information, see Synchronize VM Time with the Host Time (p. 14). • 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 4.5 GBs of RAM. Gateway allocation fails if there is less than 4.5 GBs of RAM. For more information, see Requirements (p. 4).
<p>You allocated a local disk to your VM, but are not able to see it in the AWS Storage Gateway console.</p>	<p>Check the paravirtualization setting of the VM. If you don't see the disk, it may be because paravirtualization is turned off. To turn paravirtualization on for a gateway that has already been activated, use the following steps:</p> <ul style="list-style-type: none"> • Delete your gateway's storage volumes (as they will go into an IRRECOVERABLE (p. 68) state). • Shut down your gateway in the AWS Storage Gateway console. For more information, see Shutting Down and Turning On a Gateway Using the AWS Storage Gateway Console (p. 99). • Power off the gateway VM in the vSphere client. • Turn on paravirtualization. For more information, see Configure AWS Storage Gateway VM to Use Paravirtualization (p. 56). • Power on the gateway VM. • Turn on the gateway in the AWS Storage Gateway console.
<p>You need to remove a working storage disk because you want to reduce the amount of working storage for a gateway or you need to replace a disk used as working storage that has failed.</p>	<p>For step-by-step instructions about removing a working storage disk, see Removing Working Storage (p. 74).</p>

Scenario	Action to Take
You need to improve bandwidth between your gateway and AWS.	You can improve the bandwidth from your gateway to AWS by setting up your internet connection to AWS on a separate NIC than 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. 122)). Improving the internet connectivity helps to ensure that your working storage does not fill up.
Throughput to or from your gateway drops to zero.	<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 that appear in the IP Addresses field are the same as those that appear in the vSphere client on the Summary tab. 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. 99). After the restart, the IP Addresses field in the Gateway tab of the AWS Storage Gateway console should match the IP addresses shown in the vSphere client on the Summary tab for the gateway. • Check your gateway's connectivity to AWS as described in Test Your AWS Storage Gateway Connection to the Internet (p. 108). • 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 Static IP Addresses (p. 105) 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. 122).</p>

Troubleshooting Storage Volume Issues

The following table lists the most typical issues that you might encounter working with storage volumes.

Scenario	Action to Take
The AWS Storage Gateway console indicates that your volume has a status of WORKING STORAGE NOT CONFIGURED (p. 68).	Add working storage. You cannot use a gateway to store your application data if working storage for the gateway is not configured. For more information, see Resources: To configure working storage for your gateway (p. 71) .

AWS Storage Gateway User Guide
Troubleshooting Storage Volume Issues

Scenario	Action to Take
<p>The AWS Storage Gateway console indicates that your volume has a status of IRRECOVERABLE (p. 68).</p>	<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. 61). Delete pending snapshots of the volume before deleting the storage volume. For more information, see Deleting a Snapshot Using in the AWS Storage Gateway Console (p. 93).</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>
<p>The AWS Storage Gateway console indicates that your volume has a status of PASS THROUGH (p. 68).</p>	<p>A volume can be in PASS THROUGH (p. 68) for several reasons. Some of the reasons are a cause for action and some are not.</p> <p>An example of when your storage volume is in PASS THROUGH and you should take action is because your gateway has run out of working storage. To verify if your working storage was exceeded in the past you can view the <code>WorkingStoragePercentUsed</code> metric in the Amazon CloudWatch console (see Monitoring AWS Storage Gateway's Working Storage (p. 124)). If your gateway is in PASS THROUGH because it has run out of working storage, you should allocate more working storage to your gateway. Adding more working storage will cause your storage volume to transition from PASS THROUGH to BOOTSTRAPPING (p. 68) to AVAILABLE (p. 68) automatically. During BOOTSTRAPPING, the gateway reads data off the storage volume's disk and uploads this data to Amazon S3 and catches up as needed. Once, the gateway is 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 on adding more working storage, see Ongoing Management of Working Storage for a Gateway (p. 73).</p> <p>To take action before working storage is exceeded, you can set a threshold alarm on a gateway's working storage. For more information, see To set an upper threshold alarm for a gateway's working storage (p. 125).</p> <p>Another example of when a storage volume is in PASS THROUGH but you do not need to take action 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 working storage disk has failed. In this is the case, you should remove the disk. For more information, see Removing Working Storage (p. 74).</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 working storage for the gateway. For more information, see To configure working storage for your gateway (p. 71).</p>

Scenario	Action to Take
You want to change the iSCSI target name of your storage volume.	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 To create a storage volume (p. 63) .
Your scheduled snapshot of a storage volume did not occur.	Check if your volume is in PASS THROUGH (p. 68) , or if the gateway's working storage was filled just prior to the time the snapshot was scheduled to be taken. You can check the <code>WorkingStoragePercentUsed</code> metric for the gateway in the Amazon Cloudwatch console and create an alarm for it. For more information, see Monitoring AWS Storage Gateway's Working Storage (p. 124) and To set an upper threshold alarm for a gateway's working storage (p. 125) .
You need to remove a storage volume because it isn't needed or you need to replace a storage volume disk that has failed.	For step-by-step instructions about removing a storage volume, see To remove a storage volume (p. 65) .
Throughput from your application to a storage volume has dropped to zero.	<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. 99). 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. 61). • 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 Managing Your Application Access to Storage Volumes (p. 76). <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. 120).</p>

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. 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 rule of thumb, 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. 122).
- **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 rule of thumb, 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. 120)).

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. 122)) 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. 119)
- [Measuring Performance Between Your Application and Gateway](#) (p. 120)
- [Measuring Performance Between Your Gateway and AWS](#) (p. 122)
- [Monitoring AWS Storage Gateway's Working Storage](#) (p. 124)
- [Understanding AWS Storage Gateway Metrics](#) (p. 128)

In this section, we discuss how to monitor your gateway, including its storage volumes and working storage. 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 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 specific 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	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> . Throughput and latency data of a gateway is based on all the volumes in the gateway. Data is available automatically in 5-minute periods at no charge.
	VolumeId	This dimension filters for metric data that is specific to a storage volume. Identify a storage volume to work with by <code>VolumeId</code> . Data is available automatically in 5-minute periods at no charge.

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.

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.

AWS Storage Gateway User Guide
Measuring Performance Between Your Application and Gateway

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.

AWS Storage Gateway User Guide
Measuring Performance Between Your Gateway and
AWS



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 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 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.

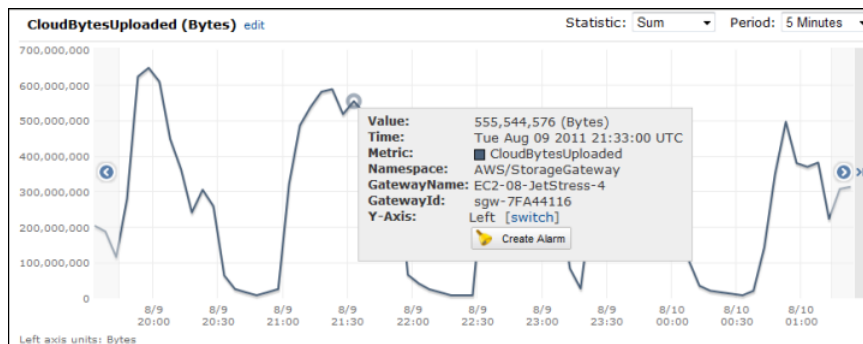
To measure the upload data throughput from a gateway to AWS

1	Select the StorageGateway: Gateway Metrics dimension and find the gateway that you want to work with.
---	--

AWS Storage Gateway User Guide
Measuring Performance Between Your Gateway and AWS

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, 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 AWS Storage Gateway's Working Storage

The following section discusses how to monitor working storage and how to create a working storage alarm so that you get a notification when working storage exceeds a specified threshold. This enables you to proactively add working storage to a gateway before the working storage fills completely and your storage application stops backing up to AWS.

Item of Interest	How to Measure
Working storage usage	Use the <code>WorkingStoragePercentUsed</code> , <code>WorkingStorageUsed</code> , <code>WorkingStorageFree</code> metrics with the Sum and Average statistics. For example, use the <code>WorkingStorageUsed</code> with the Average 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 working storage percent used

1	Select the StorageGateway: Gateway Metrics dimension and find the gateway that you want to work with.
2	Select the <code>WorkingStoragePercentUsed</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 working storage.

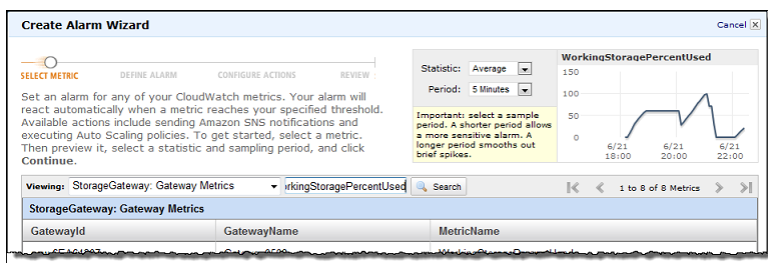
The working storage task below shows you how to create an alarm using the Amazon CloudWatch console and the **Create Alarm Wizard**. To learn more about alarms and thresholds, see [Creating CloudWatch Alarms](#).

To set an upper threshold alarm for a gateway's working storage

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 `WorkingStoragePercentUsed` 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 `WorkingStoragePercentUsed` greater than or equal to 50 percent for 5 minutes.

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:
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 `WorkingStoragePercentUsed` is \geq 50 for 5 minutes.

Metric: WorkingStoragePercentUsed
Period: 5 Minutes
Statistic: Average

WorkingStoragePercentUsed (Percent)

Time	WorkingStoragePercentUsed (Percent)
6/21 17:00	0
6/21 18:00	50
6/21 19:00	50
6/21 20:00	50
6/21 21:00	100
6/21 22:00	50

[< Back](#) [Continue >](#)

- c. Click **Continue**.

4. Configure an email action for the alarm.

- a. In the **CONFIGURE ACTIONS** page of the **Create Alarm Wizard**, select **ALARM** from the **Alarm State** drop-down list.
- b. 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](#).

- c. In the **Topic** field, enter a descriptive name for the topic.
- d. Click **ADD ACTION**.

Create Alarm Wizard Cancel X

SELECT METRIC DEFINE ALARM **CONFIGURE ACTIONS** REVIEW

Define what actions are taken when your 'Alarm Name' 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.

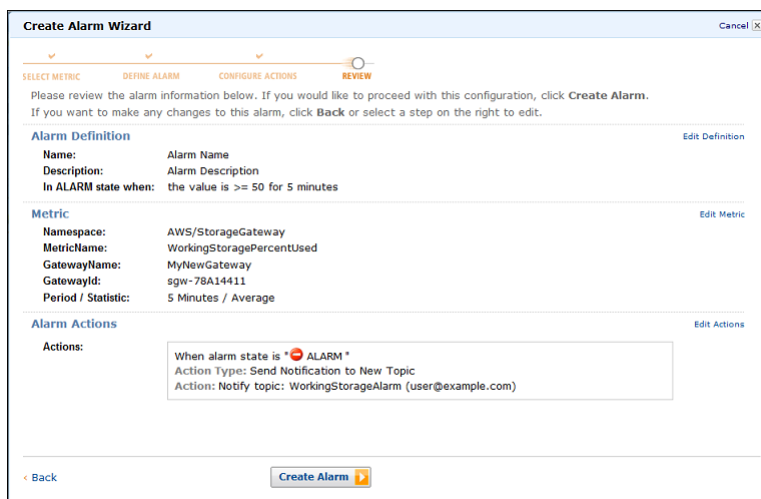
Define Your Actions
Actions define what steps you want to automate when the alarm state changes. For example, you can send a message using email via the Simple Notification Service (SNS). You can also execute an Auto Scaling Policy, if you have one configured ([learn about policies](#)).

When Alarm state is	Take action	Action details	
ALARM	Send Notification	Topic: Select or create email topic...	ADD ACTION

[< Back](#) [Continue >](#)

- e. Click **Continue**.

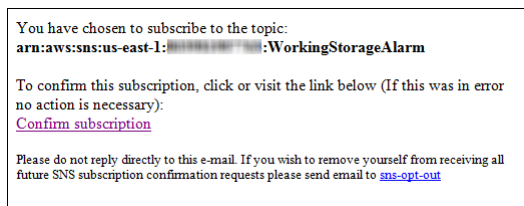
5. Review the alarm settings and create the alarm.
 - a. In the **REVIEW** page of the **Create Alarm Wizard**, review the alarm definition, metric, and associated actions from this step.



- b. After reviewing the alarm summary, click **Create Alarm**.

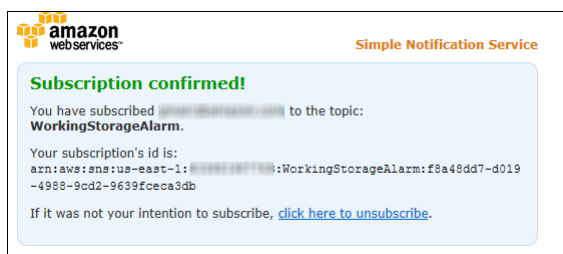
6. Confirm your subscription to the alarm topic.
 - a. Open the Simple Notification Service email that is sent to the email address that you specified when creating the topic.

The example below shows a notification.



- b. Confirm your subscription by clicking the link in the email.

A subscription confirmation displays.



Understanding AWS Storage Gateway Metrics

Topics

- [Gateway Metrics \(p. 128\)](#)
- [Storage Volume Metrics \(p. 129\)](#)

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 \(p. 119\)](#).

The following table describes the AWS Storage Gateway metrics that you can use to get information about your gateway.

Metric	Description
<code>ReadBytes</code>	<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 <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>
<code>WriteBytes</code>	<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>
<code>ReadTime</code>	<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 <code>Average</code> statistic to measure latency.</p> <p>Units: Milliseconds</p>
<code>WriteTime</code>	<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>

Metric	Description
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>
CloudBytesDownloaded	<p>The total number of pre-compressed bytes that the gateway downloaded from AWS during 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>
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 <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>
CloudDownloadLatency	<p>The total number of milliseconds spent reading data from AWS during the reporting period.</p> <p>Use this metric with the <code>Average</code> statistic to measure latency.</p> <p>Units: Milliseconds</p>
WorkingStoragePercentUsed	<p>Percent utilization of the gateway's working storage. The sample is taken at the end of the reporting period.</p> <p>Units: Percent</p>
WorkingStorageUsed	<p>The total number of bytes being used in the gateway's working storage. The sample is taken at the end of the reporting period.</p> <p>Units: Bytes</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>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. One thing to note is 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 \(p. 119\)](#).

The following table describes the AWS Storage Gateway metrics that you can use to get information about your storage volumes.

AWS Storage Gateway User Guide
Understanding AWS Storage Gateway Metrics

Metric	Description
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>
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>
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>
WriteTime	<p>The total number of milliseconds spent to do writes 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>
QueuedWrites	<p>The number of bytes waiting to be written to AWS, sampled at the end of the reporting period.</p> <p>Units: Bytes</p>

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 \(p. 137\)](#)) 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 \(p. 153\)](#).

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 \(p. 134\)](#).

API Reference for AWS Storage Gateway

Topics

- [AWS Storage Gateway Required Request Headers \(p. 137\)](#)
- [Signing Requests \(p. 139\)](#)
- [Error Responses \(p. 141\)](#)
- [Operations in AWS Storage Gateway \(p. 153\)](#)

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 \(p. 154\)](#) operation.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
```

AWS Storage Gateway User Guide

Required Request Headers

```
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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.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
<code>Authorization</code>	<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. 139).</p>
<code>Content-Type</code>	<p>Use application/x-amz-json-1.1 as the content type for all requests to AWS Storage Gateway.</p> <pre>Content-Type: application/x-amz-json-1.1</pre>
<code>Host</code>	<p>Use the host header to specify the AWS Storage Gateway endpoint where you send your request. For example, storagegateway.us-east-1.amazonaws.com 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>
<code>x-amz-date</code>	<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</i>'T'<i>HHMMSS</i>'Z'</pre>

Description
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. 137).

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 re-calculates 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. 194). 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, 07 May 2012 00:00:00" GMT.

- The endpoint is the US East Region.

The general request syntax (including the JSON body) is:

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
x-amz-Date: 20120507T000000Z
Authorization: SignatureToBeCalculated
Content-type: application/x-amz-json-1.1
x-amz-target: StorageGateway_20120430.ListGateways
{ }
```

The canonical form of the request calculated for [Task 1: Create a Canonical Request \(p. 139\)](#) is:

```
POST
/

content-type:application/x-amz-json-1.1
host:storagegateway.us-east-1.amazonaws.com
x-amz-date:20120507T000000Z
x-amz-target:StorageGateway_20120430.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 \(p. 139\)](#) is:

```
AWS4-HMAC-SHA256
20120507T000000Z
20120507/us-east-1/storagegateway/aws4_request
61f67415f7d34bca3d43b007db06e85e2cc835436df3bd64e232c3556f47b940
```

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 \(p. 139\)](#), the *derived key* can be represented as:

```
derived key = HMAC(HMAC(HMAC(HMAC("AWS4" + YourSecretAccessKey, "20120507"), "us-east-1"), "storagegateway"), "aws4_request")
```

If the secret access key, wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY, is used, then the calculated signature is:

```
a2a495a059d7a96bab837b2278c68b31221a1f31f4d03f0294296d5a26f095d0
```

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/20120507/us-east-1/storagegateway/aws4_request,
SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=a2a495a059d7a96bab837b2278c68b31221a1f31f4d03f0294296d5a26f095d0
```

Error Responses

Topics

- [Exceptions \(p. 141\)](#)
- [Operation Error Codes \(p. 142\)](#)
- [Error Responses \(p. 151\)](#)

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 \(p. 154\)](#) 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 \(p. 151\)](#).

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 \(p. 142\)](#) 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. 142) .	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. 142) .	400 Bad Request

AWS Storage Gateway User Guide
Operation Error Codes

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. 153) .	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 \(p. 141\)](#).

Operation Error Code	Message	Operations That Return this Error Code
ActivationKeyExpired	The specified activation key has expired.	ActivateGateway (p. 154)
ActivationKeyInvalid	The specified activation key is invalid.	ActivateGateway (p. 154)

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. 154)
BandwidthThrottleScheduleNotFound	The specified bandwidth throttle was not found.	DeleteBandwidthRateLimit (p. 166)
CannotExportSnapshot	The specified snapshot cannot be exported.	CreateStorediSCSIVolume (p. 162)
InitiatorNotFound	The specified initiator was not found.	DeleteChapCredentials (p. 168)
DiskAlreadyAllocated	The specified disk is already allocated.	AddWorkingStorage (p. 156) CreateStorediSCSIVolume (p. 162)
DiskDoesNotExist	The specified disk does not exist.	AddWorkingStorage (p. 156) CreateStorediSCSIVolume (p. 162)
DiskSizeNotGigAligned	The specified disk is not gigabyte-aligned.	CreateStorediSCSIVolume (p. 162)
DiskSizeGreaterThanVolumeMaxSize	The specified disk size is greater than the maximum volume size.	CreateStorediSCSIVolume (p. 162)
DiskSizeLessThanVolumeSize	The specified disk size is less than the volume size.	CreateStorediSCSIVolume (p. 162)
DuplicateCertificateInfo	The specified certificate information is a duplicate.	ActivateGateway (p. 154)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
GatewayInternalError	A gateway internal error occurred.	AddWorkingStorage (p. 156) CreateSnapshot (p. 159) CreateStorediSCSIVolume (p. 162) DeleteBandwidthRateLimit (p. 166) DeleteChapCredentials (p. 168) DeleteVolume (p. 172) DescribeBandwidthRateLimit (p. 175) DescribeChapCredentials (p. 177) DescribeGatewayInformation (p. 180) DescribeMaintenanceStartTime (p. 183) DescribeSnapshotSchedule (p. 185) DescribeStorediSCSIVolumes (p. 188) DescribeWorkingStorage (p. 192) ListLocalDisks (p. 197) ListVolumes (p. 200) ShutdownGateway (p. 203) StartGateway (p. 206) UpdateBandwidthRateLimit (p. 208) UpdateChapCredentials (p. 210) UpdateMaintenanceStartTime (p. 218) UpdateGatewaySoftwareNow (p. 216) UpdateSnapshotSchedule (p. 220)

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.	AddWorkingStorage (p. 156) CreateSnapshot (p. 159) CreateStorediSCSIVolume (p. 162) DeleteBandwidthRateLimit (p. 166) DeleteChapCredentials (p. 168) DeleteVolume (p. 172) DescribeBandwidthRateLimit (p. 175) DescribeChapCredentials (p. 177) DescribeGatewayInformation (p. 180) DescribeMaintenanceStartTime (p. 183) DescribeSnapshotSchedule (p. 185) DescribeStorediSCSIVolumes (p. 188) DescribeWorkingStorage (p. 192) ListLocalDisks (p. 197) ListVolumes (p. 200) ShutdownGateway (p. 203) StartGateway (p. 206) UpdateBandwidthRateLimit (p. 208) UpdateChapCredentials (p. 210) UpdateMaintenanceStartTime (p. 218) UpdateGatewaySoftwareNow (p. 216) UpdateSnapshotSchedule (p. 220)

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.	AddWorkingStorage (p. 156) CreateSnapshot (p. 159) CreateStorediSCSIVolume (p. 162) DeleteBandwidthRateLimit (p. 166) DeleteChapCredentials (p. 168) DeleteGateway (p. 170) DeleteVolume (p. 172) DescribeBandwidthRateLimit (p. 175) DescribeChapCredentials (p. 177) DescribeGatewayInformation (p. 180) DescribeMaintenanceStartTime (p. 183) DescribeSnapshotSchedule (p. 185) DescribeStorediSCSIVolumes (p. 188) DescribeWorkingStorage (p. 192) ListLocalDisks (p. 197) ListVolumes (p. 200) ShutdownGateway (p. 203) StartGateway (p. 206) UpdateBandwidthRateLimit (p. 208) UpdateChapCredentials (p. 210) UpdateMaintenanceStartTime (p. 218) UpdateGatewaySoftwareNow (p. 216) UpdateSnapshotSchedule (p. 220)

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.	AddWorkingStorage (p. 156) CreateSnapshot (p. 159) CreateStorediSCSIVolume (p. 162) DeleteBandwidthRateLimit (p. 166) DeleteChapCredentials (p. 168) DeleteVolume (p. 172) DescribeBandwidthRateLimit (p. 175) DescribeChapCredentials (p. 177) DescribeGatewayInformation (p. 180) DescribeMaintenanceStartTime (p. 183) DescribeSnapshotSchedule (p. 185) DescribeStorediSCSIVolumes (p. 188) DescribeWorkingStorage (p. 192) ListLocalDisks (p. 197) ListVolumes (p. 200) ShutdownGateway (p. 203) StartGateway (p. 206) UpdateBandwidthRateLimit (p. 208) UpdateChapCredentials (p. 210) UpdateMaintenanceStartTime (p. 218) UpdateGatewaySoftwareNow (p. 216) UpdateSnapshotSchedule (p. 220)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
InternalError	An internal error occurred.	ActivateGateway (p. 154) AddWorkingStorage (p. 156) CreateSnapshot (p. 159) CreateStorediSCSIVolume (p. 162) DeleteBandwidthRateLimit (p. 166) DeleteChapCredentials (p. 168) DeleteGateway (p. 170) DeleteVolume (p. 172) DescribeBandwidthRateLimit (p. 175) DescribeChapCredentials (p. 177) DescribeGatewayInformation (p. 180) DescribeMaintenanceStartTime (p. 183) DescribeSnapshotSchedule (p. 185) DescribeStorediSCSIVolumes (p. 188) DescribeWorkingStorage (p. 192) ListLocalDisks (p. 197) ListGateways (p. 194) ListVolumes (p. 200) ShutdownGateway (p. 203) StartGateway (p. 206) UpdateBandwidthRateLimit (p. 208) UpdateChapCredentials (p. 210) UpdateMaintenanceStartTime (p. 218) UpdateGatewayInformation (p. 213) UpdateGatewaySoftwareNow (p. 216) UpdateSnapshotSchedule (p. 220)

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.	ActivateGateway (p. 154) AddWorkingStorage (p. 156) CreateSnapshot (p. 159) CreateStorediSCSIVolume (p. 162) DeleteBandwidthRateLimit (p. 166) DeleteChapCredentials (p. 168) DeleteGateway (p. 170) DeleteVolume (p. 172) DescribeBandwidthRateLimit (p. 175) DescribeChapCredentials (p. 177) DescribeGatewayInformation (p. 180) DescribeMaintenanceStartTime (p. 183) DescribeSnapshotSchedule (p. 185) DescribeStorediSCSIVolumes (p. 188) DescribeWorkingStorage (p. 192) ListLocalDisks (p. 197) ListGateways (p. 194) ListVolumes (p. 200) ShutdownGateway (p. 203) StartGateway (p. 206) UpdateBandwidthRateLimit (p. 208) UpdateChapCredentials (p. 210) UpdateMaintenanceStartTime (p. 218) UpdateGatewayInformation (p. 213) UpdateGatewaySoftwareNow (p. 216) UpdateSnapshotSchedule (p. 220)
LocalStorageLimitExceeded	The local storage limit was exceeded.	AddWorkingStorage (p. 156)
LunInvalid	The specified LUN is invalid.	CreateStorediSCSIVolume (p. 162)

AWS Storage Gateway User Guide
Operation Error Codes

Operation Error Code	Message	Operations That Return this Error Code
MaximumVolumeCountExceeded	The maximum volume count was exceeded.	CreateStorediSCSIVolume (p. 162) DescribeStorediSCSIVolumes (p. 188)
NetworkConfigurationChanged	The gateway network configuration has changed.	CreateStorediSCSIVolume (p. 162)
NotSupported	The specified operation is not supported.	ActivateGateway (p. 154) AddWorkingStorage (p. 156) CreateSnapshot (p. 159) CreateStorediSCSIVolume (p. 162) DeleteBandwidthRateLimit (p. 166) DeleteChapCredentials (p. 168) DeleteGateway (p. 170) DeleteVolume (p. 172) DescribeBandwidthRateLimit (p. 175) DescribeChapCredentials (p. 177) DescribeGatewayInformation (p. 180) DescribeMaintenanceStartTime (p. 183) DescribeSnapshotSchedule (p. 185) DescribeStorediSCSIVolumes (p. 188) DescribeWorkingStorage (p. 192) ListLocalDisks (p. 197) ListGateways (p. 194) ListVolumes (p. 200) ShutdownGateway (p. 203) StartGateway (p. 206) UpdateBandwidthRateLimit (p. 208) UpdateChapCredentials (p. 210) UpdateMaintenanceStartTime (p. 218) UpdateGatewayInformation (p. 213) UpdateGatewaySoftwareNow (p. 216) UpdateSnapshotSchedule (p. 220)

Operation Error Code	Message	Operations That Return this Error Code
OutdatedGateway	The specified gateway is out of date.	ActivateGateway (p. 154)
SnapshotInProgressException	The specified snapshot is in progress.	DeleteVolume (p. 172)
SnapshotIdInvalid	The specified snapshot is invalid.	CreateStorediSCSIVolume (p. 162)
StagingAreaFull	The staging area is full.	CreateStorediSCSIVolume (p. 162)
TargetAlreadyExists	The specified target already exists.	CreateStorediSCSIVolume (p. 162)
TargetInvalid	The specified target is invalid.	CreateStorediSCSIVolume (p. 162) DeleteChapCredentials (p. 168) DescribeChapCredentials (p. 177) UpdateChapCredentials (p. 210)
TargetNotFound	The specified target was not found.	CreateStorediSCSIVolume (p. 162) DeleteChapCredentials (p. 168) DescribeChapCredentials (p. 177) DeleteVolume (p. 172) UpdateChapCredentials (p. 210)
VolumeAlreadyExists	The specified volume already exists.	CreateStorediSCSIVolume (p. 162)
VolumeIdInvalid	The specified volume is invalid.	DeleteVolume (p. 172)
VolumeInUse	The specified volume is already in use.	DeleteVolume (p. 172)
VolumeNotFound	The specified volume was not found.	CreateSnapshot (p. 159) DeleteVolume (p. 172) DescribeSnapshotSchedule (p. 185) DescribeStorediSCSIVolumes (p. 188) UpdateSnapshotSchedule (p. 220)
VolumeNotReady	The specified volume is not ready.	CreateSnapshot (p. 159)

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 \(p. 141\)](#).

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 \(p. 142\)](#).

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 \(p. 142\)](#).

Type: String

Error Response Examples

The following JSON body is returned if you use the [DescribeStorediSCSIVolumes \(p. 188\)](#) 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. 154)
- [AddWorkingStorage](#) (p. 156)
- [CreateSnapshot](#) (p. 159)
- [CreateStorediSCSIVolume](#) (p. 162)
- [DeleteBandwidthRateLimit](#) (p. 166)
- [DeleteChapCredentials](#) (p. 168)
- [DeleteGateway](#) (p. 170)
- [DeleteVolume](#) (p. 172)
- [DescribeBandwidthRateLimit](#) (p. 175)
- [DescribeChapCredentials](#) (p. 177)
- [DescribeGatewayInformation](#) (p. 180)
- [DescribeMaintenanceStartTime](#) (p. 183)
- [DescribeSnapshotSchedule](#) (p. 185)
- [DescribeStorediSCSIVolumes](#) (p. 188)
- [DescribeWorkingStorage](#) (p. 192)
- [ListGateways](#) (p. 194)
- [ListLocalDisks](#) (p. 197)
- [ListVolumes](#) (p. 200)
- [ShutdownGateway](#) (p. 203)
- [StartGateway](#) (p. 206)
- [UpdateBandwidthRateLimit](#) (p. 208)
- [UpdateChapCredentials](#) (p. 210)
- [UpdateGatewayInformation](#) (p. 213)
- [UpdateGatewaySoftwareNow](#) (p. 216)
- [UpdateMaintenanceStartTime](#) (p. 218)
- [UpdateSnapshotSchedule](#) (p. 220)
- [Data Types](#) (p. 223)
- [Enumeration Types](#) (p. 228)

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.

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 \(p. 46\)](#). In the activation process you specify information such as 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 timezone after activation (see [UpdateGatewayInformation \(p. 213\)](#)).

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_20120430.ActivateGateway

{
  "ActivationKey": "String",
  "GatewayName": "String",
  "GatewayTimezone": "String",
  "GatewayRegion": "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`.

Required: Yes

Type: String

GatewayName

A unique identifier for your gateway. This name becomes part of the gateway Amazon Resources Name (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 \(p. 229\)](#) 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 \(p. 228\)](#) 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

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 Amazon Resource Name (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 \(p. 141\)](#).

- ActivationKeyExpired
- ActivationKeyInvalid
- ActivationKeyNotFound
- DuplicateCertificateInfo
- InternalError
- InvalidParameters
- NotSupported
- OutdatedGateway

Examples

Example Request

The following example shows a request that activates 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.ActivateGateway

{
  "ActivationKey": "29AV1-30FV9-VVIUB-NKT0I-LR06V",
  "GatewayName": "mygateway",
  "GatewayTimezone": "GMT-12:00",
  "GatewayRegion": "us-east-1"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 206)
- [ListGateways](#) (p. 194)
- [ShutdownGateway](#) (p. 203)
- [DescribeGatewayInformation](#) (p. 180)
- [DeleteGateway](#) (p. 170)

AddWorkingStorage

Description

This operation configures one or more gateway local disks as working storage.

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_20120430.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 \(p. 197\)](#) API.

Required: Yes

Type: Array

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 194\)](#) 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 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 \(p. 141\)](#).

- 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 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.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, 25 Apr 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. 192)
- [ListLocalDisks](#) (p. 197)

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 in the AWS Storage Gateway Console](#) (p. 88).

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_20120430.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 \(p. 200\)](#) 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 \(p. 162\)](#)).

Type: String

VolumeARN

The Amazon Resource Name (ARN) of the volume of which the snapshot was taken. Obtain volume ARNs from the

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 \(p. 141\)](#).

- 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=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120425T120000Z
x-amz-target: StorageGateway_20120430.CreateSnapshot

{
  "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, 25 Apr 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",
  "SnapshotId": "snap-78e22663"
}
```

Related Actions

- [UpdateSnapshotSchedule](#) (p. 220)
- [DescribeSnapshotSchedule](#) (p. 185)

CreateStorediSCSIVolume

Description

This operation creates a volume on a specified gateway. The size of the volume 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 ARN, its size, and the iSCSI target Amazon Resource Name (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_20120430.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. 197) to list disk IDs for a gateway.

Required: Yes

Type: String

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 194) 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 \(p. 180\)](#) 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 Amazon Resource Name (ARN) of the volume target that includes the iSCSI name that initiators can use to connect to the target.

Type: String

VolumeARN

The Amazon Resource Name (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 \(p. 141\)](#).

- CannotExportSnapshot
- DiskAlreadyAllocated
- DiskDoesNotExist
- DiskSizeNotGigAligned
- DiskSizeGreaterThanVolumeMaxSize
- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- LunInvalid
- MaximumVolumeCountExceeded
- NetworkConfigurationChanged
- NotSupported
- SnapshotIdInvalid
- StagingAreaFull
- TargetAlreadyExists
- TargetInvalid
- TargetNotFound
- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.CreateStorediSCSIVolume

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
  "DiskId": "pci-0000:03:00.0-scsi-0:0:0:0",
  "TargetName": "myvolume",
  "NetworkInterfaceId": "10.1.1.1"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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/mygateway/volume/vol-1122AABB",
  "VolumeSizeInBytes": 1099511627776,
  "TargetARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/target/iqn.1997-05.com.amazon:myvolume"
}
```

Related Actions

- [ListVolumes](#) (p. 200)
- [ListLocalDisks](#) (p. 197)
- [DeleteVolume](#) (p. 172)
- [DescribeStorediSCSIVolumes](#) (p. 188)

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_20120430.DeleteBandwidthRateLimit

{
  "GatewayARN": "String",
  "BandwidthType": "String"
}
```

JSON Fields

BandwidthType

One of the [BandwidthType](#) (p. 228) 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. 194) 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 \(p. 141\)](#).

- 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: 20120425T120000Z  
x-amz-target: StorageGateway_20120430.DeleteBandwidthRateLimit  
  
{  
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",  
  "BandwidthType": "ALL"  
}
```

Example Response

```
HTTP/1.1 200 OK  
x-amzn-RequestId: gur28r2rq1gb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
```

```
Date: Wed, 25 Apr 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. 175)
- [UpdateBandwidthRateLimit](#) (p. 208)

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_20120430.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. 188) 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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.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: gur28r2rqlgb8vvs0mq17hlgijlq8glllelqeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 210)
- [DescribeChapCredentials](#) (p. 177)

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 from your environment.

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_20120430.DeleteGateway

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 194\)](#) 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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.DeleteGateway

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 194)
- [ShutdownGateway](#) (p. 203)

DeleteVolume

Description

This operation deletes the specified gateway volume that you created previously using the [CreateStorediSCSIVolume](#) (p. 162) API. The gateway 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_20120430.DeleteVolume

{
  "VolumeARN": "String"
}
```

JSON Fields

VolumeARN

The Amazon Resource Name (ARN) of the volume. Use the [ListVolumes \(p. 200\)](#) 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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.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, 25 Apr 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. 162)
- [ListLocalDisks](#) (p. 197)

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_20120430.DescribeBandwidthRateLimit

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 194) 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 \(p. 141\)](#).

- 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=9cd5a3584d1d67d57e61f120f35102d6b3649066ab  
dd4bf4bbcf05bd9f2f8fe2  
x-amz-date: 20120425T120000Z  
x-amz-target: StorageGateway_20120430.DescribeBandwidthRateLimit  
  
{
```

```
"GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygate  
way"  
}
```

Example Response

```
HTTP/1.1 200 OK  
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0  
Date: Wed, 25 Apr 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. 208)
- [DeleteBandwidthRateLimit](#) (p. 166)

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_20120430.DescribeChapCredentials  
  
{  
  "TargetARN": "String"  
}
```

JSON Fields

TargetARN

The Amazon Resource Name (ARN) of the iSCSI volume target. Use the [DescribeStorediSCSIVolumes](#) (p. 188) 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. 223) 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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.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: gur28r2rqlgb8vvs0mq17hlgijlq8gllelqeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 168)
- [UpdateChapCredentials](#) (p. 210)

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_20120430.DescribeGatewayInformation  
  
{  
  "GatewayARN": "String"  
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 194) 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",
  "NextUpdateAvailabilityDate": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (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. 225) array that contains descriptions of the gateway network interfaces.

Type: Array

GatewayState

One of the [GatewayState](#) (p. 228) values that indicates the operating state of the gateway.

Type: String

GatewayTimezone

One of the [GatewayTimezone](#) (p. 228) values that indicates the time zone configured for the gateway.

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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.DescribeGatewayInformation

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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"
}
```

Related Actions

- [ListGateways](#) (p. 194)

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_20120430.DescribeMaintenanceStartTime

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 194) 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 \(p. 228\)](#) 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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.DescribeMaintenanceStartTime

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllelqeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 218)

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_20120430.DescribeSnapshotSchedule

{
```

```
"VolumeARN": "String"  
}
```

JSON Fields

VolumeARN

The Amazon Resource Name (ARN) of the volume. Use the [ListVolumes \(p. 200\)](#) 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 \(p. 228\)](#) 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. 141).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- NotSupported
- VolumeNotFound

Examples

The following example shows a request that deletes the volume *myvolume*.

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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.DescribeSnapshotSchedule

{
  "VolumeARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway/volume/vol-1122AABB"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllelqeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 220)

DescribeStorediSCSIVolumes

Description

This operation returns description of the gateway volumes specified in the request. 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_20120430.DescribeStorediSCSIVolumes

{
  "VolumeARNs": [ "String", ... ]
}
```

JSON Fields

VolumeARNs

An array of strings, where each string represents the Amazon Resource Name (ARN) of a stored volume. All of the specified stored volumes must be from the same gateway. Use [ListVolumes](#) (p. 200) 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, e.g. snap-1122aabb. Otherwise, this field is not included.

Type: String

StorediSCSIVolumes

An array of [StorediSCSIVolume](#) (p. 226) objects where each object contains metadata about one stored volume.

Type: Array

TargetARN

The Amazon Resource Name (ARN) of the volume target.

Type: String

VolumeARN

The Amazon Resource Name (ARN) of the stored volume.

Type: String

VolumeDiskId

The disk ID of the local disk that was specified in the [CreateStorediSCSIVolume \(p. 162\)](#) operation.

Type: String

VolumeId

The unique identifier of the storage volume, e.g. vol-1122AABB.

Type: String

VolumeiSCSIAttributes

An [VolumeiSCSIAttributes \(p. 227\)](#) 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 \(p. 229\)](#) 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 \(p. 229\)](#)

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 \(p. 141\)](#).

- GatewayInternalError
- GatewayNotConnected
- GatewayNotFound
- GatewayProxyNetworkConnectionBusy
- InternalError
- InvalidParameters
- MaximumVolumeCountExceeded
- NotSupported

- VolumeNotFound

Examples

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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.DescribeStorediSCSIVolumes

{
  "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, 25 Apr 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. 162)
- [DescribeWorkingStorage](#) (p. 192)
- [ListLocalDisks](#) (p. 197)

DescribeWorkingStorage

Description

This operation returns information about the working storage of a 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_20120430.DescribeWorkingStorage

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 194) 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 Amazon Resource Name (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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.DescribeWorkingStorage

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 156)
- [ListLocalDisks](#) (p. 197)

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_20120430.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. 225) 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. 141).

- `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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.ListGateways
{}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 180)

ListLocalDisks

Description

This operation returns a list of the local disks of a gateway. To specify which gateway to describe 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_20120430.ListLocalDisks
```

```
{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 194) 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. 228) 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 VMSphere client 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. 224) 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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.ListLocalDisks

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 200)
- [DescribeGatewayInformation](#) (p. 180)

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. 188) 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_20120430.ListVolumes

{
  "GatewayARN": "String",
  "Marker": "String",
  "Limit": Number
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 194\)](#) 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 \(p. 227\)](#) 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 \(p. 229\)](#) 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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.ListVolumes

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 197)
- [DescribeStorediSCSIVolumes](#) (p. 188)
- [CreateStorediSCSIVolume](#) (p. 162)

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.

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 shutdown, you cannot call any other API except [StartGateway](#) (p. 206), [DescribeGatewayInformation](#) (p. 180), and [ListGateways](#) (p. 194). For more information, see [ActivateGateway](#) (p. 154). 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 \(p. 154\)](#).

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_20120430.ShutdownGateway

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 194\)](#) 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. 141).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.ShutdownGateway

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 206)
- [DeleteGateway](#) (p. 170)
- [ActivateGateway](#) (p. 154)

StartGateway

Description

This operation starts a gateway that you previously shut down (see [ShutdownGateway \(p. 203\)](#)). 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 \(p. 154\)](#).

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_20120430.StartGateway

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 194\)](#) 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 \(p. 141\)](#).

- 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=9cd5a3584d1d67d57e61f120f35102d6b3649066ab  
dd4bf4bbcf05bd9f2f8fe2  
x-amz-date: 20120425T120000Z  
x-amz-target: StorageGateway_20120430.StartGateway  
  
{  
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygate  
way"  
}
```

Example Response

```
HTTP/1.1 200 OK  
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glle1qeu3kpgg6f0kstauu0  
Date: Wed, 25 Apr 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. 203)
- [DeleteGateway](#) (p. 170)

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_20120430.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 \(p. 194\)](#) 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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.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: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 175)
- [DeleteBandwidthRateLimit](#) (p. 166)

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_20120430.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 \(p. 188\)](#) 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 \(p. 141\)](#).

- 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=9cd5a3584d1d67d57e61f120f35102d6b3649066ab
dd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120425T120000Z
x-amz-target: StorageGateway_20120430.UpdateChapCredentials

{
  "TargetARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygate
way/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: gur28r2rqlgb8vvs0mq17hlgijlq8gllle1qeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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/mygate
way/target/iqn.1997-05.com.amazon:myvolume",
  "InitiatorName": "iqn.1991-05.com.microsoft:computername.domain.example.com"
}
```

Related Actions

- [DeleteChapCredentials](#) (p. 168)
- [DescribeChapCredentials](#) (p. 177)
- [Configuring CHAP Authentication for Your Storage Volume](#) (p. 80)

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_20120430.UpdateGatewayInformation
```

```
{  
  "GatewayARN": "String",  
  "GatewayName": "String",  
  "GatewayTimezone": "String"  
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 194\)](#) 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 \(p. 228\)](#) 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. 141).

- `InternalServerError`
- `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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.UpdateGatewayInformation

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway",
  "GatewayName": "mygateway2"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8glllelqeu3kpgg6f0kstauu0
Date: Wed, 25 Apr 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. 180)
- [ListGateways](#) (p. 194)

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 \(p. 180\)](#) to verify the gateway is in the STATE_RUNNING state. For more information, see [DescribeGatewayInformation \(p. 180\)](#).

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 \(p. 78\)](#) and [Customizing Your Linux iSCSI Settings \(p. 79\)](#), 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_20120430.UpdateGatewaySoftwareNow

{
  "GatewayARN": "String"
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways \(p. 194\)](#) 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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.UpdateGatewaySoftwareNow

{
  "GatewayARN": "arn:aws:storagegateway:us-east-1:111122223333:gateway/mygateway"
}
```

Example Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: gur28r2rqlgb8vvs0mq17hlgijlq8gll1qeu3kpgg6f0kstauu0
```

```
Date: Wed, 25 Apr 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. 183)

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_20120430.UpdateMaintenanceStartTime

{
  "GatewayARN": "String",
  "HourOfDay": "Number",
  "MinuteOfHour": "Number",
  "DayOfWeek": Number
}
```

JSON Fields

GatewayARN

The Amazon Resource Name (ARN) of the gateway. Use the [ListGateways](#) (p. 194) 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 (0 to 23). The hour of the day is in the time zone of the gateway.

Length: 2

Valid Values: *hh*, where *hh* is the hours (00 to 23).

Required: Yes

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.

Length: 2

Valid Values: *mm*, where *mm* the minutes (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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.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, 25 Apr 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. 183)

UpdateSnapshotSchedule

Description

This operation updates a snapshot schedule configured for a gateway volume.

The default snapshot schedule for volume 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_20120430.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 \(p. 200\)](#) 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 \(p. 141\)](#).

- 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: 20120425T120000Z
x-amz-target: StorageGateway_20120430.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, 25 Apr 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. 185)

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:

- [ChapInfo](#) (p. 223)
- [Disk](#) (p. 224)
- [GatewayInfo](#) (p. 225)
- [NetworkInterface](#) (p. 225)
- [StoriedSCSIVolume](#) (p. 226)
- [VolumeInfo](#) (p. 227)
- [VolumeiSCSIAttributes](#) (p. 227)

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 Amazon Resource Name (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. 228) 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 Amazon Resource Name (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. 227) 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, e.g. snap-78e22663. Otherwise, this field is not included.

Type: String

VolumeARN

The Amazon Resource Name (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. 162) operation.

Valid Values: TBD

Type: String

VolumeId

The unique identifier of the volume, e.g. 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. 229) 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. 229) enumeration values describing the type of the volume.

Type: String

VolumeInfo

Describes a storage volume.

```
{ "VolumeARN": "String",  
  "VolumeType": "String"  
}
```

VolumeARN

The Amazon Resource Name (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. 229) values that indicates the configuration of the storage volume, for example as a storage volume.

Type: String

VolumeiSCSIAttributes

Lists iSCSI information about a volume.

```
{ "TargetARN": "String",  
  "NetworkInterfaceId": "String",  
  "NetworkInterfacePort": Number,  
  "LunNumber": Number,  
  "ChapEnabled": Boolean  
}
```

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 Amazon Resource Name (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

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. 228)
- [DiskAllocationType](#) (p. 228)
- [GatewayState](#) (p. 228)
- [GatewayTimezone](#) (p. 228)
- [Regions](#) (p. 229)
- [VolumeStatus](#) (p. 229)
- [VolumeType](#) (p. 229)

BandwidthType

The bandwidth rate limit type.

Valid Values: UPLOAD | DOWNLOAD | ALL

DiskAllocationType

The configuration of a gateway local disk.

Valid Values: AVAILABLE | STORED iSCSI VOLUME | WORKING STORAGE

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

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 | WORKING STORAGE NOT CONFIGURED

VolumeType

The type of storage volume. Currently only STORED iSCSI is supported.

Valid Values: 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**—2012-04-30
- **Latest product release**—January 2012
- **Last document update**—25 Jun 2012

Change	Description	Release Date
API and IAM Support	<p>In this release of 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. 137) 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. 131). For more information about IAM, go to AWS Identity and Access Management (IAM) detail page. 	In this release.
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 Static IP Addresses (p. 105).</p>	05 Mar 2012

Change	Description	Release Date
New Guide	This is the first release of <i>AWS Storage Gateway User Guide</i> .	24 Jan 2012

Appendices for AWS Storage Gateway

This AWS Storage Gateway guide appendix includes the following sections.

Topics

- [Appendix A: The Components in Your vSphere Environment for AWS Storage Gateway \(p. 232\)](#)
- [Appendix B: Configuring a VMware ESXi Host for AWS Storage Gateway \(p. 234\)](#)
- [Appendix C: About AWS Storage Gateway \(p. 238\)](#)

Appendix A: The Components in Your vSphere Environment for AWS Storage Gateway

The AWS Storage Gateway uses VMware to create virtual machines which host the gateway and storage volumes. 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. 46).
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](#).

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 step-by-step instructions, see [Synchronize VM Time with the Host Time \(p. 14\)](#).


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**.

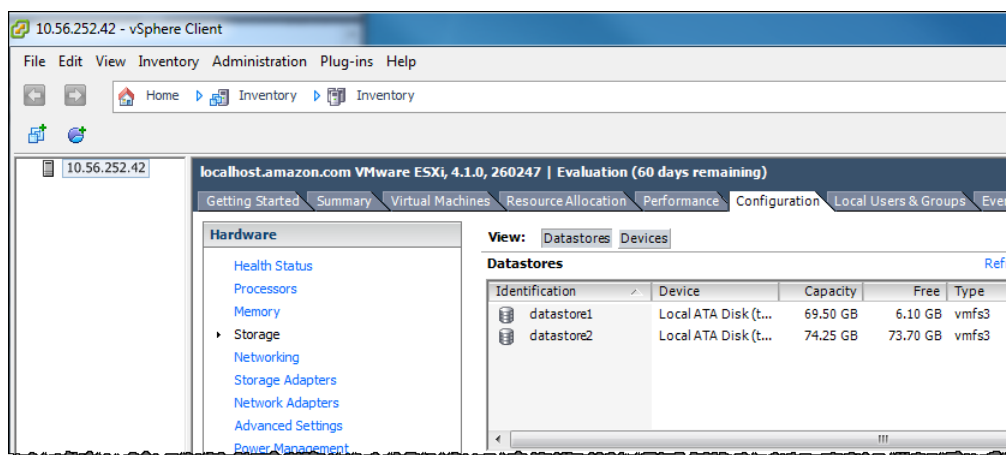
This connects your client to the host. You are now ready to configure the host.

To add a new iSCSI target

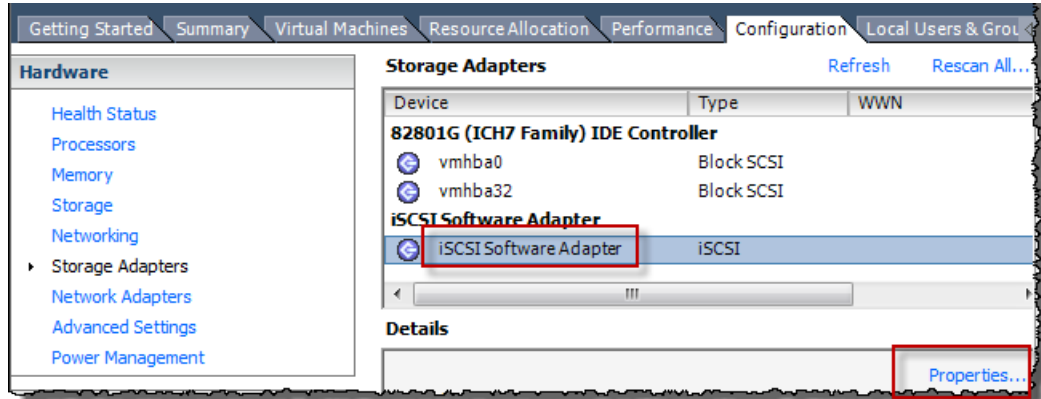
1. After you have connected to your remote device through the hypervisor, go to the **Configuration** tab of the host and click **Storage** in the **Hardware** list.

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.



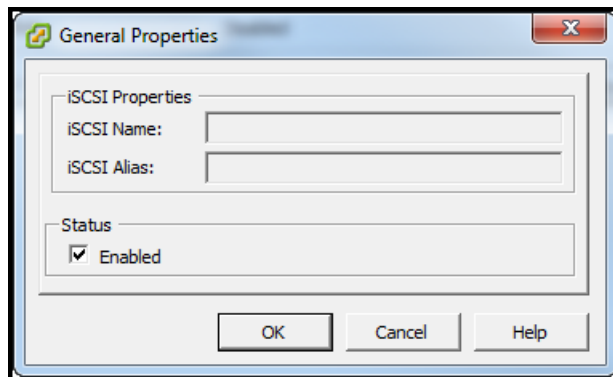
2. In the **Hardware** list, click **Storage Adapters**.
3. In the **Storage Adapters** pane, select **iSCSI Software Adapter**, and then click the **Properties** link in the **Details** pane.



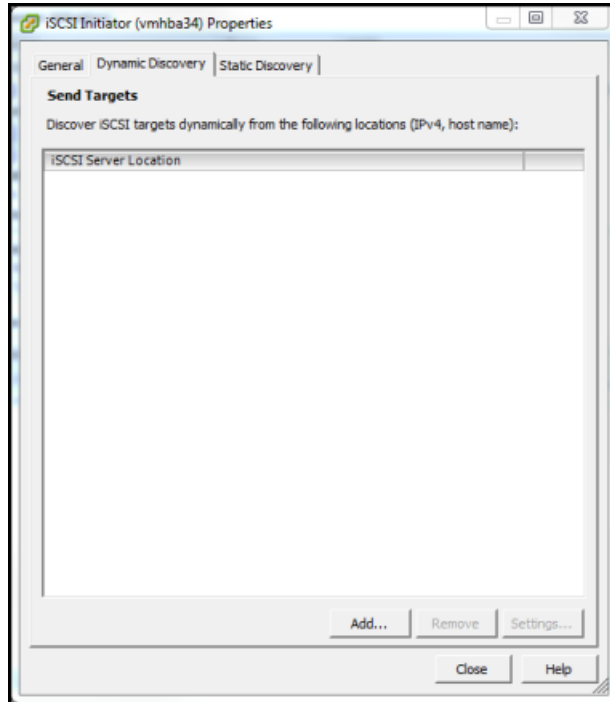
4. In the **iSCSI Initiator (iSCSI Software Adapter) Properties** dialog box, click **Configure**.




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**.



The name that you enter should be globally unique and properly formatted or some storage devices might not recognize the hardware iSCSI initiator.

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: About AWS Storage Gateway

The source code for certain open source software components that are included with the Software is available for download at <https://s3.amazonaws.com/aws-storage-gateway-terms/sources.tar>.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).

The AWS Storage Gateway uses CentOS 5 as the basis for gateway virtual machines (VMs). The packages comprising the 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. Similar to other Linux operating systems, these package updates do not cause the Major Version number of CentOS to increment. For more information about managing updates, see [Managing Gateway Updates Using the AWS Storage Gateway Console](#) (p. 100).