



Volume Gateway User Guide

AWS Storage Gateway



API Version 2013-06-30

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

AWS Storage Gateway: Volume Gateway User Guide

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

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

Table of Contents

What is Volume Gateway?	1
How Volume Gateway works	2
Volume Gateways	2
Getting started with AWS Storage Gateway	7
Sign Up for AWS Storage Gateway	7
Create an IAM user with administrator privileges	8
Accessing AWS Storage Gateway	9
AWS Regions that support Storage Gateway	10
Volume Gateway setup requirements	11
Hardware and storage requirements	11
Hardware requirements for VMs	11
Requirements for Amazon EC2 instance types	11
.....	12
Storage requirements	12
Network and firewall requirements	13
Port requirements	14
Networking and firewall requirements for the hardware appliance	19
Allowing gateway access through firewall and routers	22
Configuring security group	24
Supported hypervisors and host requirements	25
Supported iSCSI initiators	26
Using the hardware appliance	28
Setting up your hardware appliance	29
Physically installing your hardware appliance	30
Accessing the hardware appliance console	33
Configuring hardware appliance network parameters	34
Activating your hardware appliance	35
Creating a gateway on your hardware appliance	37
Configuring a gateway IP address on the hardware appliance	37
Removing gateway software from your hardware appliance	40
Deleting your hardware appliance	40
Creating your gateway	42
Overview - Gateway Activation	42
Set up gateway	42

Connect to AWS	43
Review and activate	43
Overview - Gateway Configuration	43
Overview - Storage Resources	43
Creating a Volume Gateway	43
Set up a Volume Gateway	44
Connect your Volume Gateway to AWS	45
Review settings and activate your Volume Gateway	46
Configure your Volume Gateway	47
Creating a volume	49
Configure CHAP authentication for your volumes	51
Connecting your volumes to your client	52
Connecting to a Microsoft Windows client	53
Connecting to a Red Hat Enterprise Linux client	53
Initializing and formatting your volume	54
Initializing and formatting on Windows	55
Initializing and formatting on RHEL	56
Testing your gateway	57
Backing up your volumes	59
Using Storage Gateway to back up your volumes	59
Using AWS Backup to back up your volumes	59
Where do I go from here?	62
Sizing Your Volume Gateway's Storage for Real-World Workloads	63
Activating your gateway in a virtual private cloud	64
Creating a VPC endpoint for Storage Gateway	65
Managing Your Volume Gateway	67
Editing Gateway Information	68
Adding and expanding volumes	69
Cloning a volume	70
Viewing volume usage	71
Deleting storage volumes	72
Moving Your Volumes to a Different Gateway	73
Creating a recovery snapshot	75
Editing a snapshot schedule	76
Deleting Snapshots	77
Using the AWS SDK for Java	77

Using the AWS SDK for .NET	81
Using the AWS Tools for Windows PowerShell	87
Understanding Volume Status and Transitions	90
Understanding Volume Status	90
Understanding Volume Status	94
Understanding Cached Volume Status Transitions	95
Understanding Stored Volume Status Transitions	98
Moving your data to a new gateway	100
Moving stored volumes to a new stored Volume Gateway	101
Moving cached volumes to a new gateway virtual machine	103
Monitoring Storage Gateway	107
Understanding gateway metrics	107
Dimensions for Storage Gateway metrics	113
Monitoring the upload buffer	114
Monitoring cache storage	116
Understanding CloudWatch alarms	118
Creating recommended CloudWatch alarms	119
Creating a custom CloudWatch alarm	120
Monitoring your Volume Gateway	122
Getting Volume Gateway health logs	123
Using Amazon CloudWatch Metrics	124
Measuring Performance Between Your Application and Gateway	126
Measuring Performance Between Your Gateway and AWS	127
Understanding volume metrics	131
Maintaining Your Gateway	138
Managing local disks	138
Deciding the amount of local disk storage	139
Add upload buffer or cache storage	142
Managing Bandwidth	143
Changing Bandwidth Throttling Using the Storage Gateway Console	144
Scheduling Bandwidth Throttling	144
Using the AWS SDK for Java	146
Using the AWS SDK for .NET	148
Using the AWS Tools for Windows PowerShell	150
Managing gateway updates	151
Update frequency and expected behavior	152

Turn maintenance updates on or off	152
Modify the gateway maintenance window schedule	153
Apply an update manually	154
Shutting Down Your Gateway VM	155
Starting and Stopping a Volume Gateway	156
Deleting your gateway and removing resources	157
Deleting Your Gateway by Using the Storage Gateway Console	157
Removing Resources from a Gateway Deployed On-Premises	158
Removing Resources from a Gateway Deployed on an Amazon EC2 Instance	159
Performing maintenance tasks using the local console	160
Accessing the Gateway Local Console	160
Accessing the Gateway Local Console with Linux KVM	161
Accessing the Gateway Local Console with VMware ESXi	161
Access the Gateway Local Console with Microsoft Hyper-V	162
Performing Tasks on the VM Local Console	163
Logging in to the Volume Gateway local console	164
Configuring a SOCKS5 proxy for your on-premises gateway	166
Configuring Your Gateway Network	167
Testing your gateway connectivity to the internet	172
Running storage gateway commands in the local console for an on-premises gateway	173
Viewing your gateway system resource status	175
Performing Tasks on the EC2 Local Console	176
Logging In to Your EC2 Gateway Local Console	177
Configuring an HTTP proxy	178
Testing gateway network connectivity	179
Viewing your gateway system resource status	180
Running Storage Gateway commands on the local console	181
Performance and optimization for Volume Gateway	183
Optimizing gateway performance	183
Recommended Configuration	183
Add Resources to Your Gateway	184
Optimize iSCSI Settings	187
Add Resources to Your Application Environment	187
Security	189
Data protection	190
Data encryption	191

Configuring CHAP authentication	192
Identity and Access Management	194
Audience	194
Authenticating with identities	195
Managing access using policies	198
How AWS Storage Gateway works with IAM	200
Identity-based policy examples	207
Troubleshooting	210
Compliance validation	211
Resilience	212
Infrastructure Security	213
AWS Security Best Practices	214
Logging and Monitoring	214
Storage Gateway Information in CloudTrail	214
Understanding Storage Gateway Log File Entries	215
Troubleshooting gateway issues	218
Troubleshooting: gateway offline issues	218
Check the associated firewall or proxy	219
Check for an ongoing SSL or deep-packet inspection of your gateway's traffic	219
Check for a power outage or hardware failure on the hypervisor host	219
Check for issues with an associated cache disk	219
Troubleshooting: gateway activation issues	220
Resolve errors when activating your gateway using a public endpoint	221
Resolve errors when activating your gateway using an Amazon VPC endpoint	224
Resolve errors when activating your gateway using a public endpoint and there is a Storage Gateway VPC endpoint in the same VPC	228
Troubleshooting on-premises gateway issues	228
Activating AWS Support to help troubleshoot your gateway	233
Troubleshooting Microsoft Hyper-V setup issues	234
Troubleshooting Amazon EC2 gateway issues	238
Gateway activation hasn't occurred after a few moments	238
Can't find the EC2 gateway instance in the instance list	239
Can't attach a an Amazon EBS volume to the EC2 gateway instance	239
Can't attach an initiator to a volume target of the EC2 gateway	239
No disks available when you try to add storage volumes message	240
How to remove a disk allocated as upload buffer space to reduce upload buffer space	240

Throughput to or from the EC2 gateway drops to zero	240
Activating AWS Support to help troubleshoot the gateway	240
Connect to your Amazon EC2 gateway using the serial console	242
Troubleshooting hardware appliance issues	242
How to determine service IP address	243
How to perform a factory reset	243
How to perform a remote restart	243
How to obtain Dell iDRAC support	243
How to find the hardware appliance serial number	243
How to get hardware appliance support	244
Troubleshooting volume issues	244
The Console Says That Your Volume Is Not Configured	245
The Console Says That Your Volume Is Irrecoverable	245
Your Cached Gateway is Unreachable And You Want to Recover Your Data	246
The Console Says That Your Volume Has PASS THROUGH Status	246
You Want to Verify Volume Integrity and Fix Possible Errors	247
Your Volume's iSCSI Target Doesn't Appear in Windows Disk Management Console	247
You Want to Change Your Volume's iSCSI Target Name	247
Your Scheduled Volume Snapshot Did Not Occur	247
You Need to Remove or Replace a Disk That Has Failed	248
Throughput from Your Application to a Volume Has Dropped to Zero	248
A Cache Disk in Your Gateway Encounters a Failure	249
A Volume Snapshot Has PENDING Status Longer Than Expected	249
High Availability Health Notifications	249
Troubleshooting high availability issues	250
Health notifications	250
Metrics	251
Best practices	252
Best practices: recovering your data	252
Recovering from an unexpected VM shutdown	253
Recovering data from malfunctioning gateway or VM	253
Recovering data from an irrecoverable volume	254
Recovering data from a malfunctioning cache disk	254
Recovering data from a corrupted file system	254
Recovering data from an inaccessible data center	256
Cleaning up unnecessary resources	256

Reducing the amount of billed storage on a volume	257
Additional Resources	258
Host setup	259
Deploy a default Amazon EC2 host for Volume Gateway	259
Deploy a customized Amazon EC2 instance for Volume Gateway	262
Modify Amazon EC2 instance metadata options	266
Synchronize VM time with Hyper-V or Linux KVM host time	266
Synchronize VM time with VMware host time	267
Configure paravirtualized disk controllers	269
Configuring network adapters for your gateway	269
Using VMware High Availability with Storage Gateway	274
Working with Volume Gateway storage resources	279
Removing Disks from Your Gateway	280
EBS Volumes for EC2 Gateways	281
Getting Activation Key	282
Linux (curl)	283
Linux (bash/zsh)	284
Microsoft Windows PowerShell	285
Using your local console	285
Connecting iSCSI Initiators	286
Connecting to your volumes from a Windows client	287
Connecting volumes to a Linux client	290
Customizing iSCSI Settings	292
Configuring CHAP Authentication	298
Using AWS Direct Connect with Storage Gateway	303
Port requirements for Volume Gateway	304
Getting the gateway IP address	311
Getting an IP Address from an Amazon EC2 Host	311
Understanding Resources and Resource IDs	313
Working with Resource IDs	313
Tagging Your Resources	314
Working with Tags	314
Open-Source Components	316
Storage Gateway quotas	316
Quotas for volumes	316
Recommended local disk sizes for your gateway	317

API Reference	319
Required Request Headers	319
Signing Requests	322
Example Signature Calculation	322
Error Responses	324
Exceptions	325
Operation Error Codes	327
Error Responses	346
Operations	348
Document history	349
Earlier updates	365
Release notes	384

What is Volume Gateway?

AWS Storage Gateway connects an on-premises software appliance with cloud-based storage to provide seamless integration with data security features between your on-premises IT environment and the AWS storage infrastructure. You can use the service to store data in the Amazon Web Services Cloud for scalable and cost-effective storage that helps maintain data security.

You can deploy Storage Gateway either on-premises as a VM appliance running on VMware ESXi, KVM, or Microsoft Hyper-V hypervisor, as a hardware appliance, or in AWS as an Amazon EC2 instance. You can use gateways hosted on EC2 instances for disaster recovery, data mirroring, and providing storage for applications hosted on Amazon EC2.

To see the wide range of use cases that AWS Storage Gateway helps make possible, see [AWS Storage Gateway](#). For current information about pricing, see [Pricing](#) on the AWS Storage Gateway details page.

AWS Storage Gateway offers file-based (S3 File Gateway and FSx File Gateway), volume-based (Volume Gateway), and tape-based (Tape Gateway) storage solutions.

This User Guide provides information related to Volume Gateway.

Volume Gateway provides cloud-backed storage volumes that you can mount as Internet Small Computer System Interface (iSCSI) devices from your on-premises application servers.

Volume Gateway supports the following volume configurations:

- **Cached volumes** – You store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally. Cached volumes offer a substantial cost savings on primary storage and minimize the need to scale your storage on-premises. You also retain low-latency access to your frequently accessed data.
- **Stored volumes** – If you need low-latency access to your entire dataset, first configure your on-premises gateway to store all your data locally. Then asynchronously back up point-in-time snapshots of this data to Amazon S3. This configuration provides durable and inexpensive offsite backups that you can recover to your local data center or Amazon Elastic Compute Cloud (Amazon EC2). For example, if you need replacement capacity for disaster recovery, you can recover the backups to Amazon EC2.

For an architectural overview, see [How Volume Gateway works](#).

In this User Guide, you can find a Getting Started section that covers setup information common to all gateway types. You can also find Volume Gateway setup requirements, and sections that describe how to deploy, activate, configure, and manage your Volume Gateway.

The procedures in this User Guide primarily focus on performing gateway operations by using the AWS Management Console. If you want to perform these operations programmatically, see the [AWS Storage Gateway API Reference](#).

How Volume Gateway works

Following, you can find an architectural overview of the Volume Gateway solution.

Volume Gateways

For Volume Gateways, you can use either cached volumes or stored volumes.

Topics

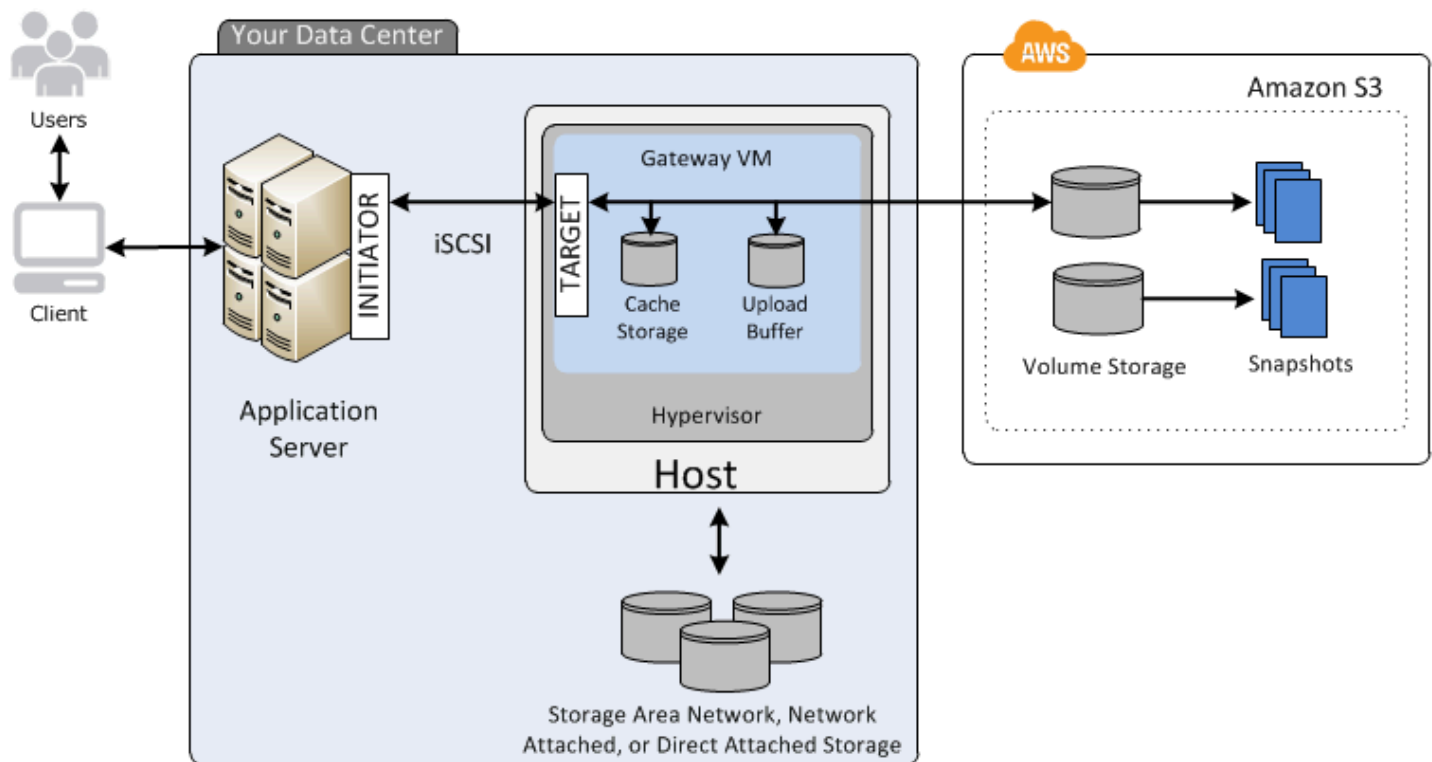
- [Cached volumes architecture](#)
- [Stored volumes architecture](#)

Cached volumes architecture

By using cached volumes, you can use Amazon S3 as your primary data storage, while retaining frequently accessed data locally in your Storage Gateway. Cached volumes minimize the need to scale your on-premises storage infrastructure, while still providing your applications with low-latency access to their frequently accessed data. You can create storage volumes up to 32 TiB in size and attach to them as iSCSI devices from your on-premises application servers. Your gateway stores data that you write to these volumes in Amazon S3 and retains recently read data in your on-premises Storage Gateway's cache and upload buffer storage.

Cached volumes can range from 1 GiB to 32 TiB in size and must be rounded to the nearest GiB. Each gateway configured for cached volumes can support up to 32 volumes for a total maximum storage volume of 1,024 TiB (1 PiB).

In the cached volumes solution, Storage Gateway stores all your on-premises application data in a storage volume in Amazon S3. The following diagram provides an overview of the cached volumes deployment.



After you install the Storage Gateway software appliance—the VM—on a host in your data center and activate it, you use the AWS Management Console to provision storage volumes backed by Amazon S3. You can also provision storage volumes programmatically using the Storage Gateway API or the AWS SDK libraries. You then mount these storage volumes to your on-premises application servers as iSCSI devices.

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

- **Disks for use by the gateway as cache storage** – As your applications write data to the storage volumes in AWS, the gateway first stores the data on the on-premises disks used for cache storage. Then the gateway uploads the data to Amazon S3. The cache storage acts as the on-premises durable store for data that is waiting to upload to Amazon S3 from the upload buffer.

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

You can use the following guidelines to determine the amount of disk space to allocate for cache storage. Generally, you should allocate at least 20 percent of your existing file store size as cache

storage. Cache storage should also be larger than the upload buffer. This guideline helps make sure that cache storage is large enough to persistently hold all data in the upload buffer that has not yet been uploaded to Amazon S3.

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

You can take incremental backups, called *snapshots*, of your storage volumes in Amazon S3. These point-in-time snapshots are also stored in Amazon S3 as Amazon EBS snapshots. When you take a new snapshot, only the data that has changed since your last snapshot is stored. When the snapshot is taken, the gateway uploads the changes up to the snapshot point, then creates the new snapshot using Amazon EBS. You can initiate snapshots on a scheduled or one-time basis. A single volume supports queueing multiple snapshots in rapid succession, but each snapshot must finish being created before the next can be taken. When you delete a snapshot, only the data not needed for any other snapshots is removed. For information about Amazon EBS snapshots, see [Amazon EBS snapshots](#).

You can restore an Amazon EBS snapshot to a gateway storage volume if you need to recover a backup of your data. Alternatively, for snapshots up to 16 TiB in size, you can use the snapshot as a starting point for a new Amazon EBS volume. You can then attach this new Amazon EBS volume to an Amazon EC2 instance.

All gateway data and snapshot data for cached volumes is stored in Amazon S3 and encrypted at rest using server-side encryption (SSE). However, you can't access this data with the Amazon S3 API or other tools such as the Amazon S3 Management Console.


Stored volumes architecture

By using stored volumes, you can store your primary data locally, while asynchronously backing up that data to AWS. Stored volumes provide your on-premises applications with low-latency access to their entire datasets. At the same time, they provide durable, offsite backups. You can create storage volumes and mount them as iSCSI devices from your on-premises application servers. Data written to your stored volumes is stored on your on-premises storage hardware. This data is asynchronously backed up to Amazon S3 as Amazon Elastic Block Store (Amazon EBS) snapshots.

Stored volumes can range from 1 GiB to 16 TiB in size and must be rounded to the nearest GiB. Each gateway configured for stored volumes can support up to 32 volumes and a total volume storage of 512 TiB (0.5 PiB).

With stored volumes, you maintain your volume storage on-premises in your data center. That is, you store all your application data on your on-premises storage hardware. Then, using features that help maintain data security, the gateway uploads data to the Amazon Web Services Cloud for cost-effective backup and rapid disaster recovery. This solution is ideal if you want to keep data locally on-premises, because you need to have low-latency access to all your data, and also to maintain backups in AWS.

The following diagram provides an overview of the stored volumes deployment.



After you install the Storage Gateway software appliance—the VM—on a host in your data center and activated it, you can create gateway *storage volumes*. You then map them to on-premises direct-attached storage (DAS) or storage area network (SAN) disks. You can start with either new disks or disks already holding data. You can then mount these storage volumes to your on-premises application servers as iSCSI devices. As your on-premises applications write data to and read data from a gateway's storage volume, this data is stored and retrieved from the volume's assigned disk.

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

You can take incremental backups, called *snapshots*, of your storage volumes. The gateway stores these snapshots in Amazon S3 as Amazon EBS snapshots. When you take a new snapshot, only the data that has changed since your last snapshot is stored. When the snapshot is taken, the gateway uploads the changes up to the snapshot point, then creates the new snapshot using Amazon EBS. You can initiate snapshots on a scheduled or one-time basis. A single volume supports queueing multiple snapshots in rapid succession, but each snapshot must finish being created before the next can be taken. 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 if you need to recover a backup of your data. You can also use the snapshot as a starting point for a new Amazon EBS volume, which you can then attach to an Amazon EC2 instance.

Getting started with AWS Storage Gateway

This section provides instructions for getting started with AWS. You need an AWS account before you can start using AWS Storage Gateway. You can use an existing AWS account, or sign up for a new account. You also need an IAM user in your AWS account that belongs to a group with the necessary administrative permissions to perform Storage Gateway tasks. Users with the appropriate privileges can access the Storage Gateway console and Storage Gateway API to perform gateway deployment, configuration, and maintenance tasks. If you are a first-time user, we recommend that you review the [Supported AWS regions](#) and [Volume Gateway setup requirements](#) sections before you begin working with Storage Gateway.

This section contains the following topics, which provide additional information about getting started with AWS Storage Gateway:

Topics

- [Sign Up for AWS Storage Gateway](#) - Learn how to sign up for AWS and create an AWS account.
- [Create an IAM user with administrator privileges](#) - Learn how to create an IAM user with administrative privileges for your AWS account.
- [Accessing AWS Storage Gateway](#) - Learn how to access AWS Storage Gateway through the Storage Gateway console or programmatically using the AWS SDKs.
- [AWS Regions that support Storage Gateway](#) - Learn which AWS Regions you can use to store your data when you activate your gateway in Storage Gateway.

Sign Up for AWS Storage Gateway

An AWS account is a fundamental requirement for accessing AWS services. Your AWS account is the basic container for all of the AWS resources you create as an AWS user. Your AWS account is also the basic security boundary for your AWS resources. Any resources that you create in your account are available to users who have credentials for the account. Before you can start using AWS Storage Gateway, you need to sign up for an AWS account.

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

1. Open <https://portal.aws.amazon.com/billing/signup>.

2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform [tasks that require root user access](#).

We also recommend that you require your users to use temporary credentials when accessing AWS. To provide temporary credentials, you can use federation and an identity provider, such as AWS IAM Identity Center. If your company already uses an identity provider, you can use it with federation to simplify how you provide access to the resources in your AWS account.

Create an IAM user with administrator privileges

After you create your AWS account, use the following steps to create an AWS Identity and Access Management (IAM) user for yourself, and then add that user to a group that has administrative permissions. For more information about using the AWS Identity and Access Management service to control access to Storage Gateway resources, see [Identity and Access Management for AWS Storage Gateway](#).

To create an administrator user, choose one of the following options.

Choose one way to manage your administrator	To	By	You can also
In IAM Identity Center	Use short-term credentials to access AWS.	Following the instructions in Getting started in the <i>AWS IAM Identity Center User Guide</i> .	Configure programmatic access by Configuring the AWS CLI to use AWS IAM Identity Center in the AWS

Choose one way to manage your administrator	To	By	You can also
(Recommended)	This aligns with the security best practices . For information about best practices , see Security best practices in IAM in the <i>IAM User Guide</i> .		<i>Command Line Interface User Guide</i> .
In IAM (Not recommended)	Use long-term credentials to access AWS.	Following the instructions in Creating your first IAM admin user and user group in the <i>IAM User Guide</i> .	Configure programmatic access by Managing access keys for IAM users in the <i>IAM User Guide</i> .

Warning

IAM users have long-term credentials which present a security risk. To help mitigate this risk, we recommend that you provide these users with only the permissions they require to perform the task and that you remove these users when they are no longer needed.

Accessing AWS Storage Gateway

You can use the [AWS Storage Gateway console](#) to perform various gateway configuration and maintenance tasks, including activating or removing Storage Gateway hardware appliances from your deployment, creating, managing, and deleting the different types of gateways, creating, managing, and deleting storage volumes, and monitoring the health and status of various elements of the Storage Gateway service. For simplicity and ease of use, this guide focuses on performing tasks using the Storage Gateway console web interface. You can access the Storage

Gateway console through your web browser at: <https://console.aws.amazon.com/storagegateway/home/>.

If you prefer a programmatic approach, you can use the AWS Storage Gateway Application Programming Interface (API) or Command Line Interface (CLI) to set up and manage the resources in your Storage Gateway deployment. For more information about actions, data types, and required syntax for the Storage Gateway API, see the [Storage Gateway API Reference](#). For more information about the Storage Gateway CLI, see the [AWS CLI Command Reference](#).

You can also use the AWS SDKs to develop applications that interact with Storage Gateway. The AWS SDKs for Java, .NET, and PHP wrap the underlying Storage Gateway API to simplify your programming tasks. For information about downloading the SDK libraries, see the [AWS Developer Center](#).

For information about pricing, see [AWS Storage Gateway pricing](#).

AWS Regions that support Storage Gateway

An AWS Region is a physical location in the world where AWS has multiple Availability Zones. Availability Zones consist of one or more discrete AWS data centers, each with redundant power, networking, and connectivity, housed in separate facilities. This means that each AWS Region is physically isolated and independent of the other Regions. Regions provide fault tolerance, stability, and resilience, and can also reduce latency. The resources that you create in one Region do not exist in any other Region unless you explicitly use a replication feature offered by an AWS service. For example, Amazon S3 and Amazon EC2 support cross-Region replication. Some services, such as AWS Identity and Access Management, do not have Regional resources. You can launch AWS resources in locations that meet your business requirements. For example, you might want to launch Amazon EC2 instances to host your AWS Storage Gateway appliances in an AWS Region in Europe to be closer to your European users, or to meet legal requirements. Your AWS account determines which of the Regions supported by a specific service are available for you to use.

- Storage Gateway—For supported AWS Regions and a list of AWS service endpoints you can use with Storage Gateway, see [AWS Storage Gateway Endpoints and Quotas](#) in the *AWS General Reference*.
- Storage Gateway Hardware Appliance—For supported AWS Regions you can use with the hardware appliance, see [AWS Storage Gateway Hardware Appliance Regions](#) in the *AWS General Reference*.

Requirements for setting up Volume Gateway

Unless otherwise noted, the following requirements are common to all gateway configurations.

Topics

- [Hardware and storage requirements](#)
- [Network and firewall requirements](#)
- [Supported hypervisors and host requirements](#)
- [Supported iSCSI initiators](#)

Hardware and storage requirements

This section describes the minimum hardware and settings for your gateway and the minimum amount of disk space to allocate for the required storage.

Hardware requirements for VMs

When deploying your gateway, you must make sure that the underlying hardware on which you deploy the gateway VM can dedicate the following minimum resources:

- Four virtual processors assigned to the VM.
- For Volume Gateway, your hardware should dedicate the following amounts of RAM:
 - 16 GiB of reserved RAM for gateways with cache size up to 16 TiB
 - 32 GiB of reserved RAM for gateways with cache size 16 TiB to 32 TiB
 - 48 GiB of reserved RAM for gateways with cache size 32 TiB to 64 TiB
- 80 GiB of disk space for installation of VM image and system data.

For more information, see [Optimizing gateway performance](#). For information about how your hardware affects the performance of the gateway VM, see [AWS Storage Gateway quotas](#).

Requirements for Amazon EC2 instance types

When deploying your gateway on Amazon Elastic Compute Cloud (Amazon EC2), the instance size must be at least **xlarge** for your gateway to function. However, for the compute-optimized instance family the size must be at least **2xlarge**.

Note

The Storage Gateway AMI is only compatible with x86-based instances that use Intel or AMD processors. ARM-based instances that use Graviton processors are not supported.

For Volume Gateway, your Amazon EC2 instance should dedicate the following amounts of RAM depending on the cache size you plan to use for your gateway:

- 16 GiB of reserved RAM for gateways with cache size up to 16 TiB
- 32 GiB of reserved RAM for gateways with cache size 16 TiB to 32 TiB
- 48 GiB of reserved RAM for gateways with cache size 32 TiB to 64 TiB

Use one of the following instance types recommended for your gateway type.

Recommended for cached volumes and Tape Gateway types

- General-purpose instance family – **m4, m5, or m6** instance type.

Note

We don't recommend using the **m4.16xlarge** instance type.

- Compute-optimized instance family – **c4, c5, or c6** instance types. Choose the **2xlarge** instance size or higher to meet the required RAM requirements.
- Memory-optimized instance family – **r3, r5, or r6** instance types.
- Storage-optimized instance family – **i3 or i4** instance types.

Storage requirements

In addition to 80 GiB disk space for the VM, you also need additional disks for your gateway.

The following table recommends sizes for local disk storage for your deployed gateway.

Gateway Type	Cache (Minimum)	Cache (Maximum)	Upload Buffer (Minimum)	Upload Buffer (Maximum)	Other Required Local Disks
Cached Volume Gateway	150 GiB	64 TiB	150 GiB	2 TiB	—
Stored Volume Gateway	—	—	150 GiB	2 TiB	1 or more for stored volume or volumes

Note

You can configure one or more local drives for your cache and upload buffer, up to the maximum capacity.

When adding cache or upload buffer to an existing gateway, it's important to create new disks in your host (hypervisor or Amazon EC2 instance). Don't change the size of existing disks if the disks have been previously allocated as either a cache or upload buffer.

For information about gateway quotas, see [AWS Storage Gateway quotas](#).

Network and firewall requirements

Your gateway requires access to the internet, local networks, Domain Name Service (DNS) servers, firewalls, routers, and so on. Following, you can find information about required ports and how to allow access through firewalls and routers.

Note

In some cases, you might deploy Storage Gateway on Amazon EC2 or use other types of deployment (including on-premises) with network security policies that restrict AWS IP address ranges. In these cases, your gateway might experience service connectivity issues when the AWS IP range values changes. The AWS IP address range values that you need to

use are in the Amazon service subset for the AWS Region that you activate your gateway in. For the current IP range values, see [AWS IP address ranges](#) in the *AWS General Reference*.

Note

Network bandwidth requirements vary based on the quantity of data that is uploaded and downloaded by the gateway. A minimum of 100Mbps is required to successfully download, activate, and update the gateway. Your data transfer patterns will determine the bandwidth necessary to support your workload. In some cases, you might deploy Storage Gateway on Amazon EC2 or use other types of deployment

Topics

- [Port requirements](#)
- [Networking and firewall requirements for the Storage Gateway Hardware Appliance](#)
- [Allowing AWS Storage Gateway access through firewalls and routers](#)
- [Configuring security groups for your Amazon EC2 gateway instance](#)

Port requirements

Storage Gateway requires certain ports to be allowed for its operation. The following illustrations show the required ports that you must allow for each type of gateway. Some ports are required by all gateway types, and others are required by specific gateway types. For more information about port requirements, see [Port requirements for Volume Gateway](#).

Common ports for all gateway types

The following ports are common to all gateway types and are required by all gateway types.

Protocol	Port	Direction	Source	Destination	How Used
TCP	443 (HTTPS)	Outbound	Storage Gateway	AWS	For communication from Storage

Protocol	Port	Direction	Source	Destination	How Used
					Gateway to the AWS service endpoint. For information about service endpoints, see Allowing AWS Storage Gateway access through firewalls and routers.

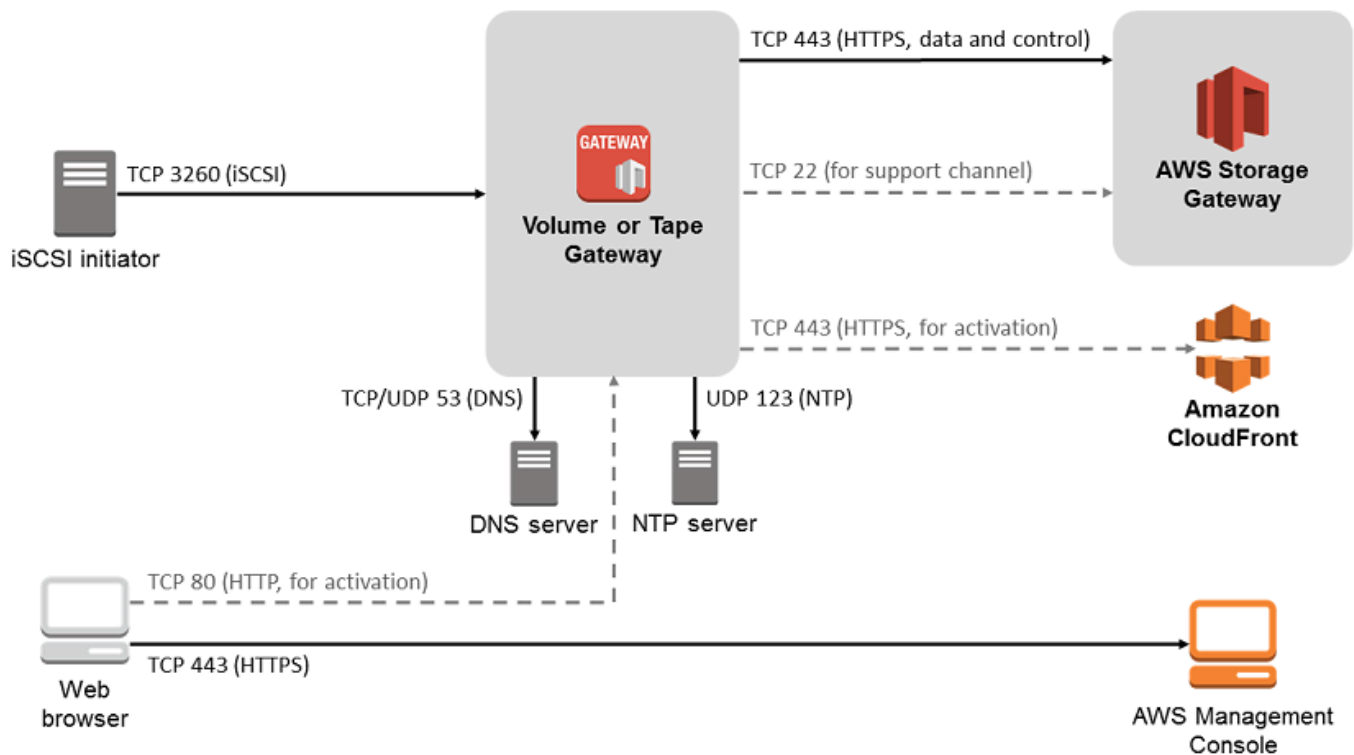
Protocol	Port	Direction	Source	Destination	How Used
TCP	80 (HTTP)	Inbound	The host from which you connect to the AWS Management Console.	Storage Gateway	<p>By local systems to obtain the Storage Gateway activation key. Port 80 is only used during activation of the Storage Gateway appliance.</p> <p>Storage Gateway does not require port 80 to be publicly accessible. The required level of access to port 80 depends on your network configuration. If you activate your gateway from the Storage Gateway Management</p>

Protocol	Port	Direction	Source	Destination	How Used
					Console, the host from which you connect to the console must have access to your gateway's port 80.
TCP/UDP	53 (DNS)	Outbound	Storage Gateway	Domain Name Service (DNS) server	For communication between Storage Gateway and the DNS server.

Protocol	Port	Direction	Source	Destination	How Used
TCP	22 (Support channel)	Outbound	Storage Gateway	AWS Support	Allows AWS Support to access your gateway to help you with troubleshooting gateway issues. You don't need this port open for the normal operation of your gateway, but it is required for troubleshooting.
UDP	123 (NTP)	Outbound	NTP client	NTP server	Used by local systems to synchronize VM time to the host time.

Ports for volume and Tape Gateways

The following illustration shows the ports to open for Volume Gateway.



In addition to the common ports, Volume Gateway requires the following port.

Protocol	Port	Direction	Source	Destination	How Used
TCP	3260 (iSCSI)	Inbound	iSCSI Initiators	Storage Gateway	By local systems to connect to iSCSI targets exposed by the gateway.

For detailed information about port requirements, see [Port requirements for Volume Gateway](#) in the *Additional Storage Gateway resources* section.

Networking and firewall requirements for the Storage Gateway Hardware Appliance


Each Storage Gateway Hardware Appliance requires the following network services:

- **Internet access** – an always-on network connection to the internet through any network interface on the server.
- **DNS services** – DNS services for communication between the hardware appliance and DNS server.
- **Time synchronization** – an automatically configured Amazon NTP time service must be reachable.
- **IP address** – A DHCP or static IPv4 address assigned. You cannot assign an IPv6 address.

There are five physical network ports at the rear of the Dell PowerEdge R640 server. From left to right (facing the back of the server) these ports are as follows:

1. iDRAC
2. em1
3. em2
4. em3
5. em4

You can use the iDRAC port for remote server management.



A hardware appliance requires the following ports to operate.

Protocol	Port	Direction	Source	Destination	How Used
SSH	22	Outbound	Hardware appliance	54.201.223.107	Support channel

Protocol	Port	Direction	Source	Destination	How Used
DNS	53	Outbound	Hardware appliance	DNS servers	Name resolution
UDP/NTP	123	Outbound	Hardware appliance	*.amazon.pool.ntp.org	Time synchronization
HTTPS	443	Outbound	Hardware appliance	*.amazonaws.com	Data transfer
HTTP	8080	Inbound	AWS	Hardware appliance	Activation (only briefly)


To perform as designed, a hardware appliance requires network and firewall settings as follows:

- Configure all connected network interfaces in the hardware console.
- Make sure that each network interface is on a unique subnet.
- Provide all connected network interfaces with outbound access to the endpoints listed in the diagram preceding.
- Configure at least one network interface to support the hardware appliance. For more information, see [Configuring hardware appliance network parameters](#).

Note

For an illustration showing the back of the server with its ports, see [Physically installing your hardware appliance](#)

All IP addresses on the same network interface (NIC), whether for a gateway or a host, must be on the same subnet. The following illustration shows the addressing scheme.



For more information on activating and configuring a hardware appliance, see [Using the Storage Gateway Hardware Appliance](#).

Allowing AWS Storage Gateway access through firewalls and routers

Your gateway requires access to the following service endpoints to communicate with AWS. If you use a firewall or router to filter or limit network traffic, you must configure your firewall and router to allow these service endpoints for outbound communication to AWS.

Note

If you configure private VPC endpoints for your Storage Gateway to use for connection and data transfer to and from AWS, your gateway does not require access to the public internet. For more information, see [Activating a gateway in a virtual private cloud](#).

Important

Depending on your gateway's AWS Region, replace *region* in the service endpoint with the correct region string.

The following service endpoint is required by all gateways for head-bucket operations.


```
s3.amazonaws.com:443
```

The following service endpoints are required by all gateways for control path (anon-cp, client-cp, proxy-app) and data path (dp-1) operations.

```
anon-cp.storagegateway.region.amazonaws.com:443  
client-cp.storagegateway.region.amazonaws.com:443  
proxy-app.storagegateway.region.amazonaws.com:443  
dp-1.storagegateway.region.amazonaws.com:443
```

The following gateway service endpoint is required to make API calls.

```
storagegateway.region.amazonaws.com:443
```

The following example is a gateway service endpoint in the US West (Oregon) Region (us-west-2).

```
storagegateway.us-west-2.amazonaws.com:443
```

The Amazon S3 service endpoint, shown following, is used by File Gateways only. A File Gateway requires this endpoint to access the S3 bucket that a file share maps to.

```
bucketname.s3.region.amazonaws.com
```

The following example is an S3 service endpoint in the US East (Ohio) Region (us-east-2).

```
s3.us-east-2.amazonaws.com
```

Note

If your gateway can't determine the AWS Region where your S3 bucket is located, this service endpoint defaults to `s3.us-east-1.amazonaws.com`. We recommend that you allow access to the US East (N. Virginia) Region (us-east-1) in addition to AWS Regions where your gateway is activated, and where your S3 bucket is located.

The following are S3 service endpoints for AWS GovCloud (US) Regions.

```
s3-fips.us-gov-west-1.amazonaws.com (AWS GovCloud (US-West) Region (FIPS))
s3-fips.us-gov-east-1.amazonaws.com (AWS GovCloud (US-East) Region (FIPS))
s3.us-gov-west-1.amazonaws.com (AWS GovCloud (US-West) Region (Standard))
s3.us-gov-east-1.amazonaws.com (AWS GovCloud (US-East) Region (Standard))
```

The following example is a FIPS service endpoint for an S3 bucket in the AWS GovCloud (US-West) Region.

```
bucket-name.s3-fips.us-gov-west-1.amazonaws.com
```

A Storage Gateway VM is configured to use the following NTP servers.

```
0.amazon.pool.ntp.org
1.amazon.pool.ntp.org
2.amazon.pool.ntp.org
3.amazon.pool.ntp.org
```

- Storage Gateway—For supported AWS Regions and a list of AWS service endpoints you can use with Storage Gateway, see [AWS Storage Gateway endpoints and quotas](#) in the *AWS General Reference*.
- Storage Gateway Hardware Appliance—For supported AWS Regions you can use with the hardware appliance see [Storage Gateway hardware appliance regions](#) in the *AWS General Reference*.

Configuring security groups for your Amazon EC2 gateway instance

A security group controls traffic to your Amazon EC2 gateway instance. When you configure a security group, we recommend the following:

- The security group should not allow incoming connections from the outside internet. It should allow only instances within the gateway security group to communicate with the gateway. If you need to allow instances to connect to the gateway from outside its security group, we recommend that you allow connections only on ports 3260 (for iSCSI connections) and 80 (for activation).
- If you want to activate your gateway from an Amazon EC2 host outside the gateway security group, allow incoming connections on port 80 from the IP address of that host. If you cannot

determine the activating host's IP address, you can open port 80, activate your gateway, and then close access on port 80 after completing activation.

- Allow port 22 access only if you are using AWS Support for troubleshooting purposes. For more information, see [You want AWS Support to help troubleshoot your EC2 gateway](#).

In some cases, you might use an Amazon EC2 instance as an initiator (that is, to connect to iSCSI targets on a gateway that you deployed on Amazon EC2). In such a case, we recommend a two-step approach:

1. You should launch the initiator instance in the same security group as your gateway.
2. You should configure access so the initiator can communicate with your gateway.

For information about the ports to open for your gateway, see [Port requirements for Volume Gateway](#).

Supported hypervisors and host requirements

You can run Storage Gateway on-premises as either a virtual machine (VM) appliance, or a physical hardware appliance, or in AWS as an Amazon EC2 instance.

Note

When a manufacturer ends general support for a hypervisor version, Storage Gateway also ends support for that hypervisor version. For detailed information about support for specific versions of a hypervisor, see the manufacturer's documentation.

Storage Gateway supports the following hypervisor versions and hosts:

- VMware ESXi Hypervisor (version 7.0 or 8.0) – For this setup, you also need a VMware vSphere client to connect to the host.
- Microsoft Hyper-V Hypervisor (version 2012 R2, 2016, 2019, or 2022) – A free, standalone version of Hyper-V is available at the [Microsoft Download Center](#). For this setup, you need a Microsoft Hyper-V Manager on a Microsoft Windows client computer to connect to the host.
- Linux Kernel-based Virtual Machine (KVM) – A free, open-source virtualization technology. KVM is included in all versions of Linux version 2.6.20 and newer. Storage Gateway is tested and

supported for the CentOS/RHEL 7.7, Ubuntu 16.04 LTS, and Ubuntu 18.04 LTS distributions. Any other modern Linux distribution may work, but function or performance is not guaranteed. We recommend this option if you already have a KVM environment up and running and you are already familiar with how KVM works.

- Amazon EC2 instance – Storage Gateway provides an Amazon Machine Image (AMI) that contains the gateway VM image. Only file, cached volume, and Tape Gateway types can be deployed on Amazon EC2. For information about how to deploy a gateway on Amazon EC2, see [Deploy a customized Amazon EC2 instance for Volume Gateway](#).
- Storage Gateway Hardware Appliance – Storage Gateway provides a physical hardware appliance as a on-premises deployment option for locations with limited virtual machine infrastructure.

Note

Storage Gateway doesn't support recovering a gateway from a VM that was created from a snapshot or clone of another gateway VM or from your Amazon EC2 AMI. If your gateway VM malfunctions, activate a new gateway and recover your data to that gateway. For more information, see [Recovering from an unexpected virtual machine shutdown](#).

Storage Gateway doesn't support dynamic memory and virtual memory ballooning.

Supported iSCSI initiators

When you deploy a cached volume or stored Volume Gateway, you can create iSCSI storage volumes on your gateway.

To connect to these iSCSI devices, Storage Gateway supports the following iSCSI initiators:

- Windows Server 2019
- Windows Server 2016
- Windows Server 2012 R2
- Windows 10
- Windows 8.1
- Red Hat Enterprise Linux 5
- Red Hat Enterprise Linux 6
- Red Hat Enterprise Linux 7

- VMware ESX Initiator, which provides an alternative to using initiators in the guest operating systems of your VMs

 **Important**

Storage Gateway doesn't support Microsoft Multipath I/O (MPIO) from Windows clients. Storage Gateway supports connecting multiple hosts to the same volume if the hosts coordinate access by using Windows Server Failover Clustering (WSFC). However, you can't connect multiple hosts to that same volume (for example, sharing a nonclustered NTFS/ ext4 file system) without using WSFC.

Using the Storage Gateway Hardware Appliance

The Storage Gateway Hardware Appliance is a physical hardware appliance with the Storage Gateway software preinstalled on a validated server configuration. You can manage the hardware appliances in your deployment from the **Hardware appliance overview** page in the AWS Storage Gateway console.

The hardware appliance is a high-performance 1U server that you can deploy in your data center, or on-premises inside your corporate firewall. When you buy and activate your hardware appliance, the activation process associates the hardware appliance with your AWS account. After activation, your hardware appliance appears in the console as a gateway on the **Hardware appliance overview** page. You can configure the hardware appliance as an S3 File Gateway, FSx File Gateway, Tape Gateway, or Volume Gateway type. The procedure that you use to deploy and activate these gateway types on a hardware appliance is same as on a virtual platform.

For a list of supported AWS Regions where the Storage Gateway Hardware Appliance is available for activation and use, see [Storage Gateway Hardware Appliance Regions](#) in the *AWS General Reference*.

In the sections that follow, you can find instructions about how to set up, rack mount, power, configure, activate, launch, use, and delete an Storage Gateway Hardware Appliance.

Topics

- [Setting up your Storage Gateway Hardware Appliance](#)
- [Physically installing your hardware appliance](#)
- [Accessing the hardware appliance console](#)
- [Configuring hardware appliance network parameters](#)
- [Activating your Storage Gateway Hardware Appliance](#)
- [Creating a gateway on your hardware appliance](#)
- [Configuring a gateway IP address on the hardware appliance](#)
- [Removing gateway software from your hardware appliance](#)
- [Deleting your Storage Gateway Hardware Appliance](#)

Setting up your Storage Gateway Hardware Appliance

After you receive your Storage Gateway Hardware Appliance, you use the hardware appliance console to configure networking to provide an always-on connection to AWS and activate your appliance. Activation associates your appliance with the AWS account that is used during the activation process. After the appliance is activated, you can launch a file, volume, or Tape Gateway from the Storage Gateway console.

To install and configure your hardware appliance

1. Rack-mount the appliance, and plug in power and network connections. For more information, see [Physically installing your hardware appliance](#).
2. Set the Internet Protocol version 4 (IPv4) addresses for both the hardware appliance (the host) and Storage Gateway (the service). For more information, see [Configuring hardware appliance network parameters](#).
3. Activate the hardware appliance on the console **Hardware appliance overview** page in the AWS Region of your choice. For more information, see [Activating your Storage Gateway Hardware Appliance](#).
4. Install the Storage Gateway on your hardware appliance. For more information, see [Creating a Volume Gateway](#).

You set up gateways on your hardware appliance the same way that you set up gateways on VMware ESXi, Microsoft Hyper-V, Linux Kernel-based Virtual Machine (KVM), or Amazon EC2.

Increasing the usable cache storage

You can increase the usable storage on the hardware appliance from 5 TB to 12 TB. Doing this provides a larger cache for low latency access to data in AWS. If you ordered the 5 TB model, you can increase the usable storage to 12 TB by buying five 1.92 TB SSDs (solid state drives).

You can then add them to the hardware appliance before you activate it. If you have already activated the hardware appliance and want to increase the usable storage on the appliance to 12 TB, do the following:

1. Reset the hardware appliance to its factory settings. Contact AWS Support for instructions on how to do this.
2. Add five 1.92 TB SSDs to the appliance.

Network interface card options


Depending on the model of appliance you ordered, it may come with a 10G-Base-T copper network card or a 10G DA/SFP+ network card.

- 10G-Base-T NIC configuration:
 - Use CAT6 cables for 10G or CAT5(e) for 1G
- 10G DA/SFP+ NIC configuration:
 - Use Twinax copper Direct Attach Cables up to 5 meters
 - Dell/Intel compatible SFP+ optical modules (SR or LR)
 - SFP/SFP+ copper transceiver for 1G-Base-T or 10G-Base-T

Physically installing your hardware appliance

Your appliance has a 1U form factor and fits in a standard International Electrotechnical Commission (IEC) compliant 19-inch rack. The following image shows the dimensions of the hardware appliance:

hardware appliance dimensions including mounting brackets and bezel.



System	Xa	Xb	Y	Za (with bezel)	Za (without bezel)	Zb*	Zc
10 x 2.5-inches	482.0 mm (18.97-inches)	434.0 mm (17.08-inches)	42.8 mm (1.68-inches)	35.84 mm (1.41-inches)	22.0 mm (0.87-inches)	733.82 mm (29.61-inches)	772.67 mm (30.42-inches)

hardware appliance dimensions including mounting brackets and bezel.

Prerequisites

To install your hardware appliance, you need the following components:

- Power cables: one required, two recommended.

- Supported network cabling (depending on which Network Interface Card (NIC) is included in the hardware appliance). Twinax Copper DAC, SFP+ optical module (Intel compatible) or SFP to Base-T copper transceiver.
- Keyboard and monitor, or a keyboard, video, and mouse (KVM) switch solution.

Note


Before you perform the following procedure, make sure that you meet all of the requirements for the Storage Gateway Hardware Appliance as described in [Networking and firewall requirements for the Storage Gateway Hardware Appliance](#).

To physically install your hardware appliance

1. Unbox your hardware appliance and follow the instructions contained in the box to rack-mount the server.

The following image shows the back of the hardware appliance with ports for connecting power, ethernet, USB, monitor, keyboard, and iDRAC.

hardware appliance rear with network and power connector labels.



hardware appliance rear with network and power connector labels.

2. Plug in a power connection to each of the two power supplies. It's possible to plug in to only one power connection, but we recommend power connections to both power supplies.
3. Plug an Ethernet cable into the em1 port to provide an always-on internet connection. The em1 port is the first of the four physical network ports on the rear, from left to right.

Note

The hardware appliance doesn't support VLAN trunking. Set up the switch port to which you are connecting the hardware appliance as a non-trunked VLAN port.

4. Plug in the keyboard and monitor.
5. Power on the server by pressing the **Power** button on the front panel, as shown in the following image.
hardware appliance front with power button label.



hardware appliance front with power button label.

Next step

[Accessing the hardware appliance console](#)

Accessing the hardware appliance console

When you power on your hardware appliance, the hardware appliance console appears on the monitor. The hardware appliance console presents a user interface specific to AWS that you can use to set an administrator password, configure initial network parameters, and open a support channel to AWS.

To work with the hardware appliance console, enter text from the keyboard and use the Up, Down, Right, and Left Arrow keys to move about the screen in the indicated direction. Use the Tab key to move forward in order through items on-screen. On some setups, you can use the Shift +Tab keystroke to move sequentially backward. Use the Enter key to save selections, or to choose a button on the screen.

The first time the hardware appliance console appears, the **Welcome** page is displayed, and you are prompted to set a password for the *admin* user account before you can access the console.

To set an admin password

- At the **Please set your login password** prompt, do the following:
 - a. For **Set Password**, enter a password, and then press Down arrow.

- b. For **Confirm**, re-enter your password, and then choose **Save Password**.

After you set your password, the hardware console **Home** page appears. The **Home** page displays network information for the **em1**, **em2**, **em3**, and **em4** network interfaces, and has the following menu options:

- Configure Network
- Open Guest Console
- Change Password
- Logout
- Open Support Console

Next step

[Configuring hardware appliance network parameters](#)

Configuring hardware appliance network parameters

After the hardware appliance boots up and you set your admin user password in the hardware console as described in [Accessing the hardware appliance console](#), use the following procedure to configure network parameters so your hardware appliance can connect to AWS.

To set a network address


1. From the **Home** page, choose **Configure Network** and then press **Enter**. The **Configure Network** page appears. The **Configure Network** page shows IP and DNS information for each of the 4 network interfaces on the hardware appliance, and includes menu options to configure **DHCP** or **Static** addresses for each.
2. For the **em1** interface, do one of the following:
 - Choose **DHCP** and press **Enter** to use the IPv4 address assigned by your Dynamic Host Configuration Protocol (DHCP) server to your physical network port.

Note this address for later use in the activation step.

- Choose **Static** and press **Enter** to configure a static IPv4 address.

Enter a valid **IP Address**, **Subnet Mask**, **Gateway**, and **DNS** server address for the **em1** network interface.

When finished, choose **Save** and then press **Enter** to save the configuration.

 **Note**

You can use this procedure to configure other network interfaces in addition to **em1** for redundancy. If you configure other interfaces, they must provide the same always-on connection to the AWS endpoints listed in the requirements.

Network bonding and Link Aggregation Control Protocol (LACP) are not supported by the hardware appliance or by Storage Gateway.

We do not recommend configuring multiple network interfaces on the subnet as this can sometimes cause routing issues.

To log out of the hardware console

1. Choose **Back** and press **Enter** to return to the **Home** page.
2. Choose **Logout** and press **Enter** to return to the **Welcome** page.

Next step

[Activating your Storage Gateway Hardware Appliance](#)

Activating your Storage Gateway Hardware Appliance

After configuring your IP address, you enter this IP address on the **Hardware** page of the AWS Storage Gateway console to activate your hardware appliance. The activation process validates that your hardware appliance has the appropriate security credentials and registers the appliance to your AWS account.

You can choose to activate your hardware appliance in any of the supported AWS Regions. For a list of supported AWS Regions, see [Storage Gateway Hardware Appliance Regions](#) in the *AWS General Reference*.

To activate your Storage Gateway Hardware Appliance

1. Open the [AWS Storage Gateway Management Console](#) and sign in with the account credentials you want to use to activate your hardware.

Note

For activation only, the following must be true:

- Your browser must be on the same network as your hardware appliance.
- Your firewall must allow HTTP access on port 8080 to the appliance for inbound traffic.

2. Choose **Hardware** from the navigation menu on the left side of the page.
3. Choose **Activate appliance**.
4. For **IP Address**, enter the IP address that you configured for your hardware appliance, then choose **Connect**.

For more information about configuring the IP address, see [Configuring network parameters](#).

5. For **Name**, enter a name for your hardware appliance. Names can be up to 255 characters long and can't include a slash character.
6. For **Hardware appliance time zone**, enter the local time zone from which most of the workload for the gateway will be generated., then choose **Next**.

The time zone controls when hardware updates take place, with 2 a.m. used as the default scheduled time to perform updates. Ideally, if the time zone is set properly, updates will take place outside of the local working day window by default.

7. Review the activation parameters in the Hardware appliance detail section. You can choose **Previous** to go back and make changes if necessary. Otherwise, choose **Activate** to finish the activation.

A banner appears on the **Hardware appliance overview** page, indicating that the hardware appliance has been successfully activated.

At this point, the appliance is associated with your account. The next step is to configure and launch an S3 File Gateway, FSx File Gateway, Tape Gateway, or Volume Gateway on the new appliance.

Next step

[Creating a gateway on your hardware appliance](#)

Creating a gateway on your hardware appliance

You can install gateway appliance software to create an S3 File Gateway, FSx File Gateway, Tape Gateway, or Volume Gateway on any Storage Gateway Hardware Appliance in your deployment.

To create a gateway on your hardware appliance

1. Sign in to the AWS Management Console and open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. From the navigation pane on the left side of the console page, choose **Hardware**.
3. From the **Hardware appliance** list, select the activated hardware appliance on which you want to create your gateway, and then choose **Create Gateway**.
4. Follow the procedures described in [Creating Your Gateway](#) to set up, connect, and configure the type of Storage Gateway that you want to deploy.

When you finish creating your gateway in the Storage Gateway console, the Storage Gateway software automatically starts installing on the hardware appliance. If you use Dynamic Host Configuration Protocol (DHCP), it can take 5 to 10 minutes for a gateway to display as online in the console. To assign a static IP address to your installed gateway, see [Configuring an IP address for the gateway](#).

To assign a static IP address to your installed gateway, you next configure the gateway's network interfaces so your applications can use it.

Next step

[Configuring a gateway IP address on the hardware appliance](#)

Configuring a gateway IP address on the hardware appliance

Before you activated your hardware appliance, you assigned an IP address to its physical network interface. Now that you have activated the appliance and launched your Storage Gateway on it, you need to assign another IP address to the Storage Gateway virtual machine that runs on

the hardware appliance. To assign a static IP address to a gateway installed on your hardware appliance, configure the IP address from the local console for that gateway. Your applications (such as your NFS or SMB client) connect to this IP address. You can access the gateway local console from the hardware appliance console.

To configure an IP address on your appliance to work with applications

1. On the hardware console, choose **Open Service Console** and then press `Enter` to open the login page for the gateway local console.
2. The AWS Storage Gateway local console login page prompts you to login to change your network configuration and other settings.

For **localhost login**, type the account name and press `Enter`. Then, type the password and press `Enter`.

The default account is `admin` and the default password is `password`.

Note

We recommend changing the default password by entering the corresponding numeral for **Gateway Console** from the **AWS Appliance Activation - Configuration** main menu, then running the `passwd` command. For information about how to run the command, see [Running storage gateway commands in the local console for an on-premises gateway](#). You can also set the password from the Storage Gateway console. For more information, see [Setting the Local Console Password from the Storage Gateway Console](#).

3. The **AWS Appliance Activation - Configuration** page includes the following menu options:
 - HTTP/SOCKS Proxy Configuration
 - Network Configuration
 - Test Network Connectivity
 - View System Resource Check
 - System Time Management
 - License Information
 - Command Prompt
 - Get activation key

Note

Some options appear only for specific gateway types or host platforms.

Enter the corresponding numeral to navigate to the **Network Configuration** page.

4. Do one of the following to configure the gateway IP address:

- To use the IP address assigned by your Dynamic Host Configuration Protocol (DHCP) server, enter the corresponding numeral for **Configure DHCP**, and then enter valid DHCP configuration information on the following page.
- To assign a static IP address, enter the corresponding numeral for **Configure Static IP**, and then enter valid IP address and DNS information on the following page.

Note

The IP address you specify here must be on the same subnet as the IP address used during hardware appliance activation.

To exit the gateway local console

- Press the `Ctrl+] (close bracket)` keystroke. The hardware console appears.

Note

The keystroke preceding is the only way to exit the gateway local console.

After your hardware appliance has been activated and configured, your appliance appears in the console. Now you can continue the setup and configuration procedure for your gateway in the Storage Gateway console. For instructions, see .

Removing gateway software from your hardware appliance

If you no longer need a specific Storage Gateway that you have deployed on a hardware appliance, you can remove the gateway software from the hardware appliance. After you remove the gateway software, you can choose to deploy a new gateway in its place, or delete the hardware appliance itself from the Storage Gateway console. To remove gateway software from your hardware appliance, use the following procedure.

To remove a gateway from a hardware appliance

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **Hardware** from the navigation pane on the left side of the console page, and then choose the **Hardware appliance name** for the appliance from which you want to remove gateway software.
3. From the **Actions** drop down menu, choose **Remove gateway**.

The confirmation dialog box appears.

4. Verify that you want to remove the gateway software from the specified hardware appliance, and then type the word `remove` in the confirmation box.
5. Choose **Remove** to permanently remove the gateway software.

Note

After you remove the gateway software, you can't undo the action. For certain gateway types, you can lose data on deletion, particularly cached data. For more information on deleting a gateway, see [Deleting your gateway and removing associated resources](#).

Removing the gateway doesn't delete the hardware appliance from the console. The hardware appliance remains for future gateway deployments.

Deleting your Storage Gateway Hardware Appliance

If you no longer need an Storage Gateway Hardware Appliance that you have already activated, you can delete the appliance completely from your AWS account.

Note

To move your appliance to a different AWS account or AWS Region, you must first delete it using the following procedure, then open the gateway's support channel and contact AWS Support to perform a soft reset. For more information, see [Turning on AWS Support access to help troubleshoot your gateway hosted on-premises](#).

To delete your hardware appliance

1. If you have installed a gateway on the hardware appliance, you must first remove the gateway before you can delete the appliance. For instructions on how to remove a gateway from your hardware appliance, see [Removing gateway software from your hardware appliance](#).
2. On the Hardware page of the Storage Gateway console, choose the hardware appliance you want to delete.
3. For **Actions**, choose **Delete Appliance**. The confirmation dialog box appears.
4. Verify that you want to delete the specified hardware appliance, then type the word *delete* in the confirmation box and choose **Delete**.

When you delete the hardware appliance, all resources associated with the gateway that is installed on the appliance are deleted, but the data on the hardware appliance itself is not deleted.

Creating your gateway

The overview sections on this page provide a high-level synopsis of how the Storage Gateway creation process works. For step-by-step procedures to create a specific type of gateway using the Storage Gateway console, see the following topics:

- [Create and activate an Amazon S3 File Gateway](#)
- [Create and activate an Amazon FSx File Gateway](#)
- [Create and activate a Tape Gateway](#)
- [Create and activate a Volume Gateway](#)

Important

AWS Storage Gateway's FSx File Gateway will no longer be available to new customers after 10/28/24. To use the service, you must sign up prior to that date. Existing customers of FSx File Gateway can continue to use the service normally. For capabilities similar to FSx File Gateway, visit [this blog post](#).

Overview - Gateway Activation

Gateway activation involves setting up your gateway, connecting it to AWS, then reviewing your settings and activating it.

Set up gateway

To set up your Storage Gateway, you first choose the type of gateway you want to create and the host platform on which you will run the gateway virtual appliance. You then download the gateway virtual appliance template for the platform of your choice and deploy it in your on-premises environment. You can also deploy your Storage Gateway as a physical hardware appliance that you order from your preferred reseller, or as an Amazon EC2 instance in your AWS cloud environment. When you deploy the gateway appliance, you allocate local physical disk space on the virtualization host.

Connect to AWS

The next step is to connect your gateway to AWS. To do this, you first choose the type of service endpoint you want to use for communications between the gateway virtual appliance and AWS services in the cloud. This endpoint can be accessible from the public internet, or only from within your Amazon VPC, where you have full control over the network security configuration. You then specify the gateway's IP address or its activation key, which you can obtain by connecting to the local console on the gateway appliance.

Review and activate

At this point, you'll have an opportunity to review the gateway and connection options you chose, and make changes if necessary. When everything is set up the way you want you can activate the gateway. Before you can start using your activated gateway, you will need to configure some additional settings and create your storage resources.

Overview - Gateway Configuration

After you activate your Storage Gateway, you need to perform some additional configuration. In this step, you allocate the physical storage you provisioned on the gateway host platform to be used as either the cache or the upload buffer by the gateway appliance. You then configure settings to help monitor the health of your gateway using Amazon CloudWatch Logs and CloudWatch alarms, and add tags to help identify the gateway, if desired. Before you can start using your activated and configured gateway, you will need to create your storage resources.

Overview - Storage Resources

After you activate and configure your Storage Gateway, you need to create cloud storage resources for it to use. Depending on the type of gateway you created, you will use the Storage Gateway console to create Volumes, Tapes, or Amazon S3 or Amazon FSx files shares to associate with it. Each gateway type uses its respective resources to emulate the related type of network storage infrastructure, and transfers the data you write to it into the AWS cloud.

Creating a Volume Gateway

In this section, you can find instructions on how to download, deploy, and activate a Volume Gateway.

Topics

- [Set up a Volume Gateway](#)
- [Connect your Volume Gateway to AWS](#)
- [Review settings and activate your Volume Gateway](#)
- [Configure your Volume Gateway](#)

Set up a Volume Gateway

To set up a new Volume Gateway

1. Open the AWS Management Console at <https://console.aws.amazon.com/storagegateway/home/>, and choose the AWS Region where you want to create your gateway.
2. Choose **Create gateway** to open the **Set up gateway** page.
3. In the **Gateway settings** section, do the following:
 - a. For **Gateway name**, enter a name for your gateway. You can search for this name to find your gateway on list pages in the Storage Gateway console.
 - b. For **Gateway time zone**, choose the local time zone for the part of the world where you want to deploy your gateway.
4. In the **Gateway options** section, for **Gateway type**, choose **Volume Gateway**, then choose the volume type your gateway will use. You can choose from the following options:
 - **Cached volumes** - Stores your primary data in Amazon S3 and retains frequently accessed data locally in cache for faster access.
 - **Stored volumes** - Stores all of your data locally while also backing it up asynchronously to Amazon S3. Gateways using this volume type cannot be deployed on Amazon EC2.
5. In the **Platform options** section, do the following:
 - a. For **Host platform**, choose the platform on which you want to deploy your gateway, then follow the platform-specific instructions displayed on the Storage Gateway console page to set up your host platform. You can choose from the following options:
 - **VMware ESXi** - Download, deploy, and configure the gateway virtual machine using VMware ESXi.
 - **Microsoft Hyper-V** - Download, deploy, and configure the gateway virtual machine using Microsoft Hyper-V.

- **Linux KVM** - Download, deploy, and configure the gateway virtual machine using Linux KVM.
 - **Amazon EC2** - Configure and launch an Amazon EC2 instance to host your gateway. This option is not available for **Stored volume** gateways.
 - **Hardware appliance** - Order a dedicated physical hardware appliance from AWS to host your gateway.
- b. For **Confirm set up gateway**, select the check box to confirm that you performed the deployment steps for the host platform you chose. This step is not applicable for the **Hardware appliance** host platform.
6. Choose **Next** to proceed.

Now that your gateway is set up, you need to choose how you want it to connect and communicate with AWS. For instructions, see [Connect your Volume Gateway to AWS](#).

Connect your Volume Gateway to AWS

To connect a new Volume Gateway to AWS

1. Complete the procedure described in [Set up a Volume Gateway](#) if you have not done so already. When finished, choose **Next** to open the **Connect to AWS** page in the Storage Gateway console.
2. In the **Endpoint options** section, for **Service endpoint**, choose the type of endpoint your gateway will use to communicate with AWS. You can choose from the following options:
 - **Publicly accessible** - Your gateway communicates with AWS over the public internet. If you select this option, use the **FIPS enabled endpoint** check box to specify whether the connection should comply with Federal Information Processing Standards (FIPS).

Note

If you require FIPS 140-2 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS-compliant endpoint. For more information, see [Federal Information Processing Standard \(FIPS\) 140-2](#). The FIPS service endpoint is only available in some AWS Regions. For more information, see [Storage Gateway endpoints and quotas](#) in the *AWS General Reference*.

- **VPC hosted** - Your gateway communicates with AWS through a private connection with your VPC, allowing you to control your network settings. If you select this option, you must specify an existing VPC endpoint by choosing its VPC endpoint ID from the drop-down menu, or by providing its VPC endpoint DNS name or IP address.
3. In the **Gateway connection options** section, for **Connection options**, choose how to identify your gateway to AWS. You can choose from the following options:
 - **IP address** - Provide the IP address of your gateway in the corresponding field. This IP address must be public or accessible from within your current network, and you must be able to connect to it from your web browser.

You can obtain the gateway IP address by logging into the gateway's local console from your hypervisor client, or by copying it from your Amazon EC2 instance details page.

- **Activation key** - Provide the activation key for your gateway in the corresponding field. You can generate an activation key using the gateway's local console. Choose this option if your gateway's IP address is unavailable.
4. Choose **Next** to proceed.

Now that you have chosen how you want your gateway to connect to AWS, you need to activate the gateway. For instructions, see [Review settings and activate your Volume Gateway](#).

Review settings and activate your Volume Gateway

To activate a new Volume Gateway

1. Complete the procedures described in the following topics if you have not done so already:
 - [Set up a Volume Gateway](#)
 - [Connect your Volume Gateway to AWS](#)

When finished, choose **Next** to open the **Review and activate** page in the Storage Gateway console.

2. Review the initial gateway details for each section on the page.
3. If a section contains errors, choose **Edit** to return to the corresponding settings page and make changes.

Note

You cannot modify the gateway options or connection settings after your gateway is created.

4. Choose **Activate gateway** to proceed.

Now that you have activated your gateway, you need to perform first-time configuration to allocate local storage disks and configure logging. For instructions, see [Configure your Volume Gateway](#).

Configure your Volume Gateway

To perform first-time configuration on a new Volume Gateway

1. Complete the procedures described in the following topics if you have not done so already:
 - [Set up a Volume Gateway](#)
 - [Connect your Volume Gateway to AWS](#)
 - [Review settings and activate your Volume Gateway](#)

When finished, choose **Next** to open the **Configure gateway** page in the Storage Gateway console.

2. In the **Configure storage** section, use the drop-down menus to allocate at least one disk with at least **165 GiB** capacity for **CACHE STORAGE**, and at least one disk with at least **150 GiB** capacity for **UPLOAD BUFFER**. The local disks listed in this section correspond to the physical storage that you provisioned on your host platform.
3. In the **CloudWatch log group** section, choose how to set up Amazon CloudWatch Logs to monitor the health of your gateway. You can choose from the following options:
 - **Create a new log group** - Set up a new log group to monitor your gateway.
 - **Use an existing log group** - Choose an existing log group from the corresponding drop-down menu.
 - **Deactivate logging** - Do not use Amazon CloudWatch Logs to monitor your gateway.

Note

To receive Storage Gateway health logs, the following permissions must be present in your log group resource policy. Replace the *highlighted section* with the specific log group resourceArn information for your deployment.

```
"Sid": "AWSLogDeliveryWrite20150319",
  "Effect": "Allow",
  "Principal": {
    "Service": [
      "delivery.logs.amazonaws.com"
    ]
  },
  "Action": [
    "logs:CreateLogStream",
    "logs:PutLogEvents"
  ],
  "Resource": "arn:aws:logs:eu-west-1:1234567890:log-group:/foo/bar:log-stream:*"
```

The "Resource" element is required only if you want the permissions to apply explicitly to an individual log group.

4. In the **CloudWatch alarms** section, choose how to set up Amazon CloudWatch alarms to notify you when gateway metrics deviate from defined limits. You can choose from the following options:
 - **Create Storage Gateway's recommended alarms** – Create all recommended CloudWatch alarms automatically when the gateway is created. For more information about recommended alarms, see [Understanding CloudWatch alarms](#).

Note

This feature requires CloudWatch policy permissions, which are *not* automatically granted as part of the preconfigured Storage Gateway full access policy. Make sure your security policy grants the following permissions before you attempt to create recommended CloudWatch alarms:

- `cloudwatch:PutMetricAlarm` - create alarms

- `cloudwatch:DisableAlarmActions` - turn alarm actions off
 - `cloudwatch:EnableAlarmActions` - turn alarm actions on
 - `cloudwatch>DeleteAlarms` - delete alarms
- **Create a custom alarm** – Configure a new CloudWatch alarm to notify you about your gateway's metrics. Choose **Create alarm** to define metrics and specify alarm actions in the Amazon CloudWatch console. For instructions, see [Using Amazon CloudWatch alarms](#) in the *Amazon CloudWatch User Guide*.
 - **No alarm** – Don't receive CloudWatch notifications about your gateway's metrics.
5. (Optional) In the **Tags** section, choose **Add new tag**, then enter a case-sensitive key-value pair to help you search and filter for your gateway on list pages in the Storage Gateway console. Repeat this step to add as many tags as you need.
 6. Choose **Configure** to finish creating your gateway.

To check the status of your new gateway, search for it on the **Gateway overview** page of the Storage Gateway.

Now that you have created your gateway, you need to create a volume for it to use. For instructions, see [Creating a volume](#).

Creating a storage volume

Previously, you allocated local disks that you added to the VM cache storage and upload buffer. Now you create a storage volume to which your applications read and write data. The gateway maintains the volume's recently accessed data locally in cache storage, and asynchronously transferred data to Amazon S3. For stored volumes, you allocated local disks that you added to the VM upload buffer and your application's data.

Note

You can use AWS Key Management Service (AWS KMS) to encrypt data written to a cached volume that is stored in Amazon S3. Currently, you can do this by using the *AWS Storage Gateway API Reference*. For more information, see [CreateCachediSCSIVolume](#) or [create-cached-iscsi-volume](#).

To create a volume

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. On the Storage Gateway console, choose **Create volume**.
3. In the **Create volume** dialog box, choose a gateway for **Gateway**.
4. For the cached volumes, enter the capacity in **Capacity**.

For stored volumes, choose a **Disk ID** value from the list.

5. For **Volume content**, your choices depend on the type of gateway that you're creating the volume for.

For cached volumes, you have the following options:

- **Create a new empty volume.**
- **Create a volume based on an Amazon EBS snapshot.** If you choose this option, provide a value for **EBS snapshot ID**.

Note

Storage Gateway does not support creating cached volumes from snapshots of AWS Marketplace volumes.

- **Clone from last volume recovery point.** If you choose this option, choose a volume ID for **Source volume**. If there are no volumes in the Region, this option doesn't appear.

For stored volumes, you have the following options:

- **Create a new empty volume.**
 - **Create a volume based on a snapshot.** If you choose this option, provide a value for **EBS snapshot ID**.
 - **Preserve existing data on the disk**
6. Enter a name for **iSCSI target name**.

The target name can contain lowercase letters, numbers, periods (.), and hyphens (-).


This target name appears as the **iSCSI target node** name in the **Targets** tab of the **iSCSI Microsoft initiator** UI after discovery. For example, the name `target1` appears as

`iqn.1007-05.com.amazon:target1`. Make sure that the target name is globally unique within your storage area network (SAN).

7. Verify that the **Network interface** setting has IP address selected, or choose an IP address for **Network interface**. For **Network interface**, one IP address appears for each adapter that is configured for the gateway VM. If the gateway VM is configured for only one network adapter, no **Network interface** list appears because there is only one IP address.

Your iSCSI target will be available on the network adapter you choose.

If you have defined your gateway to use multiple network adapters, choose the IP address that your storage applications should use to access your volume. For information about configuring multiple network adapters, see [Configuring Your Gateway for Multiple NICs](#).

 **Note**

After you choose a network adapter, you can't change this setting.

8. (Optional) For **Tags**, enter a key and value to add tags to your volume. A tag is a case-sensitive key-value pair that helps you manage, filter, and search for your volumes.
9. Choose **Create volume**.

If you have previously created volumes in this Region, you can see them listed on the Storage Gateway console.

The **Configure CHAP Authentication** dialog box appears. At this point, you can configure Challenge-Handshake Authentication Protocol (CHAP) for your volume, or you can choose **Cancel** and configure CHAP later. For more information about CHAP setup, see [Configure CHAP authentication for your volumes](#).

If you don't want to set up CHAP, get started using your volume. For more information, see [Connecting your volumes to your client](#).

Configure CHAP authentication for your volumes

CHAP provides protection against playback attacks by requiring authentication to access your storage volume targets. In the **Configure CHAP Authentication** dialog box, you provide information to configure CHAP for your volumes.

To configure CHAP

1. Choose the volume for which you want to configure CHAP.
2. For **Actions**, choose **Configure CHAP authentication**.
3. For **Initiator Name**, enter the name of your initiator.
4. For **Initiator secret**, enter the secret phrase that you used to authenticate your iSCSI initiator.
5. For **Target secret**, enter the secret phrase used to authenticate your target for mutual CHAP.
6. Choose **Save** to save your entries.

For more information about setting up CHAP authentication, see [Configuring CHAP Authentication for Your iSCSI Targets](#).

Next step

[Connecting your volumes to your client](#)

Connecting your volumes to your client

You use the iSCSI initiator in your client to connect to your volumes. At the end of the following procedure, the volumes become available as local devices on your client.

Important

With Storage Gateway, you can connect multiple hosts to the same volume if the hosts coordinate access by using Windows Server Failover Clustering (WSFC). You can't connect multiple hosts to the same volume without using WSFC, for example by sharing a nonclustered NTFS/ext4 file system.

Topics

- [Connecting to a Microsoft Windows client](#)
- [Connecting to a Red Hat Enterprise Linux client](#)

Connecting to a Microsoft Windows client

The following procedure shows a summary of the steps that you follow to connect to a Windows client. For more information, see [Connecting iSCSI Initiators](#).

To connect to a Windows client

1. Start `iscsicpl.exe`.
2. In the **iSCSI Initiator Properties** dialog box, choose the **Discovery** tab, and then choose **Discovery Portal**.
3. In the **Discover Target Portal** dialog box, type the IP address of your iSCSI target for IP address or DNS name.
4. Connect the new target portal to the storage volume target on the gateway.
5. Choose the target, and then choose **Connect**.
6. In the **Targets** tab, make sure that the target status has the value **Connected**, indicating the target is connected, and then choose **OK**.

Connecting to a Red Hat Enterprise Linux client

The following procedure shows a summary of the steps that you follow to connect to a Red Hat Enterprise Linux (RHEL) client. For more information, see [Connecting iSCSI Initiators](#).

To connect a Linux client to iSCSI targets

1. Install the `iscsi-initiator-utils` RPM package.

You can use the following command to install the package.

```
sudo yum install iscsi-initiator-utils
```

2. Make sure that the iSCSI daemon is running.

For RHEL 5 or 6, use the following command.

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

For RHEL 7, use the following command.

```
sudo service iscsid status
```

3. Discover the volume or VTL device targets defined for a gateway. Use the following discovery command.

```
sudo /sbin/iscsiadm --mode discovery --type sendtargets --portal [GATEWAY_IP]:3260
```

The output of the discovery command should look like the following example output.

For Volume Gateways: `[GATEWAY_IP]:3260, 1 iqn.1997-05.com.amazon:myvolume`

For Tape Gateways: `iqn.1997-05.com.amazon:[GATEWAY_IP]-tapedrive-01`

4. Connect to a target.

Make sure to specify the correct `[GATEWAY_IP]` and IQN in the connect command.

Use the following command.

```
sudo /sbin/iscsiadm --mode node --targetname  
iqn.1997-05.com.amazon:[ISCSI_TARGET_NAME] --portal [GATEWAY_IP]:3260,1 --login
```

5. Verify that the volume is attached to the client machine (the initiator). To do so, use the following command.

```
ls -l /dev/disk/by-path
```

The output of the command should look like the following example output.

```
lrwxrwxrwx. 1 root root 9 Apr 16 19:31 ip-[GATEWAY_IP]:3260-iscsi-  
iqn.1997-05.com.amazon:myvolume-lun-0 -> ../../sda
```

We highly recommend that after you set up your initiator you customize your iSCSI settings as discussed in [Customizing Your Linux iSCSI Settings](#).

Initializing and formatting your volume

After you use the iSCSI initiator in your client to connect to your volumes, you initialize and format your volume.

Topics

- [Initializing and formatting your volume on Microsoft Windows](#)
- [Initializing and formatting your volume on Red Hat Enterprise Linux](#)

Initializing and formatting your volume on Microsoft Windows

Use the following procedure to initialize and format your volume on Windows.

To initialize and format your storage volume

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

Partition Style	Use in the Following Conditions
MBR (Master Boot Record)	<ul style="list-style-type: none"> • If your gateway is a stored volume and the storage volume is limited to 1 TiB in size. • If your gateway is a cached volume and the storage volume is less than 2 TiB in size.
GPT (GUID Partition Table)	If your gateway's storage volume is 2 TiB or greater in size.

3. Create a simple volume:
 - a. Bring the volume online to initialize it. All the available volumes are displayed in the disk management console.
 - b. Open the context (right-click) menu for the disk, and then choose **New Simple Volume**.

Important

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

- c. Specify the maximum disk size.
- d. Assign a drive letter or path to your volume, and format the volume by choosing **Perform a quick format**.

Important

We strongly recommend using **Perform a quick format** for cached volumes. Doing so results in less initialization I/O, smaller initial snapshot size, and the fastest time to a usable volume. It also avoids using cached volume space for the full format process.

Note

The time that it takes to format the volume depends on the size of the volume. The process might take several minutes to complete.

Initializing and formatting your volume on Red Hat Enterprise Linux

Use the following procedure to initialize and format your volume on Red Hat Enterprise Linux (RHEL).

To initialize and format your storage volume

1. Change directory to the `/dev` folder.
2. Run the `sudo cfdisk` command.
3. Identify your new volume by using the following command. To find new volumes, you can list the partition layout of your volumes.

```
$ lsblk
```

An "unrecognized volumes label" error for the new unpartitioned volume appears.

4. Initialize your new volume. When selecting the partition style, you should take into account the size and type of volume you are connecting to—cached or stored—as shown in the following table.

Partition Style	Use in the Following Conditions
MBR (Master Boot Record)	<ul style="list-style-type: none"> If your gateway is a stored volume and the storage volume is limited to 1 TiB in size. If your gateway is a cached volume and the storage volume is less than 2 TiB in size.
GPT (GUID Partition Table)	If your gateway's storage volume is 2 TiB or greater in size.

For an MBR partition, use the following command: `sudo parted /dev/your volume mklabel msdos`

For a GPT partition, use the following command: `sudo parted /dev/your volume mklabel gpt`

5. Create a partition by using the following command.

```
sudo parted -a opt /dev/your volume mkpart primary file system 0% 100%
```

6. Assign a drive letter to the partition and create a file system by using the following command.

```
sudo mkfs -L datapartition /dev/your volume
```

7. Mount the file system by using the following command.


```
sudo mount -o defaults /dev/your volume /mnt/your directory
```

Testing your gateway

You test your Volume Gateway setup by performing the following tasks:

1. Write data to the volume.
2. Take a snapshot.
3. Restore the snapshot to another volume.

You verify the setup for a gateway by taking a snapshot backup of your volume and storing the snapshot in AWS. You then restore the snapshot to a new volume. Your gateway copies the data from the specified snapshot in AWS to the new volume.

 **Note**

Restoring data from Amazon Elastic Block Store (Amazon EBS) volumes that are encrypted is not supported.

To create an Amazon EBS snapshot of a storage volume on Microsoft Windows

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

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

2. In the navigation pane of the Storage Gateway console, choose **Volumes**.
3. Choose the storage volume that you created for the gateway.

This gateway should have only one storage volume. Choose the volume displays its properties.

4. For **Actions**, choose **Create EBS snapshot** to create a snapshot of the volume.

Depending on the amount of data on the disk and the upload bandwidth, it might take a few seconds to complete the snapshot. Note the volume ID for the volume from which you create a snapshot. You use the ID to find the snapshot.

5. In the **Create EBS Snapshot** dialog box, provide a description for your snapshot.
6. (Optional) For **Tags**, enter a key and value to add tags to the snapshot. A tag is a case-sensitive key-value pair that helps you manage, filter, and search for your snapshots.
7. Choose **Create Snapshot**. Your snapshot is stored as an Amazon EBS snapshot. Note your snapshot ID. The number of snapshots created for your volume is displayed in the snapshot column.
8. In the **EBS snapshots** column, choose the link for the volume that you created the snapshot for to see your EBS snapshot on the Amazon EC2 console.

To restore a snapshot to another volume

See [Creating a storage volume](#).

Backing up your volumes

By using Storage Gateway, you can help protect your on-premises business applications that use Storage Gateway volumes for cloud-backed storage. You can back up your on-premises Storage Gateway volumes using the native snapshot scheduler in Storage Gateway or AWS Backup. In both cases, Storage Gateway volume backups are stored as Amazon EBS snapshots in Amazon Web Services.

Topics

- [Using Storage Gateway to back up your volumes](#)
- [Using AWS Backup to back up your volumes](#)

Using Storage Gateway to back up your volumes

You can use the Storage Gateway Management Console to back up your volumes by taking Amazon EBS snapshots and storing the snapshots in Amazon Web Services. You can either take a one-time snapshot or set up a snapshot schedule that is managed by Storage Gateway. You can later restore the snapshot to a new volume by using the Storage Gateway console. For information about how to back up and manage your backup from the Storage Gateway, see the following topics:

- [Testing your gateway](#)
- [Creating a recovery snapshot](#)
- [Cloning a cached volume from a recovery point](#)

Using AWS Backup to back up your volumes

AWS Backup is a centralized backup service that makes it easy and cost-effective for you to back up your application data across AWS services in both the Amazon Web Services Cloud and on-premises. Doing this helps you meet your business and regulatory backup compliance requirements. AWS Backup makes protecting your AWS storage volumes, databases, and file systems simple by providing a central place where you can do the following:

- Configure and audit the AWS resources that you want to back up.
- Automate backup scheduling.
- Set retention policies.
- Monitor all recent backup and restore activity.

Because Storage Gateway integrates with AWS Backup, it lets customers use AWS Backup to back up on-premises business applications that use Storage Gateway volumes for cloud-backed storage. AWS Backup supports backup and restore of both cached and stored volumes. For information about AWS Backup, see the AWS Backup documentation. For information about AWS Backup, see [What is AWS Backup?](#) in the *AWS Backup User Guide*.

You can manage Storage Gateway volumes' backup and recovery operations with AWS Backup and avoid the need to create custom scripts or manually manage point-in-time backups. With AWS Backup, you can also monitor your on-premises volume backups alongside your in-cloud AWS resources from a single AWS Backup dashboard. You can use AWS Backup to either create a one-time on-demand backup or define a backup plan that is managed in AWS Backup.

Storage Gateway volume backups taken from AWS Backup are stored in Amazon S3 as Amazon EBS snapshots. You can see the Storage Gateway volume backups from the AWS Backup console or the Amazon EBS console.

You can easily restore Storage Gateway volumes that are managed through AWS Backup to any on-premises gateway or in-cloud gateway. You can also restore such a volume to an Amazon EBS volume that you can use with Amazon EC2 instances.

Benefits of Using AWS Backup to Back Up Storage Gateway Volumes

The benefits of using AWS Backup to back up Storage Gateway volumes are that you can meet compliance requirements, avoid operational burden, and centralize backup management. AWS Backup allows you to do the following:

- Set customizable scheduled backup policies that meet your backup requirements.
- Set backup retention and expiration rules so you no longer need to develop custom scripts or manually manage the point-in-time backups of your volumes.
- Manage and monitor backups across multiple gateways, and other AWS resources from a central view.

To use AWS Backup to create backups of your volumes

Note

AWS Backup requires that you choose an AWS Identity and Access Management (IAM) role that AWS Backup consumes. You need to create this role because AWS Backup doesn't create it for you. You also need to create a trust relationship between AWS Backup and

this IAM role. For information about how to do this, see the *AWS Backup User Guide*. For information about how to do this, see [Creating a Backup Plan](#) in the *AWS Backup User Guide*.

1. Open the Storage Gateway console and choose **Volumes** from the navigation pane at left.
2. For **Actions**, choose **Create on-demand backup with AWS Backup** or **Create AWS backup plan**.

If you want to create an on-demand backup of the Storage Gateway volume, choose **Create on-demand backup with AWS Backup**. You are directed the AWS Backup console.

If you want to create a new AWS Backup plan, choose **Create AWS backup plan**. You are directed to the AWS Backup console.

On the AWS Backup console, you can create a backup plan, assign a Storage Gateway volume to the backup plan, and create a backup. You can also do ongoing backup management tasks.

Finding and restoring your volumes from AWS Backup

You can find and restore your backup Storage Gateway volumes from the AWS Backup console. For more information, see the *AWS Backup User Guide*. For more information, see [Recovery Points](#) in the *AWS Backup User Guide*.

To find and restore your volumes

1. Open the AWS Backup console and find the Storage Gateway volume backup that you want to restore. You can restore the Storage Gateway volume backup to an Amazon EBS volume or to a Storage Gateway volume. Choose the appropriate option for your restore requirements.
2. For **Restore type**, choose to restore a stored or cached Storage Gateway volume and provide the required information:
 - For a stored volume, provide the information for **Gateway name**, **Disk ID**, and **iSCSI target name**.
 - For a cached volume, provide the information for **Gateway name**, **Capacity**, and **iSCSI target name**.
3. Choose **Restore resource** to restore your volume.

Note

You can't use the Amazon EBS console to delete a snapshot that is created by AWS Backup.

Where do I go from here?

In the preceding sections, you created and provisioned a gateway and then connected your host to the gateway's storage volume. You added data to the gateway's iSCSI volume, took a snapshot of the volume, and restored it to a new volume, connected to the new volume, and verified that the data shows up on it.

After you finish the exercise, consider the following:

- If you plan on continuing to use your gateway, read about sizing the upload buffer more appropriately for real-world workloads. For more information, see [Sizing Your Volume Gateway's Storage for Real-World Workloads](#).

Other sections of this guide include information about how to do the following:

- To learn more about storage volumes and how to manage them, see [Managing Your Volume Gateway](#).
- If you don't plan on continuing to use your gateway, consider deleting the gateway to avoid incurring any charges. For more information, see [Cleaning up unnecessary resources](#).
- To troubleshoot gateway problems, see [Troubleshooting your gateway](#).
- To optimize your gateway, see [Optimizing gateway performance](#).
- To learn about Storage Gateway metrics and how you can monitor how your gateway performs, see [Monitoring Storage Gateway](#).
- To learn more about configuring your gateway's iSCSI targets to store data, see [Connecting to your volumes from a Windows client](#).

To learn about sizing your Volume Gateway's storage for real-world workloads and cleaning up resources you don't need, see the following sections.

Sizing Your Volume Gateway's Storage for Real-World Workloads

By this point, you have a simple, working gateway. However, the assumptions used to create this gateway are not appropriate for real-world workloads. If you want to use this gateway for real-world workloads, you need to do two things:

1. Size your upload buffer appropriately.
2. Set up monitoring for your upload buffer, if you haven't done so already.

Following, you can find how to do both of these tasks. If you activated a gateway for cached volumes, you also need to size your cache storage for real-world workloads.

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

- Use the formula shown in [Determining the size of upload buffer to allocate](#) for sizing the upload buffer. We strongly recommend that you allocate at least 150 GiB for the upload buffer. If the upload buffer formula yields a value less than 150 GiB, use 150 GiB as your allocated upload buffer.

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

To size your upload buffer for a stored setup

- Use the formula discussed in [Determining the size of upload buffer to allocate](#). We strongly recommend that you allocate at least 150 GiB for your upload buffer. If the upload buffer formula yields a value less than 150 GiB, use 150 GiB as your allocated upload buffer.

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

To monitor your upload buffer

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose the **Gateway** tab, choose the **Details** tab, and then find the **Upload Buffer Used** field to view your gateway's current upload buffer.
3. Set one or more alarms to notify you about upload buffer use.

We highly recommend that you create one or more upload buffer alarms in the Amazon CloudWatch console. For example, you can set an alarm for a level of use you want to be warned about and an alarm for a level of use that, if exceeded, is cause for action. The action might be adding more upload buffer space. For more information, see [To set an upper threshold alarm for a gateway's upload buffer](#).

Activating your gateway in a virtual private cloud

You can create a private connection between your on-premises gateway appliance and cloud-based storage infrastructure. You can use this connection to activate your gateway and allow it to transfer data to AWS storage services without communicating over the public internet. Using the Amazon VPC service, you can launch AWS resources, including private network interface endpoints, in a custom virtual private cloud (VPC). A VPC gives you control over network settings such as IP address range, subnets, route tables, and network gateways. For more information about VPCs, see [What is Amazon VPC?](#) in the *Amazon VPC User Guide*.

To activate your gateway in a VPC, use the Amazon VPC Console to create a VPC endpoint for Storage Gateway and get the VPC endpoint ID, then specify this VPC endpoint ID when you create and activate the gateway. For more information, see [Connect your Volume Gateway to AWS](#).

Note

You must activate your gateway in the same region where you create the VPC endpoint for Storage Gateway

Topics

- [Creating a VPC endpoint for Storage Gateway](#)

Creating a VPC endpoint for Storage Gateway

Follow these instructions to create a VPC endpoint. If you already have a VPC endpoint for Storage Gateway, you can use it to activate your gateway.

To create a VPC endpoint for Storage Gateway

1. Sign in to the AWS Management Console and open the Amazon VPC console at <https://console.aws.amazon.com/vpc/>.
2. In the navigation pane, choose **Endpoints**, and then choose **Create Endpoint**.
3. On the **Create Endpoint** page, choose **AWS Services** for **Service category**.
4. For **Service Name**, choose `com.amazonaws.region.storagegateway`. For example `com.amazonaws.us-east-2.storagegateway`.
5. For **VPC**, choose your VPC and note its Availability Zones and subnets.
6. Verify that **Enable Private DNS Name** is not selected.
7. For **Security group**, choose the security group that you want to use for your VPC. You can accept the default security group. Verify that all of the following TCP ports are allowed in your security group:
 - TCP 443
 - TCP 1026
 - TCP 1027
 - TCP 1028
 - TCP 1031
 - TCP 2222
8. Choose **Create endpoint**. The initial state of the endpoint is **pending**. When the endpoint is created, note the ID of the VPC endpoint that you just created.
9. When the endpoint is created, choose **Endpoints**, then choose the new VPC endpoint.
10. In **Details** tab of the selected storage gateway endpoint, under **DNS Names**, use the first DNS name that doesn't specify an Availability Zone. Your DNS name look similar to this: `vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com`

Now that you have a VPC endpoint, you can create your gateway. For more information, see [Creating a Gateway](#).

Managing Your Volume Gateway

Managing your gateway includes tasks such as configuring cache storage and upload buffer space, working with volumes, and doing general maintenance. If you haven't created a gateway, see [Getting started with AWS Storage Gateway](#).

Cached volumes are volumes in Amazon Simple Storage Service (Amazon S3) that are exposed as iSCSI targets on which you can store your application data. You can find information following about how to add and delete volumes for your cached setup. You can also learn how to add and remove Amazon Elastic Block Store (Amazon EBS) volumes in Amazon EC2 gateways.

Important

If a cached volume keeps your primary data in Amazon S3, you should avoid processes that read or write all data on the entire volume. For example, we don't recommend using virus-scanning software that scans the entire cached volume. Such a scan, whether done on demand or scheduled, causes all data stored in Amazon S3 to be downloaded locally for scanning, which results in high bandwidth usage. Instead of doing a full disk scan, you can use real-time virus scanning—that is, scanning data as it is read from or written to the cached volume.

Resizing a volume is not supported. To change the size of a volume, create a snapshot of the volume, and then create a new cached volume from the snapshot. The new volume can be bigger than the volume from which the snapshot was created. For steps describing how to remove a volume, see [To delete a volume](#). For steps describing how to add a volume and preserve existing data, see [Deleting storage volumes](#).

All cached volume data and snapshot data is stored in Amazon S3 and is encrypted at rest using server-side encryption (SSE). However, you cannot access this data by using the Amazon S3 API or other tools such as the Amazon S3 Management Console.

Following, you can find information about how to manage your Volume Gateway resources.

Topics

- [Editing Basic Gateway Information](#) - Learn how to use the Storage Gateway console to edit basic information for an existing gateway, including the gateway name, time zone, and CloudWatch log group.

- [Adding and expanding volumes](#) - Learn how to add more volumes to your gateway, or expand the size of existing volumes as your application needs grow.
- [Cloning a cached volume from a recovery point](#) - Learn how to create a new volume from an existing volume's recovery point, which is a saved point in time when all of the data on the volume is consistent.
- [Viewing volume usage](#) - Learn how to view the amount of data stored on a volume by using the Storage Gateway console.
- [Deleting storage volumes](#) - Learn how to delete a volume if your application needs change, such as if you migrate an application to use a larger storage volume.
- [Moving Your Volumes to a Different Gateway](#) - Learn how to detach and reattach volumes, which is useful if you need to move your volumes to a different Volume Gateway as your performance needs change.
- [Creating a recovery snapshot](#) - Learn how to create a recovery snapshot from a volume recovery point for a gateway, and where to find that snapshot in the Storage Gateway console after you create it.
- [Editing a snapshot schedule](#) - Learn how to customize a snapshot schedule by changing either the time the snapshot occurs each day or the frequency that snapshots are taken.
- [Deleting snapshots of your storage volumes](#) - Learn how to delete unnecessary snapshots when you no longer need them.
- [Understanding Volume Statuses and Transitions](#) - Learn about the various volume status values that Storage Gateway reports to help determine whether a volume is functioning normally, or if there is a problem that might require action on your part.
- [Moving your data to a new gateway](#) - Learn how to move data between gateways as your data and performance needs grow, or if you receive an AWS notification to migrate your gateway.


Editing Basic Gateway Information

You can use the Storage Gateway console to edit basic information for an existing gateway, including the gateway name, time zone, and CloudWatch log group.

To edit basic information for an existing gateway

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **Gateways**, then choose the gateway for which you want to edit basic information.

3. From the **Actions** dropdown menu, choose **Edit gateway information**.
4. For **Gateway name**, enter a name for your gateway. You can search for this name to find your gateway on the list pages in the Storage Gateway console.

 **Note**

Gateway names must be between 2 and 255 characters, and cannot include a slash (\ or /).

Changing a gateway's name will disconnect any CloudWatch alarms set up to monitor the gateway. To reconnect the alarms, update the **GatewayName** for each alarm in the CloudWatch console.

5. For **Gateway time zone**, choose the local time zone for the part of the world where you want to deploy your gateway.
6. For **Choose how to set up log group**, choose how to set up Amazon CloudWatch Logs to monitor the health of your gateway. You can choose from the following options:
 - **Create a new log group** – Set up a new log group to monitor your gateway.
 - **Use an existing log group** – Choose an existing log group from the corresponding dropdown list.
 - **Deactivate logging** – Do not use Amazon CloudWatch Logs to monitor your gateway.
7. When you finish modifying the settings you want to change, choose **Save changes**.

Adding and expanding volumes

As your application needs grow, you might need to add more volumes to your gateway, or expand the size of existing volumes. When you add or expand volumes, you must consider the size of the cache storage and upload buffer you allocated to the gateway. The gateway must have sufficient buffer and cache space for new volumes. For more information, see [Determining the size of upload buffer to allocate](#).

You can add volumes using the Storage Gateway console or Storage Gateway API. For instructions on how to add a volume using the Storage Gateway console, see [Creating a storage volume](#). For information about using the Storage Gateway API to add volumes, see [CreateCachediSCSIVolume](#).

You can expand the size of existing volumes using either of the following methods:

- Create a snapshot of the volume you want to expand and then use the snapshot to create a new volume of a larger size. For information about how to create a snapshot, see [Creating a recovery snapshot](#). For information about how to use a snapshot to create a new volume, see [Creating a storage volume](#).
- Use the cached volume you want to expand to clone a new volume of a larger size. For information about how to clone a volume, see [Cloning a cached volume from a recovery point](#). For information about how to create a volume, see [Creating a storage volume](#).

Cloning a cached volume from a recovery point

You can create a new volume from any existing cached volume in the same AWS Region. The new volume is created from the most recent recovery point of the selected volume. A *volume recovery point* is a point in time at which all data of the volume is consistent. To clone a volume, you choose the **Clone from last recovery point** option in the **Create volume** dialog box, then select the volume to use as the source.

Cloning from an existing volume is faster and more cost-effective than creating an Amazon EBS snapshot. Cloning does a byte-to-byte copy of your data from the source volume to the new volume, using the most recent recovery point from the source volume. Storage Gateway automatically creates recovery points for your cached volumes. To see when the last recovery point was created, check the `TimeSinceLastRecoveryPoint` metric in Amazon CloudWatch.

The cloned volume is independent of the source volume. That is, changes made to either volume after cloning have no effect on the other. For example, if you delete the source volume, it has no effect on the cloned volume. You can clone a source volume while initiators are connected and it is in active use. Doing so doesn't affect the performance of the source volume. For information about how to clone a volume, see [Creating a storage volume](#).

You can also use the cloning process in recovery scenarios. For more information, see [Your Cached Gateway is Unreachable And You Want to Recover Your Data](#).

The following procedure shows you how to clone a volume from a volume recovery point and use that volume.

To clone and use a volume from an unreachable gateway

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.

2. On the Storage Gateway console, choose **Create volume**.
3. In the **Create volume** dialog box, choose a gateway for **Gateway**.
4. For **Capacity**, type the capacity for your volume. The capacity must be at least the same size as the source volume.
5. Choose **Clone from last recovery point** and select a volume ID for **Source volume**. The source volume can be any cached volume in the selected AWS Region.
6. Type a name for **iSCSI target name**.

The target name can contain lowercase letters, numbers, periods (.), and hyphens (-).

This target name appears as the **iSCSI target node** name in the **Targets** tab of the **iSCSI Microsoft initiator** UI after discovery. For example, the name `target1` appears as `iqn.1007-05.com.amazon:target1`. Ensure that the target name is globally unique within your storage area network (SAN).

7. Verify that the **Network interface** setting is the IP address of your gateway, or choose an IP address for **Network interface**.

If you have defined your gateway to use multiple network adapters, choose the IP address that your storage applications use to access the volume. Each network adapter defined for a gateway represents one IP address that you can choose.

If the gateway VM is configured for more than one network adapter, the **Create volume** dialog box displays a list for **Network interface**. In this list, one IP address appears for each adapter configured for the gateway VM. If the gateway VM is configured for only one network adapter, no list appears because there's only one IP address.

8. Choose **Create volume**. The **Configure CHAP Authentication** dialog box appears. You can configure CHAP later. For information, see [Configuring CHAP Authentication for Your iSCSI Targets](#).

The next step is to connect your volume to your client. For more information, see [Connecting your volumes to your client](#).

Viewing volume usage

When you write data to a volume, you can view the amount of data stored on the volume in the Storage Gateway Management Console. The **Details** tab for each volume shows the volume usage information.

To view amount of data written to a volume

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose **Volumes** and then choose the volume you are interested in.
3. Choose the **Details** tab.

The following fields provide information about the volume:

- **Size:** The total capacity of the selected volume.
- **Used:** The size of data stored on the volume.

Note

These values are not available for volumes created before May 13, 2015, until you store data on the volume.

Deleting storage volumes

You might need to delete a volume as your application needs change—for example, if you migrate your application to use a larger storage volume. Before you delete a volume, make sure that there are no applications currently writing to the volume. Also, make sure that there are no snapshots in progress for the volume. If a snapshot schedule is defined for the volume, you can check it on the **Snapshot Schedules** tab of the Storage Gateway console. For more information, see [Editing a snapshot schedule](#).

You can delete volumes using the Storage Gateway console or the Storage Gateway API. For information on using the Storage Gateway API to remove volumes, see [Delete Volume](#). The following procedure demonstrates using the console.

Before you delete a volume, back up your data or take a snapshot of your critical data. For stored volumes, your local disks aren't erased. After you delete a volume, you can't get it back.

To delete a volume

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **Volumes**, then select one or more volumes to delete.

3. For **Actions** choose **Delete volume**. The confirmation dialog box appears.
4. Verify that you want to delete the specified volumes, then type the word *delete* in the confirmation box and choose **Delete**.

Moving Your Volumes to a Different Gateway

As your data and performance needs grow, you might want to move your volumes to a different Volume Gateway. To do so, you can detach and attach a volume by using the Storage Gateway console or API.

By detaching and attaching a volume, you can do the following:

- Move your volumes to better host platforms or newer Amazon EC2 instances.
- Refresh the underlying hardware for your server.
- Move your volumes between hypervisor types.

When you detach a volume, your gateway uploads and stores the volume data and metadata to the Storage Gateway service in AWS. You can easily attach a detached volume to a gateway on any supported host platform later.

Note

A detached volume is billed at the standard volume storage rate until you delete it. For information about how to reduce your bill, see [Reducing the amount of billed storage on a volume](#).

Note

There are some limitations for attaching and detaching volumes:

- Detaching a volume can take a long time. When you detach a volume, the gateway uploads all the data on the volume to AWS before the volume is detached. The time it takes for the upload to complete depends on how much data needs to be uploaded and your network connectivity into AWS.
- If you detach a cached volume, you can't reattach it as a stored volume.

- If you detach a stored volume, you can't reattach it as a cached volume.
- A detached volume can't be used until it is attached to a gateway.
- When you attach a stored volume, it needs to fully restore before you can attach it to a gateway.
- When you start attaching or detaching a volume, you need to wait till the operation completed before you use the volume.
- Currently, forcibly deleting a volume is only supported in the API.
- If you delete a gateway while your volume is detaching from that gateway, it results in data loss. Wait until the volume detach operation is complete before you delete the gateway.
- If a stored gateway is in restoring state, you can't detach a volume from it.

The following steps show you how to detach and attach a volume using the Storage Gateway console. For more information about doing this using the API, see [DetachVolume](#) or [AttachVolume](#) in the *AWS Storage Gateway API Reference*.

To detach a volume from a gateway

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **Volumes**, then select one or more volumes to detach.
3. For **Actions**, choose **Detach volume**. The confirmation dialog box appears.
4. Verify that you want to detach the specified volumes, then type the word *detach* in the confirmation box and choose **Detach**.

Note

If a volume that you detach has a lot of data on it, it transitions from **Attached** to **Detaching** status until it finishes uploading all the data. Then the status changes to **Detached**. For small amounts of data, you might not see the **Detaching** status. If the volume doesn't have data on it, the status changes from **Attached** to **Detached**.

You can now attach the volume to a different gateway.

To attach a volume to a gateway

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. On the navigation pane, choose **Volumes**. The status of each volume that is detached shows as **Detached**.
3. From the list of detached volumes, choose the volume that you want to attach. You can attach only one volume at a time.
4. For **Actions**, choose **Attach volume**.
5. In the **Attach Volume** dialog box, choose the gateway that you want to attach the volume to, and then enter the iSCSI target that you want to connect the volume to.

If you are attaching a stored volume, enter its disk identifier for **Disk ID**.

6. Choose **Attach volume**. If a volume that you attach has a lot of data on it, it transitions from **Detached** to **Attached** if the AttachVolume operation succeeds.
7. In the Configure CHAP authentication wizard that appears, enter the **Initiator name**, **Initiator secret**, and **Target secret**, and then choose **Save**. For more information about working with Challenge-Handshake Authentication Protocol (CHAP) authentication, see [Configuring CHAP Authentication for Your iSCSI Targets](#).

Creating a recovery snapshot

The following procedure shows you how to create a recovery snapshot from a volume recovery point for a gateway, and where to find that snapshot in the Storage Gateway console after you create it. You can take recovery snapshots on a one time, ad hoc basis or you can set up a snapshot schedule to take recurring snapshots of the volume at regular intervals that you specify.

To create and use a recovery snapshot of a volume from an existing gateway

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane on the left side of the console page, choose **Gateways**.
3. Choose the gateway for which you want to create a snapshot, and then choose the **Details** tab.

The **Details** tab displays a recovery snapshot message for the selected gateway.

4. Choose **Create recovery snapshot** to open the **Create recovery snapshot** dialog box.

5. From the list of volumes that appears, choose the volume that you want to recover, and then choose **Create snapshots**.

Storage Gateway initiates the snapshot process for the specified volume. When the snapshot process is complete, you can find the snapshot listed in the **Snapshots** column when viewing the volume on the **Volumes** page of the Storage Gateway console.

Editing a snapshot schedule

For stored volumes, AWS Storage Gateway creates a default snapshot schedule of once a day.

Note

You can't remove the default snapshot schedule. Stored volumes require at least one snapshot schedule. However, you can change a snapshot schedule by specifying either the time the snapshot occurs each day or the frequency (every 1, 2, 4, 8, 12, or 24 hours), or both.

For cached volumes, AWS Storage Gateway doesn't create a default snapshot schedule. No default schedule is created because your data is stored in Amazon S3, so you don't need snapshots or a snapshot schedule for disaster recovery purposes. However, you can set up a snapshot schedule at any time if you need to. Creating snapshot for your cached volume provides an additional way to recover your data if necessary.

By using the following steps, you can edit the snapshot schedule for a volume.

To edit the snapshot schedule for a volume

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose **Volumes**, and then choose the volume the snapshot was created from.
3. For **Actions**, choose **Edit snapshot schedule**.
4. In the **Edit snapshot schedule** dialog box, modify the schedule, and then choose **Save**.

Deleting snapshots of your storage volumes

You can delete a snapshot of your storage volume. For example, you might want to do this if you have taken many snapshots of a storage volume over time and you don't need the older snapshots. Because snapshots are incremental backups, if you delete a snapshot, only the data that is not needed in other snapshots is deleted.

Topics

- [Deleting Snapshots by Using the AWS SDK for Java](#)
- [Deleting Snapshots by Using the AWS SDK for .NET](#)
- [Deleting Snapshots by Using the AWS Tools for Windows PowerShell](#)

On the Amazon EBS console, you can delete snapshots one at a time. For information about how to delete snapshots using the Amazon EBS console, see [Deleting an Amazon EBS Snapshot](#) in the *Amazon EC2 User Guide*.

To delete multiple snapshots at a time, you can use one of the AWS SDKs that supports Storage Gateway operations. For examples, see [Deleting Snapshots by Using the AWS SDK for Java](#), [Deleting Snapshots by Using the AWS SDK for .NET](#), and [Deleting Snapshots by Using the AWS Tools for Windows PowerShell](#).

Deleting Snapshots by Using the AWS SDK for Java

To delete many snapshots associated with a volume, you can use a programmatic approach. The example following demonstrates how to delete snapshots using the AWS SDK for Java. To use the example code, you should be familiar with running a Java console application. For more information, see [Getting Started](#) in the *AWS SDK for Java Developer Guide*. If you need to just delete a few snapshots, use the console as described in [Deleting snapshots of your storage volumes](#).

Example : Deleting Snapshots by Using the AWS SDK for Java

The following Java code example lists the snapshots for each volume of a gateway and whether the snapshot start time is before or after a specified date. It uses the AWS SDK for Java API for Storage Gateway and Amazon EC2. The Amazon EC2 API includes operations for working with snapshots.

Update the code to provide the service endpoint, your gateway Amazon Resource Name (ARN), and the number of days back you want to save snapshots. Snapshots taken before this cutoff are deleted. You also need to specify the Boolean value `viewOnly`, which indicates whether you want to view the snapshots to be deleted or to actually perform the snapshot deletions. Run the code first with just the view option (that is, with `viewOnly` set to `true`) to see what the code deletes. For a list of AWS service endpoints you can use with Storage Gateway, see [AWS Storage Gateway Endpoints and Quotas](#) in the *AWS General Reference*.

```
import java.io.IOException;
import java.util.ArrayList;
import java.util.Calendar;
import java.util.Collection;
import java.util.Date;
import java.util.GregorianCalendar;
import java.util.List;

import com.amazonaws.auth.PropertiesCredentials;
import com.amazonaws.services.ec2.AmazonEC2Client;
import com.amazonaws.services.ec2.model.DeleteSnapshotRequest;
import com.amazonaws.services.ec2.model.DescribeSnapshotsRequest;
import com.amazonaws.services.ec2.model.DescribeSnapshotsResult;
import com.amazonaws.services.ec2.model.Filter;
import com.amazonaws.services.ec2.model.Snapshot;
import com.amazonaws.services.storagegateway.AWSStorageGatewayClient;
import com.amazonaws.services.storagegateway.model.ListVolumesRequest;
import com.amazonaws.services.storagegateway.model.ListVolumesResult;
import com.amazonaws.services.storagegateway.model.VolumeInfo;

public class ListDeleteVolumeSnapshotsExample {

    public static AWSStorageGatewayClient sgClient;
    public static AmazonEC2Client ec2Client;
    static String serviceURLSG = "https://storagegateway.us-east-1.amazonaws.com";
    static String serviceURLEC2 = "https://ec2.us-east-1.amazonaws.com";

    // The gatewayARN
    public static String gatewayARN = "**** provide gateway ARN ****";

    // The number of days back you want to save snapshots. Snapshots before this cutoff
    are deleted
    // if viewOnly = false.
    public static int daysBack = 10;
```



```
// true = show what will be deleted; false = actually delete snapshots that meet
the daysBack criteria
public static boolean viewOnly = true;

public static void main(String[] args) throws IOException {

    // Create a Storage Gateway and amazon ec2 client
    sgClient = new AWSStorageGatewayClient(new PropertiesCredentials(
ListDeleteVolumeSnapshotsExample.class.getResourceAsStream("AwsCredentials.properties")));

    sgClient.setEndpoint(serviceURLSG);

    ec2Client = new AmazonEC2Client(new PropertiesCredentials(
ListDeleteVolumeSnapshotsExample.class.getResourceAsStream("AwsCredentials.properties")));
    ec2Client.setEndpoint(serviceURLEC2);

    List<VolumeInfo> volumes = ListVolumesForGateway();
    DeleteSnapshotsForVolumes(volumes, daysBack);

}
public static List<VolumeInfo> ListVolumesForGateway()
{
    List<VolumeInfo> volumes = new ArrayList<VolumeInfo>();

    String marker = null;
    do {
        ListVolumesRequest request = new
ListVolumesRequest().withGatewayARN(gatewayARN);
        ListVolumesResult result = sgClient.listVolumes(request);
        marker = result.getMarker();

        for (VolumeInfo vi : result.getVolumeInfos())
        {
            volumes.add(vi);
            System.out.println(OutputVolumeInfo(vi));
        }
    } while (marker != null);

    return volumes;
}
private static void DeleteSnapshotsForVolumes(List<VolumeInfo> volumes,
```

```
int daysBack2) {

// Find snapshots and delete for each volume
for (VolumeInfo vi : volumes) {

    String volumeARN = vi.getVolumeARN();
    String volumeId =
volumeARN.substring(volumeARN.lastIndexOf("/")+1).toLowerCase();
    Collection<Filter> filters = new ArrayList<Filter>();
    Filter filter = new Filter().withName("volume-id").withValues(volumeId);
    filters.add(filter);

    DescribeSnapshotsRequest describeSnapshotsRequest =
        new DescribeSnapshotsRequest().withFilters(filters);
    DescribeSnapshotsResult describeSnapshotsResult =
        ec2Client.describeSnapshots(describeSnapshotsRequest);

    List<Snapshot> snapshots = describeSnapshotsResult.getSnapshots();
    System.out.println("volume-id = " + volumeId);
    for (Snapshot s : snapshots){
        StringBuilder sb = new StringBuilder();
        boolean meetsCriteria = !CompareDates(daysBack, s.getStartTime());
        sb.append(s.getSnapshotId() + ", " + s.getStartTime().toString());

        sb.append(", meets criteria for delete? " + meetsCriteria);
        sb.append(", deleted? ");
        if (!viewOnly & meetsCriteria) {
            sb.append("yes");
            DeleteSnapshotRequest deleteSnapshotRequest =
                new DeleteSnapshotRequest().withSnapshotId(s.getSnapshotId());
            ec2Client.deleteSnapshot(deleteSnapshotRequest);
        }
        else {
            sb.append("no");
        }
        System.out.println(sb.toString());
    }
}

private static String OutputVolumeInfo(VolumeInfo vi) {

    String volumeInfo = String.format(
        "Volume Info:\n" +
```

```
        " ARN: %s\n" +
        " Type: %s\n",
        vi.getVolumeARN(),
        vi.getVolumeType());
    return volumeInfo;
}

// Returns the date in two formats as a list
public static boolean CompareDates(int daysBack, Date snapshotDate) {
    Date today = new Date();
    Calendar cal = new GregorianCalendar();
    cal.setTime(today);
    cal.add(Calendar.DAY_OF_MONTH, -daysBack);
    Date cutoffDate = cal.getTime();
    return (snapshotDate.compareTo(cutoffDate) > 0) ? true : false;
}
}
```

Deleting Snapshots by Using the AWS SDK for .NET

To delete many snapshots associated with a volume, you can use a programmatic approach. The following example demonstrates how to delete snapshots using the AWS SDK for .NET version 2 and 3. To use the example code, you should be familiar with running a .NET console application. For more information, see [Getting Started](#) in the *AWS SDK for .NET Developer Guide*. If you need to just delete a few snapshots, use the console as described in [Deleting snapshots of your storage volumes](#).

Example : Deleting Snapshots by Using the AWS SDK for .NET

In the following C# code example, an AWS Identity and Access Management user can list the snapshots for each volume of a gateway. The user can then determine whether the snapshot start time is before or after a specified date (retention period) and delete snapshots that have passed the retention period. The example uses the AWS SDK for .NET API for Storage Gateway and Amazon EC2. The Amazon EC2 API includes operations for working with snapshots.

The following code example uses the AWS SDK for .NET version 2 and 3. You can migrate older versions of .NET to the newer version. For more information, see [Migrating Your Code to the Latest Version of the AWS SDK for .NET](#).

Update the code to provide the service endpoint, your gateway Amazon Resource Name (ARN), and the number of days back you want to save snapshots. Snapshots taken before this cutoff are deleted. You also need to specify the Boolean value `viewOnly`, which indicates whether you want to view the snapshots to be deleted or to actually perform the snapshot deletions. Run the code first with just the view option (that is, with `viewOnly` set to `true`) to see what the code deletes. For a list of AWS service endpoints you can use with Storage Gateway, see [AWS Storage Gateway Endpoints and Quotas](#) in the *AWS General Reference*.

First, you create a user and attach the minimum IAM policy to the user. Then you schedule automated snapshots for your gateway.

The following code creates the minimum policy that allows a user to delete snapshots. In this example, the policy is named **sgw-delete-snapshot**.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "StmtEC2Snapshots",
      "Effect": "Allow",
      "Action": [
        "ec2:DeleteSnapshot",
        "ec2:DescribeSnapshots"
      ],
      "Resource": [
        "*"
      ]
    },
    {
      "Sid": "StmtSgwListVolumes",
      "Effect": "Allow",
      "Action": [
        "storagegateway:ListVolumes"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

The following C# code finds all snapshots in the specified gateway that match the volumes and the specified cut-off period and then deletes them.

```
using System;
using System.Collections.Generic;
using System.Text;
using Amazon.EC2;
using Amazon.EC2.Model;
using Amazon.StorageGateway.Model;
using Amazon.StorageGateway;

namespace DeleteStorageGatewaySnapshotNS
{
    class Program
    {
        /*
         * Replace the variables below to match your environment.
         */

        /* IAM AccessKey */
        static String AwsAccessKey = "AKIA.....";

        /* IAM SecretKey */
        static String AwsSecretKey = "*****";

        /* Account number, 12 digits, no hyphen */
        static String OwnerID = "123456789012";

        /* Your Gateway ARN. Use a Storage Gateway ID, sgw-XXXXXXX* */
        static String GatewayARN = "arn:aws:storagegateway:ap-
southeast-2:123456789012:gateway/sgw-XXXXXXX";

        /* Snapshot status: "completed", "pending", "error" */

        static String SnapshotStatus = "completed";

        /* Region where your gateway is activated */
        static String AwsRegion = "ap-southeast-2";

        /* Minimum age of snapshots before they are deleted (retention policy) */
        static int daysBack = 30;

        /*
```

```
    * Do not modify the four lines below.
    */
static AmazonEC2Config ec2Config;
static AmazonEC2Client ec2Client;
static AmazonStorageGatewayClient sgClient;
static AmazonStorageGatewayConfig sgConfig;

static void Main(string[] args)
{
    // Create an EC2 client.
    ec2Config = new AmazonEC2Config();
    ec2Config.ServiceURL = "https://ec2." + AwsRegion + ".amazonaws.com";
    ec2Client = new AmazonEC2Client(AwsAccessKey, AwsSecretKey, ec2Config);

    // Create a Storage Gateway client.
    sgConfig = new AmazonStorageGatewayConfig();
    sgConfig.ServiceURL = "https://storagegateway." + AwsRegion +
".amazonaws.com";
    sgClient = new AmazonStorageGatewayClient(AwsAccessKey, AwsSecretKey,
sgConfig);

    List<VolumeInfo> StorageGatewayVolumes = ListVolumesForGateway();
    List<Snapshot> StorageGatewaySnapshots =
ListSnapshotsForVolumes(StorageGatewayVolumes,
                        daysBack);
    DeleteSnapshots(StorageGatewaySnapshots);
}

/*
 * List all volumes for your gateway
 * returns: A list of VolumeInfos, or null.
 */
private static List<VolumeInfo> ListVolumesForGateway()
{
    ListVolumesResponse response = new ListVolumesResponse();
    try
    {
        ListVolumesRequest request = new ListVolumesRequest();
        request.GatewayARN = GatewayARN;
        response = sgClient.ListVolumes(request);

        foreach (VolumeInfo vi in response.VolumeInfos)
        {
            Console.WriteLine(OutputVolumeInfo(vi));
        }
    }
}
```

```
        }
    }
    catch (AmazonStorageGatewayException ex)
    {
        Console.WriteLine(ex.Message);
    }
    return response.VolumeInfos;
}

/*
 * Gets the list of snapshots that match the requested volumes
 * and cutoff period.
 */
private static List<Snapshot> ListSnapshotsForVolumes(List<VolumeInfo> volumes,
int snapshotAge)
{
    List<Snapshot> SelectedSnapshots = new List<Snapshot>();
    try
    {
        foreach (VolumeInfo vi in volumes)
        {
            String volumeARN = vi.VolumeARN;
            String volumeID = volumeARN.Substring(volumeARN.LastIndexOf("/") +
1).ToLower();

            DescribeSnapshotsRequest describeSnapshotsRequest = new
DescribeSnapshotsRequest();

            Filter ownerFilter = new Filter();
            List<String> ownerValues = new List<String>();
            ownerValues.Add(OwnerID);
            ownerFilter.Name = "owner-id";
            ownerFilter.Values = ownerValues;
            describeSnapshotsRequest.Filters.Add(ownerFilter);

            Filter statusFilter = new Filter();
            List<String> statusValues = new List<String>();
            statusValues.Add(SnapshotStatus);
            statusFilter.Name = "status";
            statusFilter.Values = statusValues;
            describeSnapshotsRequest.Filters.Add(statusFilter);

            Filter volumeFilter = new Filter();
            List<String> volumeValues = new List<String>();
```

```
        volumeValues.Add(volumeID);
        volumeFilter.Name = "volume-id";
        volumeFilter.Values = volumeValues;
        describeSnapshotsRequest.Filters.Add(volumeFilter);

        DescribeSnapshotsResponse describeSnapshotsResponse =
            ec2Client.DescribeSnapshots(describeSnapshotsRequest);

        List<Snapshot> snapshots = describeSnapshotsResponse.Snapshots;
        Console.WriteLine("volume-id = " + volumeID);
        foreach (Snapshot s in snapshots)
        {
            if (IsSnapshotPastRetentionPeriod(snapshotAge, s.StartTime))
            {
                Console.WriteLine(s.SnapshotId + ", " + s.VolumeId + ",
                    " + s.StartTime + ", " + s.Description);
                SelectedSnapshots.Add(s);
            }
        }
    }
}
catch (AmazonEC2Exception ex)
{
    Console.WriteLine(ex.Message);
}
return SelectedSnapshots;
}

/*
 * Deletes a list of snapshots.
 */
private static void DeleteSnapshots(List<Snapshot> snapshots)
{
    try
    {
        foreach (Snapshot s in snapshots)
        {

            DeleteSnapshotRequest deleteSnapshotRequest = new
DeleteSnapshotRequest(s.SnapshotId);
            DeleteSnapshotResponse response =
ec2Client.DeleteSnapshot(deleteSnapshotRequest);
            Console.WriteLine("Volume: " +
                s.VolumeId +
```



```
        " => Snapshot: " +
        s.SnapshotId +
        " Response: "
        + response.HttpStatusCode.ToString());
    }
}
catch (AmazonEC2Exception ex)
{
    Console.WriteLine(ex.Message);
}
}

/*
 * Checks if the snapshot creation date is past the retention period.
 */
private static Boolean IsSnapshotPastRetentionPeriod(int daysBack, DateTime
snapshotDate)
{
    DateTime cutoffDate = DateTime.Now.Add(new TimeSpan(-daysBack, 0, 0, 0));
    return (DateTime.Compare(snapshotDate, cutoffDate) < 0) ? true : false;
}

/*
 * Displays information related to a volume.
 */
private static String OutputVolumeInfo(VolumeInfo vi)
{
    String volumeInfo = String.Format(
        "Volume Info:\n" +
        "  ARN: {0}\n" +
        "  Type: {1}\n",
        vi.VolumeARN,
        vi.VolumeType);
    return volumeInfo;
}
}
}
```

Deleting Snapshots by Using the AWS Tools for Windows PowerShell

To delete many snapshots associated with a volume, you can use a programmatic approach. The example following demonstrates how to delete snapshots using the AWS Tools for Windows PowerShell. To use the example script, you should be familiar with running a PowerShell script.

For more information, see [Getting Started](#) in the *AWS Tools for Windows PowerShell*. If you need to delete just a few snapshots, use the console as described in [Deleting snapshots of your storage volumes](#).

Example : Deleting Snapshots by Using the AWS Tools for Windows PowerShell

The following PowerShell script example lists the snapshots for each volume of a gateway and whether the snapshot start time is before or after a specified date. It uses the AWS Tools for Windows PowerShell cmdlets for Storage Gateway and Amazon EC2. The Amazon EC2 API includes operations for working with snapshots.

You need to update the script and provide your gateway Amazon Resource Name (ARN) and the number of days back you want to save snapshots. Snapshots taken before this cutoff are deleted. You also need to specify the Boolean value `viewOnly`, which indicates whether you want to view the snapshots to be deleted or to actually perform the snapshot deletions. Run the code first with just the view option (that is, with `viewOnly` set to `true`) to see what the code deletes.

```
<#
.DESCRIPTION
    Delete snapshots of a specified volume that match given criteria.

.NOTES
    PREREQUISITES:
    1) AWS Tools for Windows PowerShell from https://aws.amazon.com/powershell/
    2) Credentials and AWS Region stored in session using Initialize-AWSDefault.
    For more info see, https://docs.aws.amazon.com/powershell/latest/userguide/
    specifying-your-aws-credentials.html

.EXAMPLE
    powershell.exe .\SG_DeleteSnapshots.ps1
#>

# Criteria to use to filter the results returned.
$daysBack = 18
$gatewayARN = "**** provide gateway ARN ****"
$viewOnly = $true;

#ListVolumes
$volumesResult = Get-SGVolume -GatewayARN $gatewayARN
$volumes = $volumesResult.VolumeInfos
Write-Output("`nVolume List")
foreach ($volumes in $volumesResult)
```

```
{ Write-Output("`nVolume Info:")
  Write-Output("ARN: " + $volumes.VolumeARN)
  write-Output("Type: " + $volumes.VolumeType)
}

Write-Output("`nWhich snapshots meet the criteria?")
foreach ($volume in $volumesResult)
{
  $volumeARN = $volume.VolumeARN

  $volumeId = ($volumeARN-split"/")[3].ToLower()

  $filter = New-Object Amazon.EC2.Model.Filter
  $filter.Name = "volume-id"
  $filter.Value.Add($volumeId)

  $snapshots = get-EC2Snapshot -Filter $filter
  Write-Output("`nFor volume-id = " + $volumeId)
  foreach ($s in $snapshots)
  {
    $d = ([DateTime]::Now).AddDays(-$daysBack)
    $meetsCriteria = $false
    if ([DateTime]::Compare($d, $s.StartTime) -gt 0)
    {
      $meetsCriteria = $true
    }

    $sb = $s.SnapshotId + ", " + $s.StartTime + ", meets criteria for delete? " +
$meetsCriteria
    if (!$viewOnly -AND $meetsCriteria)
    {
      $resp = Remove-EC2Snapshot -SnapshotId $s.SnapshotId
      #Can get RequestId from response for troubleshooting.
      $sb = $sb + ", deleted? yes"
    }
    else {
      $sb = $sb + ", deleted? no"
    }
    Write-Output($sb)
  }
}
}
```

Understanding Volume Statuses and Transitions

Each volume has an associated status that tells you at a glance what the health of the volume is. Most of the time, the status indicates that the volume is functioning normally and that no action is needed on your part. In some cases, the status indicates a problem with the volume that might or might not require action on your part. You can find information following to help you decide when you need to act. You can see volume status on the Storage Gateway console or by using one of the Storage Gateway API operations, for example [DescribeCachediSCSIVolumes](#) or [DescribeStorediSCSIVolumes](#).

Topics

- [Understanding Volume Status](#)
- [Understanding Attachment Status](#)
- [Understanding Cached Volume Status Transitions](#)
- [Understanding Stored Volume Status Transitions](#)

Understanding Volume Status

The following table shows volume status on the Storage Gateway console. Volume status appears in the **Status** column for each storage volume on your gateway. A volume that is functioning normally has a status of **Available**.

In the following table, you can find a description of each storage volume status, and if and when you should act based on each status. The **Available** status is the normal status of a volume. A volume should have this status all or most of the time it's in use.

Status	Meaning
Available	<p>The volume is available for use. This status is the normal running status for a volume.</p> <p>When a Bootstrapping phase is completed, the volume returns to Available state. That is, the gateway has synchronized any changes made to the volume since it first entered Pass Through status.</p>
Bootstrapping	<p>The gateway is synchronizing data locally with a copy of the data stored in AWS. You typically don't need to take action for this status,</p>

Status	Meaning
	<p>because the storage volume automatically sees the Available status in most cases.</p> <p>The following are scenarios when a volume status is Bootstrapping:</p> <ul style="list-style-type: none"> • A gateway has unexpectedly shut down. • A gateway's upload buffer has been exceeded. In this scenario, bootstrapping occurs when your volume has the Pass Through status and the amount of free upload buffer increases sufficiently. You can provide additional upload buffer space as one way to increase the percentage of free upload buffer space. In this particular scenario, the storage volume goes from Pass Through to Bootstrapping to Available status. You can continue to use this volume during this bootstrapping period. However, you can't take snapshots of the volume at this point. • You are creating a stored Volume Gateway and preserving existing local disk data. In this scenario, your gateway starts uploading all of the data to AWS. The volume has the Bootstrapping status until all of the data from the local disk is copied to AWS. You can use the volume during this bootstrapping period. However, you can't take snapshots of the volume at this point.
Creating	The volume is currently being created and is not ready for use. The Creating status is transitional. No action is required.
Deleting	The volume is currently being deleted. The Deleting status is transitional. No action is required.
Irrecoverable	An error occurred from which the volume cannot recover. For information on what to do in this situation, see Troubleshooting volume issues .

Status	Meaning
Pass Through	<p>Data maintained locally is out of sync with data stored in AWS. Data written to a volume while the volume is in Pass Through status remains in the cache until the volume status is Bootstrapping. This data starts to upload to AWS when Bootstrapping status begins.</p> <p>The Pass Through status can occur for several reasons, listed following:</p> <ul style="list-style-type: none">• The Pass Through status occurs if your gateway has run out of upload buffer space. Your applications can continue to read from and write data to your storage volumes while the volumes have the Pass Through status. However, the gateway isn't writing any of your volume data to its upload buffer or uploading any of this data to AWS. <p>The gateway continues to upload any data written to the volume before the volume entered the Pass Through status. Any pending or scheduled snapshots of a storage volume fail while the volume has the Pass Through status. For information about what to do when your storage volume has the Pass Through status because the upload buffer has been exceeded, see Troubleshooting volume issues.</p> <p>To return to ACTIVE status, a volume in Pass Through must complete the Bootstrapping phase. During Bootstrapping, the volume re-establishes synchronization within AWS, so that it can resume the record (log) of changes to the volume, and activate CreateSnapshot functionality. During Bootstrapping, writes to the volume are recorded in upload buffer.</p> <ul style="list-style-type: none">• The Pass Through status occurs when there is more than one storage volume bootstrapping at once. Only one gateway storage volume can bootstrap at a time. For example, suppose that you create two storage volumes and choose to preserve existing data on both of them. In this case, the second storage volume has the Pass Through status until the first storage volume finishes bootstrap

Status	Meaning
	<p>ping. In this scenario, you don't need to act. Each storage volume changes to the Available status automatically when it is finished being created. You can read and write to the storage volume while it has the Pass Through or Bootstrapping status.</p> <ul style="list-style-type: none"> • Infrequently, the Pass Through status can indicate that a disk allocated for upload buffer use has failed. For information about what action to take in this scenario, see Troubleshooting volume issues. • The Pass Through status can occur when a volume is in Active or Bootstrapping state. In this case, the volume receives a write, but the upload buffer has insufficient capacity to record (log) that write. • The Pass Through status occurs when a volume is in any state and the gateway is not shut down cleanly. This type of shutdown can happen because the software crashed or the VM was powered off. In this case, a volume in any state transitions to Pass Through status.
Restoring	<p>The volume is being restored from an existing snapshot. This status applies only for stored volumes. For more information, see How Volume Gateway works.</p> <p>If you restore two storage volumes at the same time, both storage volumes show Restoring as their status. Each storage volume changes 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 has the Restoring status.</p>

Status	Meaning
Restoring Pass Through	<p>The volume is being restored from an existing snapshot and has encountered an upload buffer issue. This status applies only for stored volumes. For more information, see How Volume Gateway works.</p> <p>One reason that can cause the Restoring Pass Through status is if your gateway has run out of upload buffer space. Your applications can continue to read from and write data to your storage volumes while they have the Restoring Pass Through status. However, you can't take snapshots of a storage volume during the Restoring Pass Through status period. For information about what action to take when your storage volume has the Restoring Pass Through status because upload buffer capacity has been exceeded, see Troubleshooting volume issues.</p> <p>Infrequently, the Restoring Pass Through status can indicate that a disk allocated for an upload buffer has failed. For information about what action to take in this scenario, see Troubleshooting volume issues.</p>
Upload Buffer Not Configured	<p>You can't create or use the volume because the gateway doesn't have an upload buffer configured. For information on how to add upload buffer capacity for volumes in a cached volume setup, see Determining the size of upload buffer to allocate. For information on how to add upload buffer capacity for volumes in a stored volume setup, see Determining the size of upload buffer to allocate.</p>

Understanding Attachment Status

You can detach a volume from a gateway or attach it to a gateway by using the Storage Gateway console or API. The following table shows volume attachment status on the Storage Gateway console. Volume attachment status appears in the **Attachment status** column for each storage volume on your gateway. For example, a volume that is detached from a gateway has a status of **Detached**. For information about how to detach and attach a volume, see [Moving Your Volumes to a Different Gateway](#).

Status	Meaning
Attached	The volume is attached to a gateway.
Detached	The volume is detached from a gateway.
Detaching	The volume is being detached from a gateway. When you are detach ing a volume and the volume doesn't have data on it, you might not see this status.

Understanding Cached Volume Status Transitions


Use the following state diagram to understand the most common transitions between statuses for volumes in cached gateways. You don't need to understand the diagram in detail to use your gateway effectively. Rather, the diagram provides detailed information if you are interested in knowing more about how Volume Gateways work.

The diagram doesn't show the **Upload Buffer Not Configured** status or the **Deleting** status. Volume states in the diagram appear as green, yellow, and red boxes. You can interpret the colors as described following.

Color	Volume Status
Green	The gateway is operating normally. The volume status is Available or eventually becomes Available .
Yellow	The volume has the Pass Through status, which indicates there is a potential issue with the storage volume. If this status appears because the upload buffer space is filled, then in some cases buffer space becomes available again. At that point, the storage volume self-corrects to the Available status. In other cases, you might have to add more upload buffer space to your gateway to allow the storage volume status to become Available . For information on how

Color	Volume Status
	to troubleshoot a case when upload buffer capacity has been exceeded, see Troubleshooting volume issues . For information on how to add upload buffer capacity, see Determining the size of upload buffer to allocate .
Red	The storage volume has the Irrecoverable status. In this case, you should delete the volume. For information on how to do this, see To delete a volume .

In the diagram, a transition between two states is depicted with a labeled line. For example, the transition from the **Creating** status to the **Available** status is labeled as *Create Basic Volume* or *Create Volume from Snapshot*. This transition represents creating a cached volume. For more information about creating storage volumes, see [Adding and expanding volumes](#).



Key

Gateway Operating Normally	Temporary State or Recoverable Condition*	Irrecoverable
----------------------------	---	---------------

- * e.g. run out of upload buffer
- ** e.g. lost connectivity

Note

The volume status of **Pass Through** appears as yellow in this diagram. However, this doesn't match the color of this status icon in the **Status** box of the Storage Gateway console.

Understanding Stored Volume Status Transitions


Use the following state diagram to understand the most common transitions between statuses for volumes in stored gateways. You don't need to understand the diagram in detail to use your gateway effectively. Rather, the diagram provides detailed information if you are interested in understanding more about how Volume Gateways work.

The diagram doesn't show the **Upload Buffer Not Configured** status or the **Deleting** status. Volume states in the diagram appear as green, yellow, and red boxes. You can interpret the colors as described following.

Color	Volume Status
Green	The gateway is operating normally. The volume status is Available or eventually becomes Available .
Yellow	When you are creating a storage volume and preserving data, then the path from the Creating status to the Pass Through status occurs if another volume is bootstrapping. In this case, the volume with the Pass Through status goes to the Bootstrapping status and then to the Available status when the first volume is finished bootstrapping. Other than the specific scenario mentioned, yellow (Pass Through status) indicates that there is a potential issue with the storage volume, the most common one being an upload buffer issue. If upload buffer capacity has been exceeded, then in some cases buffer space becomes available again. At that point, the storage volume self-corrects to the Available status. In other cases, you might have to add more upload buffer capacity to your gateway to return the storage volume to the Available status. For information on how to troubleshoot a case when upload buffer capacity has been exceeded, see Troubleshoot

Color	Volume Status
	ooting volume issues . For information on how to add upload buffer capacity, see Determining the size of upload buffer to allocate .
Red	The storage volume has the Irrecoverable status. In this case, you should delete the volume. For information on how to do this, see Deleting storage volumes .

In the following diagram, a transition between two states is depicted with a labeled line. For example, the transition from the **Creating** status to the **Available** status is labeled as *Create Basic Volume*. This transition represents creating a storage volume without preserving data or creating the volume from a snapshot.



Note

The volume status of **Pass Through** appears as yellow in this diagram. However, this doesn't match the color of this status icon in the **Status** box of the Storage Gateway console.

Moving your data to a new gateway

You can move data between gateways as your data and performance needs grow, or if you receive an AWS notification to migrate your gateway. The following are some reasons for doing this:

- Move your data to better host platforms or newer Amazon EC2 instances.
- Refresh the underlying hardware for your server.

The steps that you follow to move your data to a new gateway depend on the gateway type that you have.

 **Note**

Data can only be moved between the same gateway types.

Moving stored volumes to a new stored Volume Gateway

To move your stored volume to a new stored Volume Gateway

1. Stop any applications that are writing to the old stored Volume Gateway.
2. Use the following steps to create a snapshot of your volume, and then wait for the snapshot to complete.
 - a. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
 - b. In the navigation pane, choose **Volumes**, and then choose the volume that you want to create the snapshot from.
 - c. For **Actions**, choose **Create snapshot**.
 - d. In the **Create snapshot** dialog box, enter a snapshot description, and then choose **Create snapshot**.

You can verify that the snapshot was created using the console. If data is still uploading to the volume, wait until the upload is complete before you go to the next step. To see the snapshot status and validate that none are pending, select the snapshot links on the volumes.

3. Use the following steps to stop the old stored Volume Gateway:
 - a. In the navigation pane, choose **Gateways**, and then choose the old stored Volume Gateway that you want to stop. The status of the gateway is **Running**.

- b. For **Actions**, choose **Stop gateway**. Verify the ID of the gateway from the dialog box, and then choose **Stop gateway**.

While the gateway is stopping, you might see a message that indicates the status of the gateway. When the gateway shuts down, a message and a **Start gateway** button appear in the **Details** tab. When the gateway shuts down, the status of the gateway is **Shutdown**.

- c. Shut down the VM using the hypervisor controls.

For more information about stopping a gateway, see [Starting and Stopping a Volume Gateway](#).

4. Detach the storage disks associated with your stored volumes from the gateway VM. This excludes the root disk of the VM.
5. Activate a new stored Volume Gateway with a new hypervisor VM image available from the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
6. Attach the physical storage disks that you detached from the old stored Volume Gateway VM in step 5.
7. To preserve existing data on the disk, use the following steps to create stored volumes.
 - a. On the Storage Gateway console, choose **Create volume**.
 - b. In the **Create volume** dialog box, select the stored Volume Gateway that you created in step 5.
 - c. Choose a **Disk ID** value from the list.
 - d. For **Volume content**, select the **Preserve existing data on the disk** option.

For more information about creating volumes, see [Creating a storage volume](#).

8. (Optional) In the **Configure CHAP authentication** wizard that appears, enter the **Initiator name**, **Initiator secret**, and **Target secret**, and then choose **Save**.

For more information about working with Challenge-Handshake Authentication Protocol (CHAP) authentication, see [Configuring CHAP Authentication for Your iSCSI Targets](#).

9. Start the application that writes to your stored volume.
10. When you have confirmed that your new stored Volume Gateway is working correctly, you can delete the old stored Volume Gateway.

⚠ Important

Before you delete a gateway, be sure that no applications are currently writing to that gateway's volumes. If you delete a gateway while it is in use, data loss can occur.

Use the following steps to delete the old stored Volume Gateway:

⚠ Warning

When a gateway is deleted, there is no way to recover it.

- a. In the navigation pane, choose **Gateways**, and then choose the old stored Volume Gateway that you want to delete.
 - b. For **Actions**, choose **Delete gateway**.
 - c. In the confirmation dialog box that appears, select the check box to confirm your deletion. Make sure that the gateway ID listed specifies the old stored Volume Gateway that you want to delete, and then choose **Delete**.
11. Delete the old gateway VM. For information about deleting a VM, see the documentation for your hypervisor.

Moving cached volumes to a new gateway virtual machine

To move your cached volumes to a new cached Volume Gateway virtual machine (VM)

1. Stop any applications that are writing to the old cached Volume Gateway.
2. Unmount or disconnect iSCSI volumes from any clients that are using them. This helps keep data on those volumes consistent by preventing clients from changing or adding data to those volumes.
3. Use the following steps to create a snapshot of your volume, and then wait for the snapshot to complete.
 - a. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.

- b. In the navigation pane, choose **Volumes**, and then choose the volume that you want to create the snapshot from.
- c. For **Actions**, choose **Create snapshot**.
- d. In the **Create snapshot** dialog box, enter a snapshot description, and then choose **Create snapshot**.

You can verify that the snapshot was created using the console. If data is still uploading to the volume, wait until the upload is complete before you go to the next step. To see the snapshot status and validate that none are pending, select the snapshot links on the volumes.

For more information about checking volume status in the console, see [Understanding Volume Statuses and Transitions](#). For information about cached volume status, see [Understanding Cached Volume Status Transitions](#).

4. Use the following steps to stop the old cached Volume Gateway:
 - a. In the navigation pane, choose **Gateways**, and then choose the old cached Volume Gateway that you want to stop. The status of the gateway is **Running**.
 - b. For **Actions**, choose **Stop gateway**. Verify the ID of the gateway from the dialog box, and then choose **Stop gateway**. Make a note of the gateway ID, as it is needed in a later step.

While the old gateway is stopping, you might see a message that indicates the status of the gateway. When the old gateway shuts down, a message and a **Start gateway** button appear in the **Details** tab. When the gateway shuts down, the status of the gateway is **Shutdown**.

- c. Shut down the old VM using the hypervisor controls. For more information about shutting down an Amazon EC2 instance, see [Stopping and starting your instances](#) in the *Amazon EC2 User Guide*. For more information about shutting down a KVM, VMware, or Hyper-V VM, see your hypervisor documentation.

For more information about stopping a gateway, see [Starting and Stopping a Volume Gateway](#).

5. Detach all disks, including the root disk, cache disks, and upload buffer disks, from the old gateway VM.

Note

Make a note of the root disk's volume ID, as well as the gateway ID associated with that root disk. You detach this disk from the new Storage Gateway hypervisor in a later step. (See step 11.)

If you are using an Amazon EC2 instance as the VM for your cached Volume Gateway, see [Detaching an Amazon EBS volume from a Linux instance](#) in the *Amazon EC2 User Guide*. For information about detaching disks from a KVM, VMware, or Hyper-V VM, see the documentation for your hypervisor.

6. Create a new Storage Gateway hypervisor VM instance, but don't activate it as a gateway. For more information about creating a new Storage Gateway hypervisor VM, see [Set up a Volume Gateway](#). This new gateway will assume the identity of the old gateway.

Note

Do not add disks for cache or upload buffer to the new VM. Your new VM will use the same cache disks and upload buffer disks that were used by the old VM.

7. Your new Storage Gateway hypervisor VM instance should use the same network configuration as the old VM. The default network configuration for the gateway is Dynamic Host Configuration Protocol (DHCP). With DHCP, your gateway is automatically assigned an IP address.

If you need to manually configure a static IP address for your new VM, see [Configuring Your Gateway Network](#) for more details. If your gateway must use a Socket Secure version 5 (SOCKS5) proxy to connect to the internet, see [Configuring a SOCKS5 proxy for your on-premises gateway](#) for more details.

8. Start the new VM.
9. Attach the disks that you detached from the old cached Volume Gateway VM in step 5, to the new cached Volume Gateway. Attach them in the same order to the new gateway VM as they are on the old gateway VM.

All disks must make the transition unchanged. Do not change volume sizes, as that will cause metadata to become inconsistent.

10. Initiate the gateway migration process by connecting to the new VM with a URL that uses the following format.

```
http://your-VM-IP-address/migrate?gatewayId=your-gateway-ID
```

You can re-use the same IP address for the new gateway VM as you used for the old gateway VM. Your URL should look similar to the example following.

```
http://198.51.100.123/migrate?gatewayId=sgw-12345678
```

Use this URL from a browser, or from the command line using `curl`, to initiate the migration process.

When the gateway migration process is successfully initiated, you will see the following message:

```
Successfully imported Storage Gateway information. Please refer to
Storage Gateway documentation to perform the next steps to complete the
migration.
```

11. Detach the old gateway's root disk, whose volume ID you noted in step 5.
12. Start the gateway.

Use the following steps to start the new cached Volume Gateway:

- a. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
- b. In the navigation pane, choose **Gateways** and then choose the new gateway you want to start. The status of the gateway is **Shutdown**.
- c. Choose **Details**, and then choose **Start gateway**.

For more information about starting a gateway, see [Starting and Stopping a Volume Gateway](#).

13. Your volumes should now be available to your applications at the new gateway VM's IP address.
14. Confirm that your volumes are available, and delete the old gateway VM. For information about deleting a VM, see the documentation for your hypervisor.

Monitoring Storage Gateway

This section describes how to monitor a Storage Gateway, including monitoring resources associated with the gateway, using Amazon CloudWatch. You can monitor the gateway's upload buffer and cache storage. You use the Storage Gateway console to view metrics and alarms 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 taken to retrieve data from the Amazon Web Services Cloud. With metrics, you can track the health of your gateway and set up alarms to notify you when one or more metrics fall outside a defined threshold.

Storage Gateway provides CloudWatch metrics at no additional charge. Storage Gateway metrics are recorded for a period of two weeks. By using these metrics, you can access historical information and get a better perspective on how your gateway and volumes are performing. Storage Gateway also provides CloudWatch alarms, except high-resolution alarms, at no additional charge. For more information about CloudWatch pricing, see [Amazon CloudWatch pricing](#). For more information about CloudWatch, see [Amazon CloudWatch User Guide](#).

For information specific to monitoring a Volume Gateway and its associated resources, see [Monitoring your Volume Gateway](#).

Topics


- [Understanding gateway metrics](#)
- [Monitoring the upload buffer](#)
- [Monitoring cache storage](#)
- [Understanding CloudWatch alarms](#)
- [Creating recommended CloudWatch alarms for your gateway](#)
- [Creating a custom CloudWatch alarm for your gateway](#)
- [Monitoring your Volume Gateway](#)

Understanding gateway metrics

For the discussion in this topic, we define *gateway* metrics as metrics that are scoped to the gateway—that is, they measure something about the gateway. Because 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 metric includes the activity of all the volumes on the gateway.

When working with gateway metric data, you specify the unique identification of the gateway that you are interested in viewing metrics for. To do this, you specify both the `GatewayId` and the `GatewayName` values. 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 Amazon CloudWatch Metrics](#).

 **Note**

Some metrics return data points only when new data has been generated during the most recent monitoring period.

Metric	Description
AvailabilityNotifications	<p>Number of availability-related health notifications generated by the gateway.</p> <p>Use this metric with the Sum statistic to observe whether the gateway is experiencing any availability-related events. For details about the events, check your configured CloudWatch log group.</p> <p>Unit: Number</p>
CacheHitPercent	<p>Percent of application reads served from the cache. The sample is taken at the end of the reporting period.</p> <p>Unit: Percent</p>

Metric	Description	
CacheUsed	<p>The total number of bytes being used in the gateway's cache storage. The sample is taken at the end of the reporting period.</p> <p>Unit: Bytes</p>	
IoWaitPercent	<p>Percent of time that the gateway is waiting on a response from the local disk.</p> <p>Unit: Percent</p>	
MemTotalBytes	<p>Amount of RAM provisioned to the gateway VM, in bytes.</p> <p>Unit: Bytes</p>	
MemUsedBytes	<p>Amount of RAM currently in use by the gateway VM, in bytes.</p> <p>Unit: Bytes</p>	
QueuedWrites	<p>The number of bytes waiting to be written to AWS, sampled at the end of the reporting period for all volumes in the gateway. These bytes are kept in your gateway's working storage.</p> <p>Unit: Bytes</p>	

Metric	Description	
ReadBytes	<p>The total number of bytes read from your on-premises applications in the reporting period for all volumes in the gateway.</p> <p>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.</p> <p>Unit: Bytes</p>	
ReadTime	<p>The total number of milliseconds spent to do read operations from your on-premises applications in the reporting period for all volumes in the gateway.</p> <p>Use this metric with the Average statistic to measure latency.</p> <p>Unit: Milliseconds</p>	
TimeSinceLastRecoveryPoint	<p>The time since the last available recovery point. For more information, see Your Cached Gateway is Unreachable And You Want to Recover Your Data.</p> <p>Unit: Seconds</p>	

Metric	Description	
TotalCacheSize	<p>The total size of the cache in bytes. The sample is taken at the end of the reporting period.</p> <p>Unit: Bytes</p>	
UploadBufferPercentUsed	<p>Percent use of the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Unit: Percent</p>	
UploadBufferUsed	<p>The total number of bytes being used in the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Unit: Bytes</p>	
UserCpuPercent	<p>Percent of CPU time spent on gateway processing, averaged across all cores.</p> <p>Unit: Percent</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>Unit: Bytes</p>	

Metric	Description	
WorkingStoragePercentUsed	<p>Percent use of the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Unit: Percent</p>	
WorkingStorageUsed	<p>The total number of bytes being used in the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Unit: Bytes</p>	
WriteBytes	<p>The total number of bytes written to your on-premises applications in the reporting period for all volumes in the gateway.</p> <p>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.</p> <p>Unit: Bytes</p>	

Metric	Description	
WriteTime	<p>The total number of milliseconds spent to do write operations from your on-premises applications in the reporting period for all volumes in the gateway.</p> <p>Use this metric with the Average statistic to measure latency.</p> <p>Unit: Milliseconds</p>	

Dimensions for Storage Gateway metrics

The CloudWatch namespace for the Storage Gateway service is `AWS/StorageGateway`. Data is available automatically in 5-minute periods at no charge.

Dimension	Description
GatewayId , GatewayName	<p>These dimensions filter the data that you request to gateway-specific metrics. You can identify a gateway to work by the value for <code>GatewayId</code> or <code>GatewayName</code> . If the name of your gateway was different for the time range that you are interested in viewing metrics, use the <code>GatewayId</code> .</p> <p>Throughput and latency data of a gateway is based on all the volumes for the gateway. For information about working with gateway metrics, see Measuring Performance Between Your Gateway and AWS.</p>
VolumeId	<p>This dimension filters the data you request to volume-specific metrics. Identify a storage volume to work with by its <code>VolumeId</code> value. For information about working with volume</p>

Dimension	Description
	metrics, see Measuring Performance Between Your Application and Gateway .

Monitoring the upload buffer

You can find information following about how to monitor a gateway's upload buffer and how to create an alarm so that you get a notification when the buffer exceeds a specified threshold. By using this approach, you can add buffer storage to a gateway before it fills completely and your storage application stops backing up to AWS.

You monitor the upload buffer in the same way in both the cached-volume and Tape Gateway architectures. For more information, see [How Volume Gateway works](#).

Note

The `WorkingStoragePercentUsed`, `WorkingStorageUsed`, and `WorkingStorageFree` metrics represent the upload buffer for stored volumes only before the release of the cached-volume feature in Storage Gateway. Now, use the equivalent upload buffer metrics `UploadBufferPercentUsed`, `UploadBufferUsed`, and `UploadBufferFree`. These metrics apply to both gateway architectures.

Item of Interest	How to Measure
Upload buffer usage	Use the <code>UploadBufferPercentUsed</code> , <code>UploadBufferUsed</code> , and <code>UploadBufferFree</code> metrics with the Average statistic. For example, use the <code>UploadBufferUsed</code> with the Average statistic to analyze the storage usage over a time period.

To measure the percent of the upload buffer that is used

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose the **StorageGateway: Gateway Metrics** dimension, and find the gateway that you want to work with.

3. Choose the `UploadBufferPercentUsed` metric.
4. For **Time Range**, choose a value.
5. Choose the `Average` statistic.
6. For **Period**, choose a value of 5 minutes to match the default reporting time.

The resulting time-ordered set of data points contains the percent used of the upload buffer.

Using the following procedure, you can create an alarm using the CloudWatch console. To learn more about alarms and thresholds, see [Creating CloudWatch Alarms](#) in the *Amazon CloudWatch User Guide*.

To set an upper threshold alarm for a gateway's upload buffer

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose **Create Alarm** to start the Create Alarm wizard.
3. Specify a metric for your alarm:
 - a. On the **Select Metric** page of the Create Alarm wizard, choose the **AWS/StorageGateway:GatewayId,GatewayName** dimension, and then find the gateway that you want to work with.
 - b. Choose the `UploadBufferPercentUsed` metric. Use the `Average` statistic and a period of 5 minutes.
 - c. Choose **Continue**.
4. Define the alarm name, description, and threshold:
 - a. On the **Define Alarm** page of the Create Alarm wizard, identify your alarm by giving it a name and description in the **Name** and **Description** boxes.
 - b. Define the alarm threshold.
 - c. Choose **Continue**.
5. Configure an email action for the alarm:
 - a. On the **Configure Actions** page of the Create Alarm wizard, choose **Alarm for Alarm State**.
 - b. Choose **Choose or create email topic** for **Topic**.

To create an email topic means that you set up an Amazon SNS topic. For more information about Amazon SNS, see [Set Up Amazon SNS](#) in the *Amazon CloudWatch User Guide*.

- c. For **Topic**, enter a descriptive name for the topic.
 - d. Choose **Add Action**.
 - e. Choose **Continue**.
6. Review the alarm settings, and then create the alarm:
- a. On the **Review** page of the Create Alarm wizard, review the alarm definition, metric, and associated actions to take (for example, sending an email notification).
 - b. After reviewing the alarm summary, choose **Save Alarm**.
7. Confirm your subscription to the alarm topic:
- a. Open the Amazon SNS email that was sent to the email address that you specified when creating the topic.
 - b. Confirm your subscription by clicking the link in the email.

A subscription confirmation appears.

Monitoring cache storage

You can find information following about how to monitor a gateway's cache storage and how to create an alarm so that you get a notification when parameters of the cache pass specified thresholds. Using this alarm, you know when to add cache storage to a gateway.

You only monitor cache storage in the cached volumes architecture. For more information, see [How Volume Gateway works](#).

Item of Interest	How to Measure
Total usage of cache	Use the <code>CachePercentUsed</code> and <code>TotalCacheSize</code> metrics with the <code>Average</code> statistic. For example, use the <code>CachePercentUsed</code> with the <code>Average</code> statistic to analyze the cache usage over a period of time.

Item of Interest	How to Measure
	The <code>TotalCacheSize</code> metric changes only when you add cache to the gateway.
Percent of read requests that are served from the cache	Use the <code>CacheHitPercent</code> metric with the <code>Average</code> statistic. Typically, you want <code>CacheHitPercent</code> to remain high.
Percent of the cache that is dirty—that is, it contains content that has not been uploaded to AWS	Use the <code>CachePercentDirty</code> metrics with the <code>Average</code> statistic. Typically, you want <code>CachePercentDirty</code> to remain low.

To measure the percent of a cache that is dirty for a gateway and all its volumes

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose the **StorageGateway: Gateway Metrics** dimension, and find the gateway that you want to work with.
3. Choose the `CachePercentDirty` metric.
4. For **Time Range**, choose a value.
5. Choose the `Average` statistic.
6. For **Period**, choose a value of 5 minutes to match the default reporting time.

The resulting time-ordered set of data points contains the percentage of the cache that is dirty over the 5 minutes.

To measure the percent of the cache that is dirty for a volume

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose the **StorageGateway: Volume Metrics** dimension, and find the volume that you want to work with.
3. Choose the `CachePercentDirty` metric.
4. For **Time Range**, choose a value.

5. Choose the `Average` statistic.
6. For **Period**, choose a value of 5 minutes to match the default reporting time.

The resulting time-ordered set of data points contains the percentage of the cache that is dirty over the 5 minutes.

Understanding CloudWatch alarms

CloudWatch alarms monitor information about your gateway based on metrics and expressions. You can add CloudWatch alarms for your gateway and view their statuses in the Storage Gateway console. For more information about the metrics that are used to monitor Volume Gateway, see [Understanding gateway metrics](#) and [Understanding Volume Metrics](#). For each alarm, you specify conditions that will initiate its ALARM state. Alarm status indicators in the Storage Gateway console turn red when in the ALARM state, making it easier for you to monitor status proactively. You can configure alarms to invoke actions automatically based on sustained changes in state. For more information about CloudWatch alarms, see [Using Amazon CloudWatch alarms](#) in the *Amazon CloudWatch User Guide*.

Note

If you don't have permission to view CloudWatch, you can't view the alarms.

For each activated gateway, we recommend that you create the following CloudWatch alarms:

- High IO wait: `IoWaitpercent >= 20` for 3 datapoints in 15 minutes
- Cache percent dirty: `CachePercentDirty > 80` for 4 datapoints within 20 minutes
- Health notifications: `HealthNotifications >= 1` for 1 datapoint within 5 minutes. When configuring this alarm, set **Missing data treatment** to **notBreaching**.

Note

You can set a health notification alarm only if the gateway had a previous health notification in CloudWatch.

For gateways on VMware host platforms with HA mode activated, we also recommend this additional CloudWatch alarm:

- Availability notifications: `AvailabilityNotifications >= 1` for 1 datapoint within 5 minutes. When configuring this alarm, set **Missing data treatment** to **notBreaching**.

The following table describes the state of an alarm.

State	Description
OK	The metric or expression is within the defined threshold.
Alarm	The metric or expression is outside of the defined threshold.
Insufficient data	The alarm has just started, the metric is not available, or not enough data is available for the metric to determine the alarm state.
None	No alarms are created for the gateway. To create a new alarm, see Creating a custom CloudWatch alarm for your gateway .
Unavailable	The state of the alarm is unknown. Choose Unavailable to view error information in the Monitoring tab.

Creating recommended CloudWatch alarms for your gateway

When you create a new gateway using the Storage Gateway console, you can choose to create all recommended CloudWatch alarms automatically as part of the initial setup process. For more information, see [Configure your Volume Gateway](#). If you want to add or update recommended CloudWatch alarms for an existing gateway, use the following procedure.

To add or update recommended CloudWatch alarms for an existing gateway

Note

This feature requires CloudWatch policy permissions, which are *not* automatically granted as part of the preconfigured Storage Gateway full access policy. Make sure your security policy grants the following permissions before you attempt to create recommended CloudWatch alarms:

- `cloudwatch:PutMetricAlarm` - create alarms
- `cloudwatch:DisableAlarmActions` - turn alarm actions off
- `cloudwatch:EnableAlarmActions` - turn alarm actions on
- `cloudwatch>DeleteAlarms` - delete alarms

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home/>.
2. In the navigation pane, choose **Gateways**, then choose the gateway for which you want to create recommended CloudWatch alarms.
3. On the gateway details page, choose the **Monitoring** tab.
4. Under **Alarms**, choose **Create recommended alarms**. The recommended alarms are created automatically.

The **Alarms** section lists all CloudWatch alarms for a specific gateway. From here, you can select and delete one or more alarms, turn alarm actions on or off, and create new alarms.

Creating a custom CloudWatch alarm for your gateway

CloudWatch uses Amazon Simple Notification Service (Amazon SNS) to send alarm notifications when an alarm changes state. An alarm watches a single metric over a time period that you specify, and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The action is a notification that's sent to an Amazon SNS topic. You can create an Amazon SNS topic when you create a CloudWatch alarm. For more information about Amazon SNS, see [What is Amazon SNS?](#) in the *Amazon Simple Notification Service Developer Guide*.

To create a CloudWatch alarm in the Storage Gateway console

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home/>.
2. In the navigation pane, choose **Gateways**, then choose the gateway for which you want to create an alarm.
3. On the gateway details page, choose the **Monitoring** tab.
4. Under **Alarms**, choose **Create alarm** to open the CloudWatch console.
5. Use the CloudWatch console to create the type of alarm that you want. You can create the following types of alarms:

- **Static threshold alarm:** An alarm based on a set threshold for a chosen metric. The alarm enters the ALARM state when the metric breaches the threshold for a specified number of evaluation periods.

To create a static threshold alarm, see [Creating a CloudWatch alarm based on a static threshold](#) in the *Amazon CloudWatch User Guide*.

- **Anomaly detection alarm:** Anomaly detection mines past metric data and creates a model of expected values. You set a value for the anomaly detection threshold, and CloudWatch uses this threshold with the model to determine the "normal" range of values for the metric. A higher value for the threshold produces a thicker band of "normal" values. You can choose to activate the alarm only when the metric value is above the band of expected values, only when it's below the band, or when it's above or below the band.

To create an anomaly detection alarm, see [Creating a CloudWatch alarm based on anomaly detection](#) in the *Amazon CloudWatch User Guide*.

- **Metric math expression alarm:** An alarm based one or more metrics used in a math expression. You specify the expression, threshold, and evaluation periods.

To create a metric math expression alarm, see [Creating a CloudWatch alarm based on a metric math expression](#) in the *Amazon CloudWatch User Guide*.

- **Composite alarm:** An alarm that determines its alarm state by watching the alarm states of other alarms. A composite alarm can help you reduce alarm noise.

To create a composite alarm, see [Creating a composite alarm](#) in the *Amazon CloudWatch User Guide*.

6. After you create the alarm in the CloudWatch console, return to the Storage Gateway console. You can view the alarm by doing one of the following:
 - In the navigation pane, choose **Gateways**, then choose the gateway for which you want to view alarms. On the **Details** tab, under **Alarms**, choose **CloudWatch Alarms**.
 - In the navigation pane, choose **Gateways**, choose the gateway for which you want to view alarms, then choose the **Monitoring** tab.

The **Alarms** section lists all of the CloudWatch alarms for a specific gateway. From here, you can select and delete one or more alarms, turn alarm actions on or off, and create new alarms.

- In the navigation pane, choose **Gateways**, then choose the alarm state of the gateway for which you want to view alarms.

For information about how to edit or delete an alarm, see [Editing or deleting a CloudWatch alarm](#).

Note

When you delete a gateway using the Storage Gateway console, all CloudWatch alarms associated with the gateway are also automatically deleted.

Monitoring your Volume Gateway

The topics in this section describe how to monitor Volume Gateway in either cached volume or stored volume setup, including monitoring the volumes associated with the gateway and monitoring the upload buffer. You use the AWS Management Console to view metrics for your gateway. For example, you can view the number of bytes used in read and write operations, the time spent in read and write operations, and the time taken to retrieve data from the Amazon Web Services cloud. With metrics, you can track the health of your gateway and set up alarms to notify you when one or more metrics fall outside a defined threshold.

Storage Gateway provides CloudWatch metrics at no additional charge. Storage Gateway metrics are recorded for a period of two weeks. By using these metrics, you can access historical information and get a better perspective on how your gateway and volumes are performing. For detailed information about CloudWatch, see the [Amazon CloudWatch User Guide](#).

Topics

- [Getting Volume Gateway health logs with Amazon CloudWatch Logs](#) - Learn how to use Amazon CloudWatch Logs to get information about the health of your Volume Gateway and related resources.
- [Using Amazon CloudWatch Metrics](#) - Learn how to get monitoring data for your gateway using either the AWS Management Console or the CloudWatch API.
- [Measuring Performance Between Your Application and Gateway](#) - Learn how to measure data throughput, data latency, and operations per second to understand performance between your applications and your gateway.
- [Measuring Performance Between Your Gateway and AWS](#) - Learn how to measure data throughput, data latency, and operations per second to understand performance between your gateway and the AWS cloud.
- [Understanding volume metrics](#) - Learn how to measure metrics that provide data about the volumes associated with a gateway.

Getting Volume Gateway health logs with Amazon CloudWatch Logs

You can use Amazon CloudWatch Logs to get information about the health of your Volume Gateway and related resources. You can use these logs to monitor your gateway for errors that it encounters. In addition, you can use Amazon CloudWatch subscription filters to automate processing of the log information in real time. For more information, see [Real-time Processing of Log Data with Subscriptions](#) in the *Amazon CloudWatch User Guide*.

For example, suppose that your gateway is deployed in a cluster activated with VMware High Availability (HA) and you need to know about any errors. You can configure a CloudWatch log group to monitor your gateway and get notified when your gateway encounters an error. You can either configure the group when you are activating the gateway or after your gateway is activated and up and running. For information about how to configure a CloudWatch log group when activating a gateway, see [Configure your Volume Gateway](#). For general information about CloudWatch log groups, see [Working with Log Groups and Log Streams](#) in the *Amazon CloudWatch User Guide*.

For information about how to troubleshoot and fix these types of errors, see [Troubleshooting volume issues](#).

The following procedure shows you how to configure a CloudWatch log group after your gateway is activated.

To configure a CloudWatch log group to work with your gateway

1. Sign in to the AWS Management Console and open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the left navigation pane, choose **Gateways**, and then choose the gateway that you want to configure the CloudWatch log group for.
3. For **Actions**, choose **Edit gateway information**, or on the **Details** tab, under **Health logs** and **Not Enabled**, choose **Configure log group** to open the **Edit *CustomerGatewayName*** dialog box.
4. For **Gateway health log group**, choose one of the following:
 - **Disable logging** if you don't want to monitor your gateway using CloudWatch log groups.
 - **Create a new log group** to create a new CloudWatch log group.
 - **Use an existing log group** to use a CloudWatch log group that already exists. Choose a log group from the **Existing log group list**.
5. Choose **Save changes**.
6. To see the health logs for your gateway, do the following:
 1. In the left navigation pane, choose **Gateways**, and then choose the gateway that you configured the CloudWatch log group for.
 2. Choose the **Details** tab, and under **Health logs**, choose **CloudWatch Logs**. The **Log group details** page opens in the Amazon CloudWatch console.

Using Amazon CloudWatch Metrics

You can get monitoring data for your gateway using either the AWS Management Console or the CloudWatch API. The console displays a series of graphs based on the raw data from the CloudWatch API. You can also use the CloudWatch API through one of the [AWS Software Development Kits \(SDKs\)](#) or the [Amazon CloudWatch API](#) tools. Depending on your needs, you might prefer to use either the graphs displayed in the console or retrieved from the API.

Regardless of which method you choose to use to work with metrics, you must specify the following information:

- The metric dimension to work with. A *dimension* is a name-value pair that helps you to uniquely identify a metric. The dimensions for Storage Gateway are GatewayId, GatewayName, and

VolumeId. In the CloudWatch console, you can use the Gateway Metrics and Volume Metrics views to easily select gateway-specific and volume-specific dimensions. For more information about dimensions, see [Dimensions](#) in the *Amazon CloudWatch User Guide*.

- The metric name, such as ReadBytes.

The following table summarizes the types of Storage Gateway metric data that you can use.

CloudWatch Namespace	Dimension	Description
AWS/StorageGateway	GatewayId , GatewayName	<p>These dimensions filter for metric data that describes aspects of the gateway. You can identify a gateway to work with by specifying both the GatewayId and the GatewayName dimensions.</p> <p>Throughput and latency data of a gateway are based on all the volumes in the gateway.</p> <p>Data is available automatically in 5-minute periods at no charge.</p>
	VolumeId	<p>This dimension filters for metric data that is specific to a volume. Identify a volume to work with by its VolumeId dimension.</p> <p>Data is available automatically in 5-minute periods at no charge.</p>

Working with gateway and volume metrics is similar to working with other service metrics. You can find a discussion of some of the most common metrics tasks in the CloudWatch documentation listed following:

- [Viewing 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 that is using your gateway is performing. When you use the correct aggregation statistic, you can use Storage Gateway metrics to measure these values.

A *statistic* is an aggregation of a metric over a specified period of time. When you view the values of a metric in 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 input/output operations per second (IOPS). For more information, see [Statistics](#) in the *Amazon CloudWatch User Guide*.

The following table summarizes the metrics and corresponding statistic you can 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> CloudWatch statistic. For example, the <code>Sum</code> value of the <code>ReadBytes</code> metric over a sample period of 5 minutes divided by 300 seconds gives you the throughput as a rate in bytes per second.
Latency	Use the <code>ReadTime</code> and <code>WriteTime</code> metrics with the <code>Average</code> CloudWatch statistic. For example, the <code>Average</code> value of the <code>ReadTime</code> metric 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> CloudWatch statistic. For example, the <code>Samples</code> value of the <code>ReadBytes</code> metric over a sample period of 5 minutes divided by 300 seconds gives you 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.

To measure the data throughput from an application to a volume

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose **Metrics**, then choose the **All metrics** tab and then choose **Storage Gateway**.
3. Choose the **Volume metrics** dimension, and find the volume that you want to work with.
4. Choose the `ReadBytes` and `WriteBytes` metrics.
5. For **Time Range**, choose a value.
6. Choose the `Sum` statistic.
7. For **Period**, choose a value of 5 minutes or greater.
8. In the resulting time-ordered sets of data points (one for `ReadBytes` and one for `WriteBytes`), 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.

For example, if the read throughput is 2,384,199,680 bytes over a period of 300 seconds, then the approximate throughput rate for that datapoint is 7.9 megabytes per second.

To measure the data input/output operations per second from an application to a volume

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose **Metrics**, then choose the **All metrics** tab and then choose **Storage Gateway**.
3. Choose the **Volume metrics** dimension, and find the volume that you want to work with.
4. Choose the `ReadBytes` and `WriteBytes` metrics.
5. For **Time Range**, choose a value.
6. Choose the `Samples` statistic.
7. For **Period**, choose a value of 5 minutes or greater.
8. In the resulting time-ordered sets of data points (one for `ReadBytes` and one for `WriteBytes`), divide each data point by the period (in seconds) to get IOPS.

For example, if the number of write operations is 24,373 over a period of 300 seconds, then the IOPS for that data point is 81 write operations per second.

Measuring Performance Between Your Gateway and AWS

Data throughput, data latency, and operations per second are three measures that you can use to understand how your application storage using the Storage Gateway is performing. These

three values can be measured using the 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 input/output operations per second (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> CloudWatch statistic. For example, the <code>Sum</code> value of the <code>ReadBytes</code> metric over a sample period of 5 minutes divided by 300 seconds gives you the throughput as a rate in bytes per second.
Latency	Use the <code>ReadTime</code> and <code>WriteTime</code> metrics with the <code>Average</code> CloudWatch statistic. For example, the <code>Average</code> value of the <code>ReadTime</code> metric 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> CloudWatch statistic. For example, the <code>Samples</code> value of the <code>ReadBytes</code> metric over a sample period of 5 minutes divided by 300 seconds gives you IOPS.
Throughput to AWS	Use the <code>CloudBytesDownloaded</code> and <code>CloudBytesUploaded</code> metrics with the <code>Sum</code> CloudWatch statistic. For example, the <code>Sum</code> value of the <code>CloudBytesDownloaded</code> metric over a sample period of 5 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.

To measure the upload data throughput from a gateway to AWS

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose **Metrics**, then choose the **All metrics** tab and then choose **Storage Gateway**.
3. Choose the **Gateway metrics** dimension, and find the volume that you want to work with.

4. Choose the `CloudBytesUploaded` metric.
5. For **Time Range**, choose a value.
6. Choose the Sum statistic.
7. For **Period**, choose a value of 5 minutes or greater.
8. 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.

Moving the cursor over a data point displays information about the data point, including its value and bytes uploaded. Divide this value by the **Period** value (5 minutes) to get the throughput at that sample point. For example, if the throughput from the gateway to AWS is 555,544,576 bytes over a period of 300 seconds, then the approximate throughput per second is 1.85 megabytes per second.



To measure the latency per operation of a gateway

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose **Metrics**, then choose the **All metrics** tab and then choose **Storage Gateway**.
3. Choose the **Gateway metrics** dimension, and find the volume that you want to work with.
4. Choose the `ReadTime` and `WriteTime` metrics.
5. For **Time Range**, choose a value.
6. Choose the Average statistic.
7. For **Period**, choose a value of 5 minutes to match the default reporting time.

8. In the resulting time-ordered set of points (one for ReadTime and one for WriteTime), 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. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose **Metrics**, then choose the **All metrics** tab and then choose **Storage Gateway**.
3. Choose the **Gateway metrics** dimension, and find the volume that you want to work with.
4. Choose the CloudDownloadLatency metric.
5. For **Time Range**, choose a value.
6. Choose the Average statistic.
7. For **Period**, choose a value of 5 minutes to match the default reporting time.

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. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose **Alarms**.
3. Choose **Create Alarm** to start the Create Alarm wizard.
4. Choose the **Storage Gateway** dimension, and find the gateway that you want to work with.
5. Choose the CloudBytesUploaded metric.
6. To define the alarm, define the alarm state when the CloudBytesUploaded metric is greater than or equal to a specified value for a specified time. For example, you can define an alarm state when the CloudBytesUploaded metric is greater than 10 MB for 60 minutes.
7. Configure the actions to take for the alarm state. For example, you can have an email notification sent to you.
8. Choose **Create Alarm**.

To set an upper threshold alarm for reading data from AWS

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose **Create Alarm** to start the Create Alarm wizard.

3. Choose the **StorageGateway: Gateway Metrics** dimension, and find the gateway that you want to work with.
4. Choose the `CloudDownloadLatency` metric.
5. Define the alarm by defining the alarm state when the `CloudDownloadLatency` metric is greater than or equal to a specified value for a specified time. For example, you can define an alarm state when the `CloudDownloadLatency` is greater than 60,000 milliseconds for greater than 2 hours.
6. Configure the actions to take for the alarm state. For example, you can have an email notification sent to you.
7. Choose **Create Alarm**.

Understanding volume metrics

You can find information following about the Storage Gateway metrics that cover a volume of a gateway. Each volume of a gateway has a set of metrics associated with it.

Some volume-specific metrics have the same name as certain gateway-specific metrics. These metrics represent the same kinds of measurements but are scoped to the volume instead of the gateway. Before starting work, specify whether you want to work with a gateway metric or a volume metric. Specifically, when working with volume metrics, specify the volume ID for the storage volume that you want to view metrics for. For more information, see [Using Amazon CloudWatch Metrics](#).

Note

Some metrics return data points only when new data has been generated during the most recent monitoring period.

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

Metric	Description	Cached Volumes	Stored Volumes
AvailabilityNotification	The number of availability notifications	Yes	Yes

Metric	Description	Cached Volumes	Stored Volumes
	<p>ions sent by the volume.</p> <p>Units: count</p>		
CacheHitPercent	<p>Percent of applications on read operations from the volume that are served from cache. The sample is taken at the end of the reporting period.</p> <p>When there are no application read operations from the volume, this metric reports 100 percent.</p> <p>Units: Percent</p>	Yes	No

Metric	Description	Cached Volumes	Stored Volumes
CachePercentDirty	<p>The volume's contribution to the overall percentage of the gateway's cache that isn't persisted to AWS. The sample is taken at the end of the reporting period.</p> <p>Use the CachePercentDirty metric of the gateway to view the overall percentage of the gateway's cache that isn't persisted to AWS. For more information, see Understanding gateway metrics.</p> <p>Units: Percent</p>	Yes	Yes

Metric	Description	Cached Volumes	Stored Volumes
CachePercentUsed	<p>The volume's contribution to the overall percent use of the gateway's cache storage. The sample is taken at the end of the reporting period.</p> <p>Use the CachePercentUsed metric of the gateway to view overall percent use of the gateway's cache storage. For more information, see Understanding gateway metrics.</p> <p>Units: Percent</p>	Yes	No
CloudBytesDownloaded	<p>The number of bytes downloaded from the cloud to the volume.</p> <p>Units: Bytes</p>	Yes	Yes
CloudBytesUploaded	<p>The number of bytes uploaded from the cloud to the volume.</p> <p>Units: Bytes</p>	Yes	Yes
HealthNotification	<p>The number of health notifications sent by the volume.</p> <p>Units: count</p>	Yes	Yes

Metric	Description	Cached Volumes	Stored Volumes
IoWaitPercent	<p>The percentage of IoWaitPercent units that are currently used by the volume.</p> <p>Units: Percent</p>	Yes	Yes
MemTotalBytes	<p>The percentage of total memory that is currently used by the volume.</p> <p>Units: Percent</p>	Yes	No
MemoryUsage	<p>The percentage of memory that is currently used by the volume.</p> <p>Units: Percent</p>	Yes	No
ReadBytes	<p>The total number of bytes read from your on-premises applications in the reporting period.</p> <p>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.</p> <p>Units: Bytes</p>	Yes	Yes

Metric	Description	Cached Volumes	Stored Volumes
ReadTime	<p>The total number of milliseconds spent on read operations from your on-premises applications in the reporting period.</p> <p>Use this metric with the Average statistic to measure latency.</p> <p>Units: Milliseconds</p>	Yes	Yes
UserCpuPercent	<p>The percentage of allocated CPU compute units that are currently used by the volume.</p> <p>Units: Percent</p>	Yes	Yes
WriteBytes	<p>The total number of bytes written to your on-premises applications in the reporting period.</p> <p>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.</p> <p>Units: Bytes</p>	Yes	Yes

Metric	Description	Cached Volumes	Stored Volumes
WriteTime	<p>The total number of milliseconds spent on write operations from your on-premises applications in the reporting period.</p> <p>Use this metric with the Average statistic to measure latency.</p> <p>Units: Milliseconds</p>	Yes	Yes
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>	Yes	Yes

Maintaining Your Gateway

Maintaining your Volume Gateway includes tasks such as sizing and configuring local disks for cache storage and upload buffer space, managing updates and setting an update schedule, managing bandwidth usage, and shutting down or deleting you gateway and associated resources if necessary. These tasks are common to all gateway types. If you haven't created a gateway, see [Creating your gateway](#).

Topics

- [Managing local disks for your Storage Gateway](#) - Learn how to assess disk size requirements, add cache capacity, and manage the local disks that you allocate to your Volume Gateway for buffering and storage.
- [Managing Bandwidth for Your Volume Gateway](#) - Learn how to limit the upload throughput from your gateway to AWS to control the amount of network bandwidth the gateway uses.
- [Managing gateway updates](#) - Learn how to turn maintenance updates on or off, and modify the maintenance window schedule for your Volume Gateway.
- [Shutting Down Your Gateway VM](#) - Learn about what to do if you need to shutdown or reboot your gateway virtual machine for maintenance, such as when applying a patch to your hypervisor.
- [Deleting your gateway and removing associated resources](#) - Learn how to delete your gateway using the AWS Storage Gateway console and clean up associated resources to avoid being charged for their continued use.

Managing local disks for your Storage Gateway

The gateway virtual machine (VM) uses the local disks that you allocate on-premises for buffering and storage. Gateways created on Amazon EC2 instances use Amazon EBS volumes as local disks.

Topics

- [Deciding the amount of local disk storage](#)
- [Configuring additional upload buffer or cache storage](#)

Deciding the amount of local disk storage

The number and size of disks that you want to allocate for your gateway is up to you. Depending on the storage solution you deploy, the gateway requires the following additional storage:

- Volume Gateways:
 - Stored gateways require at least one disk to use as an upload buffer.
 - Cached gateways require at least two disks. One to use as a cache, and one to use as an upload buffer.

The following table recommends sizes for local disk storage for your deployed gateway. You can add more local storage later after you set up the gateway, and as your workload demands increase.

Local storage	Description	
Upload buffer	The upload buffer provides a staging area for the data before the gateway uploads the data to Amazon S3. Your gateway uploads this buffer data over an encrypted Secure Sockets Layer (SSL) connection to AWS.	
Cache storage	The cache storage acts as the on-premises durable store for data that is pending upload to Amazon S3 from the upload buffer. When your application performs I/O on a volume or tape, the gateway saves the data to the cache storage for low-latency access. When your application requests data from a volume or tape, the gateway first checks the cache storage for the data before downloading the data from AWS.	

Note

When you provision disks, we strongly recommend that you do not provision local disks for the upload buffer and cache storage if they use the same physical resource (the same disk). Underlying physical storage resources are represented as a data store in VMware. When you deploy the gateway VM, you choose a data store on which to store the VM files. When you provision a local disk (for example, to use as cache storage or upload buffer), you have the option to store the virtual disk in the same data store as the VM or a different data store. If you have more than one data store, we strongly recommend that you choose one data store for the cache storage and another for the upload buffer. A data store that is backed by only one underlying physical disk can lead to poor performance in some situations when it is used to back both the cache storage and upload buffer. This is also true if the backup is a less-performant RAID configuration such as RAID1.

After the initial configuration and deployment of your gateway, you can adjust the local storage by adding or removing disks for an upload buffer. You can also add disks for cache storage.

Determining the size of upload buffer to allocate

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

Note

For Volume Gateways, when the upload buffer reaches its capacity, your volume goes to `PASS THROUGH` status. In this status, new data that your application writes is persisted locally but not uploaded to AWS immediately. Thus, you cannot take new snapshots. When the upload buffer capacity frees up, the volume goes through `BOOTSTRAPPING` status. In this status, any new data that was persisted locally is uploaded to AWS. Finally, the volume returns to `ACTIVE` status. Storage Gateway then resumes normal synchronization of the data stored locally with the copy stored in AWS, and you can start taking new snapshots. For more information about volume status, see [Understanding Volume Statuses and Transitions](#).

To estimate the amount of upload buffer to allocate, you can determine the expected incoming and outgoing data rates and plug them into the following formula.

Rate of incoming data

This rate refers to the application throughput, the rate at which your on-premises applications write data to your gateway over some period of time.

Rate of outgoing data

This rate refers to the network throughput, the rate at which your gateway is able to upload data to AWS. This rate depends on your network speed, utilization, and whether you've activated 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 an effective compression ratio of about 2:1. However, if you are writing videos, the gateway might not be able to achieve any data compression and might require more upload buffer for the gateway.

We strongly recommend that you allocate at least 150 GiB of upload buffer space if either of the following is true:

- Your incoming rate is higher than the outgoing rate.
- The formula returns a value less than 150 GiB.

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

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

Example

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

You can initially use this approximation to determine the disk size that you want to allocate to the gateway as upload buffer space. Add more upload buffer space as needed using the Storage

Gateway console. Also, you can use the Amazon CloudWatch operational metrics to monitor upload buffer usage and determine additional storage requirements. For information on metrics and setting the alarms, see [Monitoring the upload buffer](#).

Determining the size of cache storage to allocate

Your gateway uses its cache storage to provide low-latency access to your recently accessed data. The cache storage acts as the on-premises durable store for data that is pending upload to Amazon S3 from the upload buffer. Generally speaking, you size the cache storage at 1.1 times the upload buffer size. For more information about how to estimate your cache storage size, see [Determining the size of upload buffer to allocate](#).

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

Configuring additional upload buffer or cache storage

As your application needs change, you can increase the gateway's upload buffer or cache storage capacity. You can add storage capacity to your gateway without interrupting functionality or causing downtime. When you add more storage, you do so with the gateway VM turned on.

Important

When adding cache or upload buffer to an existing gateway, you must create new disks on the gateway host hypervisor or Amazon EC2 instance. Do not remove or change the size of existing disks that have already been allocated as cache or upload buffer.

To configure additional upload buffer or cache storage for your gateway

1. Provision one or more new disks on your gateway host hypervisor or Amazon EC2 instance. For information about how to provision a disk on a hypervisor, see your hypervisor's documentation. For information about provisioning Amazon EBS volumes for an Amazon EC2 instance, see [Amazon EBS volumes](#) in the *Amazon Elastic Compute Cloud User Guide for Linux Instances*. In the following steps, you will configure this disk as upload buffer or cache storage.
2. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.

3. In the navigation pane, choose **Gateways**.
4. Search for your gateway and select it from the list.
5. From the **Actions** menu, choose **Configure storage**.
6. In the **Configure storage** section, identify the disks you provisioned. If you don't see your disks, choose the refresh icon to refresh the list. For each disk, choose either **UPLOAD BUFFER** or **CACHE STORAGE** from the **Allocated to** drop-down menu.

 **Note**

UPLOAD BUFFER is the only available option for allocating disks on Stored Volume Gateways.

7. Choose **Save changes** to save your configuration settings.

Managing Bandwidth for Your Volume Gateway

You can limit (or throttle) the upload throughput from the gateway to AWS or the download throughput from AWS to your gateway. Using bandwidth throttling helps you to control the amount of network bandwidth used by your gateway. By default, an activated gateway has no rate limits on upload or download.

You can specify the rate limit by using the AWS Management Console, or programmatically by using either the Storage Gateway API (see [UpdateBandwidthRateLimit](#)) or an AWS Software Development Kit (SDK). By throttling bandwidth programmatically, you can change limits automatically throughout the day—for example, by scheduling tasks to change the bandwidth.

You can also define schedule-based bandwidth throttling for your gateway. You schedule bandwidth throttling by defining one or more bandwidth-rate-limit intervals. For more information, see [Schedule-Based Bandwidth Throttling Using the Storage Gateway Console](#).

Configuring a single setting for bandwidth throttling is the functional equivalent of defining a schedule with a single bandwidth-rate-limit interval set for **Everyday**, with a **Start time** of 00:00 and an **End time** of 23:59.

 **Note**

The information in this section is specific to Tape and Volume Gateways. To manage bandwidth for an Amazon S3 File Gateway, see [Managing Bandwidth for Your Amazon](#)

[S3 File Gateway](#). Bandwidth-rate limits are currently not supported for Amazon FSx File Gateway.

Topics

- [Changing Bandwidth Throttling Using the Storage Gateway Console](#)
- [Schedule-Based Bandwidth Throttling Using the Storage Gateway Console](#)
- [Updating Gateway Bandwidth-Rate Limits Using the AWS SDK for Java](#)
- [Updating Gateway Bandwidth-Rate Limits Using the AWS SDK for .NET](#)
- [Updating Gateway Bandwidth-Rate Limits Using the AWS Tools for Windows PowerShell](#)

Changing Bandwidth Throttling Using the Storage Gateway Console

The following procedure shows how to change a gateway's bandwidth throttling from the Storage Gateway console.

To change a gateway's bandwidth throttling using the console

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the left navigation pane, choose **Gateways**, and then choose the gateway that you want to manage.
3. For **Actions**, choose **Edit bandwidth limit**.
4. In the **Edit rate limits** dialog box, enter new limit values, and then choose **Save**. Your changes appear in the **Details** tab for your gateway.

Schedule-Based Bandwidth Throttling Using the Storage Gateway Console

The following procedure shows how to schedule changes to a gateway's bandwidth throttling using the Storage Gateway console.


To add or modify a schedule for gateway bandwidth throttling

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.

2. In the left navigation pane, choose **Gateways**, and then choose the gateway that you want to manage.
3. For **Actions**, choose **Edit bandwidth rate limit schedule**.

The gateway's bandwidth-rate-limit schedule is displayed in the **Edit bandwidth rate limit schedule** dialog box. By default, a new gateway bandwidth-rate-limit schedule is empty.

4. In the **Edit bandwidth rate limit schedule** dialog box, choose **Add new item** to add a new bandwidth-rate-limit interval. Enter the following information for each bandwidth-rate-limit interval:
 - **Days of week** – You can create the bandwidth-rate-limit interval for weekdays (Monday through Friday), for weekends (Saturday and Sunday), for every day of the week, or for one or more specific days of the week.
 - **Start time** – Enter the start time for the bandwidth interval in the gateway's local timezone, using the HH:MM format.

 **Note**

Your bandwidth-rate-limit interval begins at the start of the minute that you specify here.

- **End time** – Enter the end time for the bandwidth-rate-limit interval in the gateway's local time zone, using the HH:MM format.

 **Important**

The bandwidth-rate-limit interval ends at the end of the minute specified here. To schedule an interval that ends at the end of an hour, enter **59**.

To schedule consecutive continuous intervals, transitioning at the start of the hour, with no interruption between the intervals, enter **59** for the end minute of the first interval. Enter **00** for the start minute of the succeeding interval.

- **Download rate** – Enter the download rate limit, in kilobits per second (Kbps), or select **No limit** to deactivate bandwidth throttling for downloading. The minimum value for the download rate is 100 Kbps.
- **Upload rate** – Enter the upload rate limit, in Kbps, or select **No limit** to deactivate bandwidth throttling for uploading. The minimum value for the upload rate is 50 Kbps.

To modify your bandwidth-rate-limit intervals, you can enter revised values for the interval parameters.

To remove your bandwidth-rate-limit intervals, you can choose **Remove** to the right of the interval to be deleted.

When your changes are complete, choose **Save**.

5. Continue adding bandwidth-rate-limit intervals by choosing **Add new item** and entering the day, the start and end times, and the download and upload rate limits.

Important

Bandwidth-rate-limit intervals cannot overlap. The start time of an interval must occur after the end time of a preceding interval, and before the start time of a following interval.

6. After entering all bandwidth-rate-limit intervals, choose **Save changes** to save your bandwidth-rate-limit schedule.

When the bandwidth-rate-limit schedule is successfully updated, you can see the current download and upload rate limits in the **Details** panel for the gateway.

Updating Gateway Bandwidth-Rate Limits Using the AWS SDK for Java

By updating bandwidth-rate limits programmatically, you can adjust your limits automatically over a period of time—for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth-rate limits using the AWS SDK for Java. To use the example code, you should be familiar with running a Java console application. For more information, see [Getting Started](#) in the *AWS SDK for Java Developer Guide*.

Example : Updating Gateway Bandwidth-Rate Limits Using the AWS SDK for Java

The following Java code example updates a gateway's bandwidth-rate limits. To use this example code, you must provide the service endpoint, your gateway Amazon Resource Name (ARN), and the upload and download limits. For a list of AWS service endpoints that you can use with Storage Gateway, see [AWS Storage Gateway Endpoints and Quotas](#) in the *AWS General Reference*.

```
import java.io.IOException;
```

```
import com.amazonaws.AmazonClientException;
import com.amazonaws.auth.PropertiesCredentials;
import com.amazonaws.services.storagegateway.AWSStorageGatewayClient;
import com.amazonaws.services.storagegateway.model.UpdateBandwidthRateLimitRequest;
import com.amazonaws.services.storagegateway.model.UpdateBandwidthRateLimitResult;

public class UpdateBandwidthExample {

    public static AWSStorageGatewayClient sgClient;

    // The gatewayARN
    public static String gatewayARN = "**** provide gateway ARN ****";

    // The endpoint
    static String serviceURL = "https://storagegateway.us-east-1.amazonaws.com";

    // Rates
    static long uploadRate = 51200; // Bits per second, minimum 51200
    static long downloadRate = 102400; // Bits per second, minimum 102400

    public static void main(String[] args) throws IOException {

        // Create a Storage Gateway client
        sgClient = new AWSStorageGatewayClient(new PropertiesCredentials(
UpdateBandwidthExample.class.getResourceAsStream("AwsCredentials.properties")));
        sgClient.setEndpoint(serviceURL);

        UpdateBandwidth(gatewayARN, uploadRate, downloadRate);

    }

    private static void UpdateBandwidth(String gatewayARN2, long uploadRate2,
        long downloadRate2) {
        try
        {
            UpdateBandwidthRateLimitRequest updateBandwidthRateLimitRequest =
                new UpdateBandwidthRateLimitRequest()
                    .withGatewayARN(gatewayARN)
                    .withAverageDownloadRateLimitInBitsPerSec(downloadRate)
                    .withAverageUploadRateLimitInBitsPerSec(uploadRate);
```

```
UpdateBandwidthRateLimitResult updateBandwidthRateLimitResult =
sgClient.updateBandwidthRateLimit(updateBandwidthRateLimitRequest);
String returnGatewayARN = updateBandwidthRateLimitResult.getGatewayARN();
System.out.println("Updated the bandwidth rate limits of " +
returnGatewayARN);
System.out.println("Upload bandwidth limit = " + uploadRate + " bits per
second");
System.out.println("Download bandwidth limit = " + downloadRate + " bits
per second");
}
catch (AmazonClientException ex)
{
System.err.println("Error updating gateway bandwidth.\n" + ex.toString());
}
}
}
```

Updating Gateway Bandwidth-Rate Limits Using the AWS SDK for .NET

By updating bandwidth-rate limits programmatically, you can adjust your limits automatically over a period of time—for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth-rate limits by using the AWS SDK for .NET. To use the example code, you should be familiar with running a .NET console application. For more information, see [Getting Started](#) in the *AWS SDK for .NET Developer Guide*.

Example : Updating Gateway Bandwidth-Rate Limits by Using the AWS SDK for .NET

The following C# code example updates a gateway's bandwidth-rate limits. To use this example code, you must provide the service endpoint, your gateway Amazon Resource Name (ARN), and the upload and download limits. For a list of AWS service endpoints that you can use with Storage Gateway, see [AWS Storage Gateway Endpoints and Quotas](#) in the *AWS General Reference*.

```
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using Amazon.StorageGateway;
using Amazon.StorageGateway.Model;

namespace AWSStorageGateway
{
    class UpdateBandwidthExample
```

```
{
    static AmazonStorageGatewayClient sgClient;
    static AmazonStorageGatewayConfig sgConfig;

    // The gatewayARN
    public static String gatewayARN = "**** provide gateway ARN ****";

    // The endpoint
    static String serviceURL = "https://storagegateway.us-east-1.amazonaws.com";

    // Rates
    static long uploadRate = 51200; // Bits per second, minimum 51200
    static long downloadRate = 102400; // Bits per second, minimum 102400

    public static void Main(string[] args)
    {
        // Create a Storage Gateway client
        sgConfig = new AmazonStorageGatewayConfig();
        sgConfig.ServiceURL = serviceURL;
        sgClient = new AmazonStorageGatewayClient(sgConfig);

        UpdateBandwidth(gatewayARN, uploadRate, downloadRate);

        Console.WriteLine("\nTo continue, press Enter.");
        Console.Read();
    }

    public static void UpdateBandwidth(string gatewayARN, long uploadRate, long
downloadRate)
    {
        try
        {
            UpdateBandwidthRateLimitRequest updateBandwidthRateLimitRequest =
                new UpdateBandwidthRateLimitRequest()
                    .WithGatewayARN(gatewayARN)
                    .WithAverageDownloadRateLimitInBitsPerSec(downloadRate)
                    .WithAverageUploadRateLimitInBitsPerSec(uploadRate);

            UpdateBandwidthRateLimitResponse updateBandwidthRateLimitResponse =
sgClient.UpdateBandwidthRateLimit(updateBandwidthRateLimitRequest);
            String returnGatewayARN =
updateBandwidthRateLimitResponse.UpdateBandwidthRateLimitResult.GatewayARN;
            Console.WriteLine("Updated the bandwidth rate limits of " +
returnGatewayARN);
        }
    }
}
```

```
        Console.WriteLine("Upload bandwidth limit = " + uploadRate + " bits per
second");
        Console.WriteLine("Download bandwidth limit = " + downloadRate + " bits
per second");
    }
    catch (AmazonStorageGatewayException ex)
    {
        Console.WriteLine("Error updating gateway bandwidth.\n" +
ex.ToString());
    }
}
}
```

Updating Gateway Bandwidth-Rate Limits Using the AWS Tools for Windows PowerShell

By updating bandwidth-rate limits programmatically, you can adjust limits automatically over a period of time—for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth-rate limits using the AWS Tools for Windows PowerShell. To use the example code, you should be familiar with running a PowerShell script. For more information, see [Getting Started](#) in the *AWS Tools for Windows PowerShell User Guide*.

Example : Updating Gateway Bandwidth-Rate Limits by Using the AWS Tools for Windows PowerShell

The following PowerShell script example updates a gateway's bandwidth-rate limits. To use this example script, you must provide your gateway Amazon Resource Name (ARN), and the upload and download limits.

```
<#
.DESCRIPTION
    Update Gateway bandwidth limits.

.NOTES
    PREREQUISITES:
    1) AWS Tools for PowerShell from https://aws.amazon.com/powershell/
    2) Credentials and region stored in session using Initialize-AWSDefault.
    For more info, see https://docs.aws.amazon.com/powershell/latest/userguide/
specifying-your-aws-credentials.html
```



```
.EXAMPLE
powershell.exe .\SG_UpdateBandwidth.ps1
#>

$UploadBandwidthRate = 51200
$DownloadBandwidthRate = 102400
$gatewayARN = "**** provide gateway ARN ****"

#Update Bandwidth Rate Limits
Update-SGBandwidthRateLimit -GatewayARN $gatewayARN `
                             -AverageUploadRateLimitInBitsPerSec $UploadBandwidthRate `
                             -AverageDownloadRateLimitInBitsPerSec
                             $DownloadBandwidthRate

$limits = Get-SGBandwidthRateLimit -GatewayARN $gatewayARN

Write-Output("`nGateway: " + $gatewayARN);
Write-Output("`nNew Upload Rate: " + $limits.AverageUploadRateLimitInBitsPerSec)
Write-Output("`nNew Download Rate: " + $limits.AverageDownloadRateLimitInBitsPerSec)
```

Managing gateway updates

Storage Gateway consists of a managed cloud services component and a gateway appliance component that you deploy either on-premises, or on an Amazon EC2 instance in the AWS cloud. Both components receive regular updates. The topics in this section describe the cadence of these updates, how they are applied, and how to configure update-related settings on the gateways in your deployment.

Important

You should treat the Storage Gateway appliance as a managed virtual machine, and should not attempt to access or modify its installation in any way. Attempting to install or update any software packages using methods other than the normal AWS gateway update mechanism (for example, SSM or hypervisor tools) may cause the gateway to malfunction.

Update frequency and expected behavior

AWS updates the cloud services component as needed without causing disruption to deployed gateways. Your deployed gateway appliances receive monthly maintenance updates. Monthly maintenance updates can include operating system and software upgrades, fixes to address stability, performance, and security, and access to new features. All updates are cumulative, and upgrade gateways to the current version when applied. For information about the specific changes included in each update, see [Release Notes for Volume Gateway Appliance Software](#).

Monthly maintenance updates may cause a brief disruption of service. The gateway's VM host doesn't need to reboot during updates, but the gateway will be unavailable for a short period while the gateway appliance updates and restarts.

When you deploy and activate your gateway, a default weekly maintenance window schedule is set. You can modify the maintenance window schedule at any time. You can also turn off monthly maintenance updates, but we recommend leaving them turned on.

Note

Urgent updates will sometimes be applied according to the maintenance window schedule, even if regular maintenance updates are turned off.

Before any update is applied to your gateway, AWS notifies you with a message on the Storage Gateway console and your AWS Health Dashboard. For more information, see [AWS Health Dashboard](#). To modify the email address where software update notifications are sent, see [Update the alternate contacts for your AWS account](#) in the *AWS Account Management Reference Guide*.

When updates are available, the gateway **Details** tab displays a maintenance message. You can also see the date and time that the last successful update was applied on the **Details** tab.

Turn maintenance updates on or off

When maintenance updates are turned on, your gateway automatically applies these updates according to the configured maintenance window schedule. For more information, see .

If maintenance updates are turned off, the gateway will not apply these updates automatically, but you can always apply them manually using the Storage Gateway console, API, or CLI. Urgent updates will sometimes be applied during your configured maintenance window, regardless of this setting.

Note

The following procedure describes how to turn gateway updates on or off using the Storage Gateway console. To change this setting programmatically using the API, see [UpdateMaintenanceStartTime](#) in the *Storage Gateway API Reference*.

To turn maintenance updates on or off using the Storage Gateway console:

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. On the navigation pane, choose **Gateways**, and then choose the gateway for which you want to configure maintenance updates.
3. Choose **Actions**, and then choose **Edit maintenance settings**.
4. For **Maintenance updates**, select **On** or **Off**.
5. Choose **Save changes** when finished.

You can verify the updated setting on the **Details** tab for the selected gateway in the Storage Gateway console.

Modify the gateway maintenance window schedule

If maintenance updates are turned on, your gateway automatically applies these updates according the maintenance window schedule. Urgent updates will sometimes be applied during your configured maintenance window, regardless of the maintenance updates setting.

Note


The following procedure describes how to modify the maintenance window schedule using the Storage Gateway console. To change this setting programmatically using the API, see [UpdateMaintenanceStartTime](#) in the *Storage Gateway API Reference*.

To modify the maintenance window schedule using the Storage Gateway console:

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.

2. On the navigation pane, choose **Gateways**, and then choose the gateway for which you want to configure maintenance updates.
3. Choose **Actions**, and then choose **Edit maintenance settings**.
4. Under **Maintenance window start time**, do the following:
 - a. For **Schedule**, choose **Weekly** or **Monthly** to set the maintenance window cadence.
 - b. If you choose **Weekly**, modify the values for **Day of the week** and **Time** to set the specific point during each week when the maintenance window will begin.

If you choose **Monthly**, modify the values for **Day of the month** and **Time** to set the specific point during each month when the maintenance window will begin.

 **Note**

The maximum value that can be set for day of the month is 28. It is not possible to set the maintenance schedule to start on days 29 through 31.

If you receive an error while configuring this setting, it might mean that your gateway software is out of date. Considering updating your gateway manually first, and then attempt to configure the maintenance window schedule again.

5. Choose **Save changes** when finished.

You can verify the updated settings on the **Details** tab for the selected gateway in the Storage Gateway console.

Apply an update manually

If a software update is available for your gateway, you can apply it manually by following the procedure below. This manual update process ignores the maintenance window schedule and applies the update immediately, even if maintenance updates are turned off.

 **Note**

The following procedure describes how to manually apply an update using the Storage Gateway console. To perform this action programmatically using the API, see [UpdateGatewaySoftwareNow](#) in the *Storage Gateway API Reference*.

To apply a gateway software update manually using the Storage Gateway console:

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. On the navigation pane, choose **Gateways**, and then choose the gateway you want to update.

If an update is available, the console displays a blue notification banner on the gateway **Details** tab, which includes an option to apply the update.

3. Choose **Apply update now** to immediately update the gateway.

Note

This operation causes a temporary disruption to gateway functionality while the update installs. During this time, the gateway status appears **OFFLINE** in the Storage Gateway console. After the update finishes installing, the gateway resumes normal operation and its status changes to **RUNNING**.

You can verify that the gateway software was updated to the latest version by checking the **Details** tab for the selected gateway in the Storage Gateway console.

Shutting Down Your Gateway VM

You might need to shutdown or reboot your VM for maintenance, such as when applying a patch to your hypervisor. Before you shutdown the VM, you must first stop the gateway. Although this section focuses on starting and stopping your gateway using the Storage Gateway Management Console, you can also and stop your gateway by using your VM local console or Storage Gateway API. When you power on your VM, remember to restart your gateway.

Important

If you stop and start an Amazon EC2 gateway that uses ephemeral storage, the gateway will be permanently offline. This happens because the physical storage disk is replaced. There is no work-around for this issue. The only resolution is to delete the gateway and activate a new one on a new EC2 instance.

Note

If you stop your gateway while your backup software is writing or reading from a tape, the write or read task might not succeed. Before you stop your gateway, you should check your backup software and the backup schedule for any tasks in progress.

- Gateway VM local console—see [Logging in to the Volume Gateway local console](#).
- Storage Gateway API—see [ShutdownGateway](#)

Starting and Stopping a Volume Gateway

To stop a Volume Gateway

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose **Gateways**, and then choose the gateway to stop. The status of the gateway is **Running**.
3. For **Actions**, choose **Stop gateway** and verify the id of the gateway from the dialog box, and then choose **Stop gateway**.

While the gateway is stopping, you might see a message that indicates the status of the gateway. When the gateway shuts down, a message and a **Start gateway** button appears in the **Details** tab.

When you stop your gateway, the storage resources will not be accessible until you start your storage. If the gateway was uploading data when it was stopped, the upload will resume when you start the gateway.

To start a Volume Gateway

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose **Gateways** and then choose the gateway to start. The status of the gateway is **Shutdown**.
3. Choose **Details**. and then choose **Start gateway**.

Deleting your gateway and removing associated resources

If you don't plan to continue using your gateway, consider deleting the gateway and its associated resources. Removing resources avoids incurring charges for resources you don't plan to continue using and helps reduce your monthly bill.

When you delete a gateway, it no longer appears on the AWS Storage Gateway Management Console and its iSCSI connection to the initiator is closed. The procedure for deleting a gateway is the same for all gateway types; however, depending on the type of gateway you want to delete and the host it is deployed on, you follow specific instructions to remove associated resources.

You can delete a gateway using the Storage Gateway console or programmatically. You can find information following about how to delete a gateway using the Storage Gateway console. If you want to programmatically delete your gateway, see [AWS Storage Gateway API Reference](#).

Topics

- [Deleting Your Gateway by Using the Storage Gateway Console](#)
- [Removing Resources from a Gateway Deployed On-Premises](#)
- [Removing Resources from a Gateway Deployed on an Amazon EC2 Instance](#)

Deleting Your Gateway by Using the Storage Gateway Console

The procedure for deleting a gateway is the same for all gateway types. However, depending on the type of gateway you want to delete and the host the gateway is deployed on, you might have to perform additional tasks to remove resources associated with the gateway. Removing these resources helps you avoid paying for resources you don't plan to use.

Note

For gateways deployed on an Amazon EC2 instance, the instance continues to exist until you delete it.

For gateways deployed on a virtual machine (VM), after you delete your gateway the gateway VM still exists in your virtualization environment. To remove the VM, use the VMware vSphere client, Microsoft Hyper-V Manager, or Linux Kernel-based Virtual Machine (KVM) client to connect to the host and remove the VM. Note that you can't reuse the deleted gateway's VM to activate a new gateway.

To delete a gateway

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **Gateways**, then select one or more gateways to delete.
3. For **Actions**, choose **Delete gateway**. The confirmation dialog box appears.

Warning

Before you do this step, make sure that there are no applications currently writing to the gateway's volumes. If you delete the gateway while it is in use, data loss can occur. When a gateway is deleted, there is no way to get it back.

4. Verify that you want to delete the specified gateways, then type the word *delete* in the confirmation box, and choose **Delete**.
5. (Optional) If you want to provide feedback about your deleted gateway, complete the feedback dialog box, then choose **Submit**. Otherwise, choose **Skip**.

Important

You no longer pay software charges after you delete a gateway, but resources such as virtual tapes, Amazon Elastic Block Store (Amazon EBS) snapshots, and Amazon EC2 instances persist. You will continue to be billed for these resources. You can choose to remove Amazon EC2 instances and Amazon EBS snapshots by canceling your Amazon EC2 subscription. If you want to keep your Amazon EC2 subscription, you can delete your Amazon EBS snapshots using the Amazon EC2 console.

Removing Resources from a Gateway Deployed On-Premises

You can use the instructions following to remove resources from a gateway that is deployed on-premises.

Removing Resources from a Volume Gateway Deployed on a VM

If the gateway you want to delete are deployed on a virtual machine (VM), we suggest that you take the following actions to clean up resources:

- Delete the gateway. For instructions, see [Deleting Your Gateway by Using the Storage Gateway Console](#).
- Delete all Amazon EBS snapshots you don't need. For instructions, see [Deleting an Amazon EBS Snapshot](#) in the *Amazon EC2 User Guide*.

Removing Resources from a Gateway Deployed on an Amazon EC2 Instance

If you want to delete a gateway that you deployed on an Amazon EC2 instance, we recommend that you clean up the AWS resources that were used with the gateway, specifically the Amazon EC2 instance, any Amazon EBS volumes, and also tapes if you deployed a Tape Gateway. Doing so helps avoid unintended usage charges.

Removing Resources from Your Cached Volumes Deployed on Amazon EC2

If you deployed a gateway with cached volumes on EC2, we suggest that you take the following actions to delete your gateway and clean up its resources:

1. In the Storage Gateway console, delete the gateway as shown in [Deleting Your Gateway by Using the Storage Gateway Console](#).
2. In the Amazon EC2 console, stop your EC2 instance if you plan on using the instance again. Otherwise, terminate the instance. If you plan on deleting volumes, make note of the block devices that are attached to the instance and the devices' identifiers before terminating the instance. You will need these to identify the volumes you want to delete.
3. In the Amazon EC2 console, remove all Amazon EBS volumes that are attached to the instance if you don't plan on using them again. For more information, see [Clean Up Your Instance and Volume](#) in the *Amazon EC2 User Guide*.

Performing maintenance tasks using the local console

This section contains the following topics, which provide information about how to perform maintenance tasks using the gateway appliance local console. The local console runs directly on the virtualization host platform that hosts your gateway appliance. For on-premises gateways, you access the local console through your VMware, Hyper-v, or Linux KVM virtualization host. For Amazon EC2 gateways, you access the console by connecting to the Amazon EC2 instance using SSH. Most of the tasks are common across the different host platforms, but there are also some differences.

Topics

- [Accessing the Gateway Local Console](#) - Learn how to log into the local console for an on-premises gateway hosted on a Linux Kernel-based Virtual Machine (KVM), VMware ESXi, or Microsoft Hyper-V Manager platform.
- [Performing Tasks on the VM Local Console](#) - Learn how to use the local console to perform basic setup and advanced configuration tasks for an on-premises gateway, such as configuring an HTTP proxy, viewing system resource status, or running terminal commands.
- [Performing Tasks on the Amazon EC2 Local Console](#) - Learn how to log into the local console to perform basic setup and advanced configuration tasks for an Amazon EC2 gateway, such as configuring an HTTP proxy, viewing system resource status, or running terminal commands.

Accessing the Gateway Local Console

How you access your VM's local console depends on the type of the Hypervisor you deployed your gateway VM on. In this section, you can find information on how to access the VM local console using Linux Kernel-based Virtual Machine (KVM), VMware ESXi, and Microsoft Hyper-V Manager.

Topics

- [Accessing the Gateway Local Console with Linux KVM](#)
- [Accessing the Gateway Local Console with VMware ESXi](#)
- [Access the Gateway Local Console with Microsoft Hyper-V](#)

Accessing the Gateway Local Console with Linux KVM

There are different ways to configure virtual machines running on KVM, depending on the Linux distribution being used. Instructions for accessing KVM configuration options from the command line follow. Instructions might differ depending on your KVM implementation.

To access your gateway's local console with KVM

1. Use the following command to list the VMs that are currently available in KVM.

```
# virsh list
```

The command returns a list of VMs with **Id**, **Name**, and **State** information for each. Note the Id of the VM for which you want to launch the gateway local console.

2. Use the following command to access the local console.

```
# virsh console Id
```

Replace *Id* with the **Id** of the VM you noted in the previous step.

The AWS Appliance gateway local console prompts you to login to change your network configuration and other settings.

3. Enter your username and password to log into the gateway local console. For more information, see [Logging in to the Volume Gateway local console](#).

After you log in, the **AWS Appliance Activation - Configuration** menu appears. You can select from the menu options to perform gateway configuration tasks. For more information, see [Performing tasks on the virtual machine local console](#).

Accessing the Gateway Local Console with VMware ESXi

To access your gateway's local console with VMware ESXi

1. In the VMware vSphere client, select your gateway VM.
2. Make sure that the gateway VM is turned on.

Note

If your gateway VM is turned on, a green arrow icon appears with the VM icon in the VM browser panel on the left side of the application window. If your gateway VM is not turned on, you can turn it on by choosing the green **Power On** icon on the **Toolbar** at the top of the application window.

3. Choose the **Console** tab in the main information panel on the right side of the application window.

After a few moments, the AWS Appliance gateway local console prompts you to login to change your network configuration and other settings.

Note

To release the cursor from the console window, press **Ctrl+Alt**.

4. Enter your username and password to log into the gateway local console. For more information, see [Logging in to the Volume Gateway local console](#).

After you log in, the **AWS Appliance Activation - Configuration** menu appears. You can select from the menu options to perform gateway configuration tasks. For more information, see [Performing tasks on the virtual machine local console](#).

Access the Gateway Local Console with Microsoft Hyper-V

To access your gateway's local console (Microsoft Hyper-V)

1. Select your gateway appliance VM from the **Virtual Machines** panel on the left side of the Microsoft Hyper-V Manager application window.
2. Make sure that the gateway is turned on.

Note

If your gateway VM is turned on, **Running** is displayed in the **State** column for the VM in the **Virtual Machines** panel on the left side of the application window. If your

gateway VM is not turned on, you can turn it on by choosing **Start** in the **Actions** panel on the right side of the application window.

3. Choose **Connect** from the **Actions** panel.

The **Virtual Machine Connection** window appears. If an authentication window appears, type the sign-in credentials provided to you by the hypervisor administrator.

After a few moments, the AWS Appliance gateway local console prompts you to login to change your network configuration and other settings.

4. Enter your username and password to log into the gateway local console. For more information, see [Logging in to the Volume Gateway local console](#).

After you log in, the **AWS Appliance Activation - Configuration** menu appears. You can select from the menu options to perform gateway configuration tasks. For more information, see [Performing tasks on the virtual machine local console](#).

Performing Tasks on the VM Local Console

For a Volume Gateway that you deploy on-premises, you can perform the following maintenance tasks using the gateway local console that you access from your virtual machine host platform. These tasks are common to VMware, Microsoft Hyper-V, and Linux Kernel-based Virtual Machine (KVM) hypervisors.

Topics

- [Logging in to the Volume Gateway local console](#) - Learn about how to log in to the gateway local console where you can configure gateway network settings and change the default password.
- [Configuring a SOCKS5 proxy for your on-premises gateway](#) - Learn about how you can configure Storage Gateway to route all AWS endpoint traffic through a Socket Secure version 5 (SOCKS5) proxy server.
- [Configuring Your Gateway Network](#) - Learn about how you can configure your gateway to use DHCP or assign a static IP address.
- [Testing your gateway connection to the internet](#) - Learn about how you can use the gateway local console to test the connection between the gateway and the internet.

- [Running storage gateway commands in the local console for an on-premises gateway](#) - Learn about how to run local console commands that allow you to perform additional tasks such as saving routing tables, connecting to AWS Support, and more.
- [Viewing your gateway system resource status](#) - Learn about how to check the virtual CPU cores, root volume size, and RAM that are available to your gateway appliance.

Logging in to the Volume Gateway local console

When the VM is ready for you to log in, the login screen is displayed. If this is your first time logging in to the local console, you use the default sign-in credentials to log in. These default login credentials give you access to menus where you can configure gateway network settings and change the password from the local console. Storage Gateway allows you to set your own password from the AWS Storage Gateway console instead of changing the password from the local console. You don't need to know the default password to set a new password. For more information, see [Setting the Local Console Password from the Storage Gateway Console](#).

To log in to the gateway's local console

- If this is your first time logging in to the local console, log in to the VM with the default credentials. The default user name is `admin` and the password is `password`.

Otherwise, use your credentials to log in.

Note

We recommend changing the default password by entering the corresponding numeral for **Gateway Console** from the **AWS Appliance Activation - Configuration** main menu, then running the `passwd` command. For information about how to run the command, see [Running storage gateway commands in the local console for an on-premises gateway](#). You can also set your own password from the AWS Storage Gateway console. For more information, see [Setting the Local Console Password from the Storage Gateway Console](#).

⚠ Important

For older versions of the volume or Tape Gateway, the user name is `sguser` and the password is `sgpassword`. If you reset your password and your gateway is updated to a newer version, your the user name will change to `admin` but the password will be maintained.

Setting the Local Console Password from the Storage Gateway Console

When you log in to the local console for the first time, you log in to the VM with the default credentials— The user name is `admin` and the password is `password`. We recommend that you always set a new password immediately after you create your new gateway. You can set this password from the AWS Storage Gateway console rather than the local console if you want. You don't need to know the default password to set a new password.

To set the local console password on the Storage Gateway console

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. On the navigation pane, choose **Gateways** then choose the gateway for which you want to set a new password.
3. For **Actions**, choose **Set Local Console Password**.
4. In the **Set Local Console Password** dialog box, type a new password, confirm the password and then choose **Save**. Your new password replaces the default password. Storage Gateway does not save the password but rather safely transmits it to the VM.

ℹ Note

The password can consist of any character on the keyboard and can be 1 to 512 characters long.

Configuring a SOCKS5 proxy for your on-premises gateway

Volume Gateways and Tape Gateways support configuration of a Socket Secure version 5 (SOCKS5) proxy between your on-premises gateway and AWS.

Note

The only supported proxy configuration is SOCKS5.

If your gateway must use a proxy server to communicate to the internet, then you need to configure the SOCKS proxy settings for your gateway. You do this by specifying an IP address and port number for the host running your proxy. After you do so, Storage Gateway routes all traffic through your proxy server. For information about network requirements for your gateway, see [Network and firewall requirements](#).

The following procedure shows you how to configure SOCKS proxy for Volume Gateway and Tape Gateway.

To configure a SOCKS5 proxy for volume and Tape Gateways

- Log in to your gateway's local console.
 - VMware ESXi – for more information, see [Accessing the Gateway Local Console with VMware ESXi](#).
 - Microsoft Hyper-V – for more information, see [Access the Gateway Local Console with Microsoft Hyper-V](#).
 - KVM – for more information, see [Accessing the Gateway Local Console with Linux KVM](#).
- From the **AWS Storage Gateway - Configuration** main menu, enter the corresponding numeral to select **SOCKS Proxy Configuration**.
- From the **AWS Storage Gateway SOCKS Proxy Configuration** menu, enter the corresponding numeral to perform one of the following tasks:

To Perform This Task	Do This
Configure a SOCKS proxy	Enter the corresponding numeral to select Configure SOCKS Proxy .

To Perform This Task	Do This
	You will need to supply a host name and port to complete configuration.
View the current SOCKS proxy configuration	<p>Enter the corresponding numeral to select View Current SOCKS Proxy Configuration.</p> <p>If a SOCKS proxy is not configured, the message <code>SOCKS Proxy not configured</code> is displayed. If a SOCKS proxy is configured, the host name and port of the proxy are displayed.</p>
Remove a SOCKS proxy configuration	<p>Enter the corresponding numeral to select Remove SOCKS Proxy Configuration.</p> <p>The message <code>SOCKS Proxy Configuration Removed</code> is displayed.</p>

- Restart your VM to apply your HTTP configuration.

Configuring Your Gateway Network


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 might need to manually assign your gateway's IP as a static IP address, as described following.


To configure your gateway to use static IP addresses

- Log in to your gateway's local console.
 - VMware ESXi – for more information, see [Accessing the Gateway Local Console with VMware ESXi](#).
 - Microsoft Hyper-V – for more information, see [Access the Gateway Local Console with Microsoft Hyper-V](#).
 - KVM – for more information, see [Accessing the Gateway Local Console with Linux KVM](#).

2. From the **AWS Storage Gateway - Configuration** main menu, enter the corresponding numeral to select **Network Configuration**.
3. From the **AWS Storage Gateway Network Configuration** menu, perform one of the following tasks:

To Perform This Task	Do This
Describe network adapter	<p>Enter the corresponding numeral to select Describe Adapter.</p> <p>A list of adapter names appears, and you are prompted to type 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"> • Media access control (MAC) address • IP address • Netmask • Gateway IP address • DHCP activated status <p>You use the adapter names listed here when you configure a static IP address or set your gateway's default adapter.</p>
Configure DHCP	<p>Enter the corresponding numeral to select Configure DHCP.</p> <p>You are prompted to configure network interface to use DHCP.</p>

To Perform This Task	Do This
Configure a static IP address for your gateway	<p data-bbox="829 260 1442 340">Enter the corresponding numeral to select Configure Static IP.</p> <p data-bbox="829 388 1404 468">You are prompted to type the following information to configure a static IP:</p> <ul data-bbox="829 520 1396 1073" style="list-style-type: none"><li data-bbox="829 520 1198 579">• Network adapter name<li data-bbox="829 611 1011 669">• IP address<li data-bbox="829 701 992 760">• Netmask<li data-bbox="829 791 1219 850">• Default gateway address<li data-bbox="829 882 1396 982">• Primary Domain Name Service (DNS) address<li data-bbox="829 1014 1206 1073">• Secondary DNS address <div data-bbox="829 1209 1510 1619" style="border: 1px solid #f08080; border-radius: 10px; padding: 10px;"><p data-bbox="857 1247 1068 1285"> Important</p><p data-bbox="906 1304 1469 1581">If your gateway has already been activated, you must shut it down and restart it from the Storage Gateway console for the settings to take effect. For more information, see Shutting Down Your Gateway VM.</p></div> <p data-bbox="829 1719 1507 1852">If your gateway uses more than one network interface, you must set all activated interfaces to use DHCP or static IP addresses.</p>

To Perform This Task	Do This
	<p>For example, suppose your gateway VM uses two interfaces configured as DHCP. If you later set one interface to a static IP, the other interface is deactivated. To activate the interface in this case, you must set it to a static IP.</p> <p>If both interfaces are initially set to use static IP addresses and you then set the gateway to use DHCP, both interfaces will use DHCP.</p>
Configure a hostname for your gateway	<p>Enter the corresponding numeral to select Configure Hostname.</p> <p>You are prompted to choose whether the gateway will use a static hostname that you specify, or acquire one automatically through DHCP or rDNS.</p> <p>If you select Static, you are prompted to provide a static hostname, such as <code>testgateway.example.com</code>. Enter <code>y</code> to apply the configuration.</p> <div data-bbox="829 1276 1507 1734" style="border: 1px solid #00a0e3; border-radius: 10px; padding: 10px;"><p> Note</p><p>If you configure a static hostname for your gateway, ensure that the provided hostname is in the domain that gateway is joined to. You must also create an A record in your DNS system that points the gateway's IP address to its static hostname.</p></div>

To Perform This Task	Do This
Reset all your gateway's network configuration to DHCP	<p>Enter the corresponding numeral to select Reset all to DHCP.</p> <p>All network interfaces are set to use DHCP.</p> <div data-bbox="829 493 1507 905" style="border: 1px solid #f08080; border-radius: 10px; padding: 10px;"><p>⚠ Important</p><p>If your gateway has already been activated, you must shut down and restart your gateway from the Storage Gateway console for the settings to take effect. For more information, see Shutting Down Your Gateway VM.</p></div>
Set your gateway's default route adapter	<p>Enter the corresponding numeral to select Set Default Adapter.</p> <p>The available adapters for your gateway are shown, and you are prompted to select one of the adapters—for example, eth0.</p>
View your gateway's DNS configuration	<p>Enter the corresponding numeral to select View DNS Configuration.</p> <p>The IP addresses of the primary and secondary DNS name servers are displayed.</p>
View routing tables	<p>Enter the corresponding numeral to select View Routes.</p> <p>The default route of your gateway is displayed.</p>

Testing your gateway connection to the internet

You can use your gateway's local console to test your internet connection. This test can be useful when you are troubleshooting network issues with your gateway.

To test your gateway's connection to the internet

1. Log in to your gateway's local console.
 - VMware ESXi – for more information, see [Accessing the Gateway Local Console with VMware ESXi](#).
 - Microsoft Hyper-V – for more information, see [Access the Gateway Local Console with Microsoft Hyper-V](#).
 - KVM – for more information, see [Accessing the Gateway Local Console with Linux KVM](#).
2. From the **AWS Storage Gateway - Configuration** main menu, enter the corresponding numeral to select **Test Network Connectivity**.

If your gateway has already been activated, the connectivity test begins immediately. For gateways that have not yet been activated, you must specify the endpoint type and AWS Region as described in the following steps.

3. If your gateway is not yet activated, enter the corresponding numeral to select the endpoint type for your gateway.
4. If you selected the public endpoint type, enter the corresponding numeral to select the AWS Region that you want to test. For supported AWS Regions and a list of AWS service endpoints you can use with Storage Gateway, see [AWS Storage Gateway endpoints and quotas](#) in the *AWS General Reference*.

As the test progresses, each endpoint displays either **[PASSED]** or **[FAILED]**, indicating the status of the connection as follows:

Message	Description
[PASSED]	Storage Gateway has network connectivity.
[FAILED]	Storage Gateway does not have network connectivity.

Running storage gateway commands in the local console for an on-premises gateway


The VM local console in Storage Gateway helps provide a secure environment for configuring and diagnosing issues with your gateway. Using the local console commands, you can perform maintenance tasks such as saving routing tables, connecting to AWS Support, and so on.

To run a configuration or diagnostic command


1. Log in to your gateway's local console:
 - For more information on logging in to the VMware ESXi local console, see [Accessing the Gateway Local Console with VMware ESXi](#).
 - For more information on logging in to the Microsoft Hyper-V local console, see [Access the Gateway Local Console with Microsoft Hyper-V](#).
 - For more information on logging in to the KVM local console, see [Accessing the Gateway Local Console with Linux KVM](#).
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **Gateway Console**.
3. From the gateway console command prompt, enter **h**.

The console displays the **AVAILABLE COMMANDS** menu, which lists the available commands:

Command	Function
dig	Collect output from dig for DNS troubleshooting.
exit	Return to Configuration menu.
h	Display available command list.
ifconfig	View or configure network interfaces.

 **Note**

We recommend configuring network or IP settings using the Storage

Command	Function
	<p>Gateway console or the dedicated local console menu option. For instructions, see Configuring Your Gateway Network.</p>
ip	<p>Show / manipulate routing, devices, and tunnels.</p> <div data-bbox="834 594 1507 1003" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p> Note</p> <p>We recommend configuring network or IP settings using the Storage Gateway console or the dedicated local console menu option. For instructions, see Configuring Your Gateway Network.</p> </div>
iptables	Administration tool for IPv4 packet filtering and NAT.
ncport	Test connectivity to a specific TCP port on a network.
nping	Collect output from nping for network troubleshooting.
open-support-channel	Connect to AWS Support.
passwd	Update authentication tokens.
save-iptables	Persist IP tables.
save-routing-table	Save newly added routing table entry.

Command	Function
sslcheck	Returns output with certificate issuer <div data-bbox="834 302 1508 905" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>Storage Gateway uses certificate issuer verification and does not support ssl inspection. If this command returns an issuer other than <code>aws-appliance@amazon.com</code>, then it is likely that an application performing an ssl inspection. In that case, we recommend bypassing ssl inspection for the Storage Gateway appliance.</p> </div>
tcptraceroute	Collect traceroute output on TCP traffic to a destination.

- From the gateway console command prompt, enter the corresponding command for the function you want to use, and follow the instructions.

To learn about a command, enter `man + command name` at the command prompt.

Viewing your gateway system resource status

When your gateway starts, it checks its virtual CPU cores, root volume size, and RAM. It then determines whether these system resources are sufficient for your gateway to function properly. You can view the results of this check on the gateway's local console.

To view the status of a system resource check

- Log in to your gateway's local console:
 - For more information on logging in to the VMware ESXi console, see [Accessing the Gateway Local Console with VMware ESXi](#).

- For more information on logging in to the Microsoft Hyper-V local console, see [Access the Gateway Local Console with Microsoft Hyper-V](#).
 - For more information on logging in to the KVM local console, see [Accessing the Gateway Local Console with Linux KVM](#).
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **View System Resource Check**.

Each resource displays **[OK]**, **[WARNING]**, or **[FAIL]**, indicating the status of the resource as follows:

Message	Description
[OK]	The resource has passed the system resource check.
[WARNING]	The resource doesn't meet the recommended requirements, but your gateway can continue to function. Storage Gateway displays a message that describes the results of the resource check.
[FAIL]	The resource doesn't meet the minimum requirements. Your gateway might not function properly. Storage Gateway displays a message that describes the results of the resource check.

The console also displays the number of errors and warnings next to the resource check menu option.

Performing Tasks on the Amazon EC2 Local Console

Some Storage Gateway maintenance tasks require that you log in to the gateway local console for a gateway that you have deployed on an Amazon EC2 instance. You can access the gateway local console on your Amazon EC2 instance by using a Secure Shell (SSH) client. The topics in this section describes how to log in to the gateway local console and perform maintenance tasks.

Topics

- [Logging In to Your Amazon EC2 Gateway Local Console](#) - Learn about how you can connect and log in to the gateway local console your Amazon EC2 instance by using a Secure Shell (SSH) client.
- [Routing your gateway deployed on EC2 through an HTTP proxy](#) - Learn about how you can configure Storage Gateway to route all AWS endpoint traffic through a Socket Secure version 5 (SOCKS5) proxy server to your Amazon EC2 gateway instance.
- [Testing gateway network connectivity](#) - Learn about how you can use the gateway local console to test network connectivity between your gateway and various network resources.
- [Viewing your gateway system resource status](#) - Learn about how you can use the gateway local console to check the virtual CPU cores, root volume size, and RAM that are available to your gateway appliance.
- [Running Storage Gateway commands on the local console](#) - Learn about how you can run local console commands that allow you to perform additional tasks such as saving routing tables, connecting to AWS Support, and more.

Logging In to Your Amazon EC2 Gateway Local Console

You can connect to your Amazon EC2 instance by using a Secure Shell (SSH) client. For detailed information, see [Connect to Your Instance](#) in the *Amazon EC2 User Guide*. To connect this way, you will need the SSH key pair you specified when you launched the instance. For information about Amazon EC2 key pairs, see [Amazon EC2 Key Pairs](#) in the *Amazon EC2 User Guide*.

To log in to the gateway local console

1. Log in to your local console. If you are connecting to your EC2 instance from a Windows computer, log in as *admin*.
2. After you log in, you see the **AWS Storage Gateway - Configuration** main menu, from which you can perform various tasks.

To Learn About This Task

Configure a SOCKS proxy for your gateway

See This Topic

[Routing your gateway deployed on EC2 through an HTTP proxy](#)

To Learn About This Task	See This Topic
Test network connectivity	Testing gateway network connectivity
Run Storage Gateway console commands	Running Storage Gateway commands on the local console
View a system resource check	Viewing your gateway system resource status.

To shut down the gateway, enter **0**.

To exit the configuration session, enter **X**.

Routing your gateway deployed on EC2 through an HTTP proxy

Storage Gateway supports the configuration of a Socket Secure version 5 (SOCKS5) proxy between your gateway deployed on Amazon EC2 and AWS.

If your gateway must use a proxy server to communicate to the internet, then you need to configure the HTTP proxy settings for your gateway. You do this by specifying an IP address and port number for the host running your proxy. After you do so, Storage Gateway routes all AWS endpoint traffic through your proxy server. Communications between the gateway and endpoints is encrypted, even when using the HTTP proxy.

To route your gateway internet traffic through a local proxy server

1. Log in to your gateway's local console. For instructions, see [Logging In to Your Amazon EC2 Gateway Local Console](#).
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **Configure HTTP Proxy**.
3. From the **AWS Appliance Activation HTTP Proxy Configuration** menu, enter the corresponding numeral for the task you want to perform:
 - **Configure HTTP proxy** - You will need to supply a host name and port to complete configuration.

- **View current HTTP proxy configuration** - If an HTTP proxy is not configured, the message HTTP Proxy not configured is displayed. If an HTTP proxy is configured, the host name and port of the proxy are displayed.
- **Remove an HTTP proxy configuration** - The message HTTP Proxy Configuration Removed is displayed.

Testing gateway network connectivity

You can use your gateway's local console to test your network connectivity. This test can be useful when you are troubleshooting network issues with your gateway.

To test your gateway's connectivity

1. Log in to your gateway's local console. For instructions, see [Logging In to Your Amazon EC2 Gateway Local Console](#).
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **Test Network Connectivity**.

If your gateway has already been activated, the connectivity test begins immediately. For gateways that have not yet been activated, you must specify the endpoint type and AWS Region as described in the following steps.

3. If your gateway is not yet activated, enter the corresponding numeral to select the endpoint type for your gateway.
4. If you selected the public endpoint type, enter the corresponding numeral to select the AWS Region that you want to test. For supported AWS Regions and a list of AWS service endpoints you can use with Storage Gateway, see [AWS Storage Gateway endpoints and quotas](#) in the *AWS General Reference*.

As the test progresses, each endpoint displays either **[PASSED]** or **[FAILED]**, indicating the status of the connection as follows:

Message	Description
[PASSED]	Storage Gateway has network connectivity.

Message	Description
[FAILED]	Storage Gateway does not have network connectivity.

Viewing your gateway system resource status

When your gateway starts, it checks its virtual CPU cores, root volume size, and RAM. It then determines whether these system resources are sufficient for your gateway to function properly. You can view the results of this check on the gateway's local console.

To view the status of a system resource check

1. Log in to your gateway's local console. For instructions, see [Logging In to Your Amazon EC2 Gateway Local Console](#).
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **View System Resource Check**.

Each resource displays **[OK]**, **[WARNING]**, or **[FAIL]**, indicating the status of the resource as follows:

Message	Description
[OK]	The resource has passed the system resource check.
[WARNING]	The resource doesn't meet the recommended requirements, but your gateway can continue to function. Storage Gateway displays a message that describes the results of the resource check.
[FAIL]	The resource doesn't meet the minimum requirements. Your gateway might not function properly. Storage Gateway displays a message that describes the results of the resource check.

The console also displays the number of errors and warnings next to the resource check menu option.

Running Storage Gateway commands on the local console


The AWS Storage Gateway console helps provide a secure environment for configuring and diagnosing issues with your gateway. Using the console commands, you can perform maintenance tasks such as saving routing tables or connecting to AWS Support.

To run a configuration or diagnostic command


1. Log in to your gateway's local console. For instructions, see [Logging In to Your Amazon EC2 Gateway Local Console](#).
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **Gateway Console**.
3. From the gateway console command prompt, enter h.

The console displays the **AVAILABLE COMMANDS** menu, which lists the available commands:

Command	Function
dig	Collect output from dig for DNS troubleshooting.
exit	Return to Configuration menu.
h	Display available command list.
ifconfig	View or configure network interfaces.

 **Note**

We recommend configuring network or IP settings using the Storage Gateway console or the dedicated local console menu option.

Command	Function
ip	Show / manipulate routing, devices, and tunnels. <div data-bbox="834 352 1507 667" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px;"> <p> Note We recommend configuring network or IP settings using the Storage Gateway console or the dedicated local console menu option.</p> </div>
iptables	Administration tool for IPv4 packet filtering and NAT.
ncport	Test connectivity to a specific TCP port on a network.
nping	Collect output from nping for network troubleshooting.
open-support-channel	Connect to AWS Support.
save-iptables	Persist IP tables.
save-routing-table	Save newly added routing table entry.
sslcheck	Check SSL validity for network troubleshooting.
tcptraceroute	Collect traceroute output on TCP traffic to a destination.

- From the gateway console command prompt, enter the corresponding command for the function you want to use, and follow the instructions.

To learn about a command, enter the command name followed by the `-h` option, for example:
`sslcheck -h`.

Performance and optimization for Volume Gateway

This section describes Storage Gateway performance.

Topics

- [Optimizing gateway performance](#)

Optimizing gateway performance

Recommended Gateway Server Configuration

To obtain the best performance out of your gateway, Storage Gateway recommends the following gateway configuration for your gateway's host server:

- At least 24 dedicated physical CPU cores
- For Volume Gateway, your hardware should dedicate the following amounts of RAM:
 - At least 16 GiB of reserved RAM for gateways with cache size up to 16 TiB
 - At least 32 GiB of reserved RAM for gateways with cache size 16 TiB to 32 TiB
 - At least 48 GiB of reserved RAM for gateways with cache size 32 TiB to 64 TiB
- Disk 1, to be used as the gateway cache as follows:
 - SSD using an NVMe controller.
- Disk 2, to be used as the gateway upload buffer as follows:
 - SSD using an NVMe controller.
- Disk 3, to be used as the gateway upload buffer as follows:
 - SSD using an NVMe controller.
- Network adapter 1 configured on VM network 1:
 - Use VM network 1 and add VMXnet3 (10 Gbps) to be used for ingestion.
- Network adapter 2 configured on VM network 2:
 - Use VM network 2 and add a VMXnet3 (10 Gbps) to be used to connect to AWS.

Add Resources to Your Gateway

The following bottlenecks can reduce the performance of your Volume Gateway below the theoretical maximum sustained throughput (your bandwidth to AWS cloud):

- CPU core count
- Cache/Upload buffer disk throughput
- Total RAM amount
- Network bandwidth to AWS
- Network bandwidth from initiator to gateway

This section contains steps you can take in order to optimize the performance of your gateway. This guidance is based on adding resources to your gateway or your application server.

You can optimize gateway performance by adding resources to your gateway in one or more of the following ways.

Use higher-performance disks

Cache and upload buffer disk throughput can limit your gateway's upload and download performance. If your gateway is exhibiting performance significantly below what is expected, consider improving the cache and upload buffer disk throughput by:

- Using a striped RAID such as RAID 10 to improve disk throughput, ideally with a hardware RAID controller.

Note


RAID (redundant array of independent disks) or specifically disk striped RAID configurations like RAID 10, is the process of dividing a body of data into blocks and spreading the data blocks across multiple storage devices. The RAID level you use affects the exact speed and fault tolerance you can achieve. By striping IO workloads out across multiple disks, the overall throughput of the RAID device is much higher than that of any single member disk.

- Using directly attached, high performance disks

To optimize gateway performance, you can add high-performance disks such as solid-state drives (SSDs) and a NVMe controller. You can also attach virtual disks to your VM directly

from a storage area network (SAN) instead of the Microsoft Hyper-V NTFS. Improved disk performance generally results in better throughput and more 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 5 minutes divided by 300 seconds gives you the IOPS. As a general rule, when you review these metrics for a gateway, look for low throughput and low IOPS trends to indicate disk-related bottlenecks. .

 **Note**

CloudWatch metrics are not available for all gateways. For information about gateway metrics, see [Monitoring Storage Gateway](#).

Add more upload buffer disks

To achieve higher write throughput, add at least two upload buffer disks. When data is written to the gateway, it is written and stored locally on the upload buffer disks. Afterwards, the stored local data is asynchronously read from the disks to be processed and uploaded to AWS. Adding more upload buffer disks may reduce the amount of concurrent I/O operations performed to each individual disk. This can result in increased write throughput to the gateway.

Back gateway virtual disks with separate physical disks

When you provision gateway disks, we strongly recommend that you *don't* provision local disks for the upload buffer and cache storage that use the same underlying physical storage disk. For example, for VMware ESXi, the underlying physical storage resources are represented as a data store. When you deploy the gateway VM, you choose a data store on which to store the VM files. When you provision a virtual disk (for example, as an upload buffer), you can store the virtual disk in the same data store as the VM or a different data store.

If you have more than one data store, then we strongly recommend that you choose one data store for each type of local storage you are creating. A data store that is backed by only one underlying physical disk can lead to poor performance. An example is when you use such a disk to back both the cache storage and upload buffer in a gateway setup. Similarly, a data store that is backed by a less high-performing RAID configuration such as RAID 1 or RAID 6 can lead to poor performance.

Add CPU resources to your gateway host

The minimum requirement for a gateway host server is four virtual processors. To optimize gateway performance, confirm that each virtual processor that is assigned to the gateway VM is backed by a dedicated CPU core. In addition, confirm 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. Doing this allows your 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 help ensure that your gateway gets enough CPU resources when the host is shared with other VMs. Providing enough CPU resources has the general effect of improving throughput.

Increase bandwidth between your gateway and AWS cloud

Increasing your bandwidth to and from AWS will increase the maximum rate of data ingress to your gateway and egress to AWS cloud. This can improve your gateway performance if network speed is the limiting factor in your gateway configuration, rather than other factors like slow disks or poor gateway-initiator connection bandwidth.

Note

Your observed gateway performance will likely be lower than your network bandwidth due to other limiting factors listed here, such as cache/upload buffer disk throughput, CPU core count, total RAM amount, or the bandwidth between your initiator and gateway. Furthermore, your gateway's normal operation involves many actions taken to protect your data, which might cause the observed performance to be less than your network bandwidth.

Change the volumes configuration

For Volume Gateways, if you find that adding more volumes to a gateway reduces the throughput to the gateway, consider adding the volumes to a separate gateway. In particular, if a volume is used for a high-throughput application, consider creating a separate gateway for the high-throughput application. However, as a general rule, you should not use one gateway for all of your high-throughput applications and another gateway for all of your low-throughput applications. To measure your volume throughput, use the `ReadBytes` and `WriteBytes` metrics.

For more information about these metrics, see [Measuring Performance Between Your Application and Gateway](#).

Optimize iSCSI Settings

You can optimize iSCSI settings on your iSCSI initiator to achieve higher I/O performance. We recommend choosing 256 KiB for `MaxReceiveDataSegmentLength` and `FirstBurstLength`, and 1 MiB for `MaxBurstLength`. For more information about configuring iSCSI settings, see [Customizing iSCSI Settings](#).

Note

These recommended settings can facilitate overall better performance. However, the specific iSCSI settings that are needed to optimize performance vary depending on which backup software you use. For details, see your backup software documentation.

Add Resources to Your Application Environment

Increase the bandwidth between your application server and your gateway

The connection between your iSCSI initiator and gateway can limit your upload and download performance. If your gateway is exhibiting performance significantly worse than expected and you have already improved your CPU core count and disk throughput, consider:

- Upgrading your network cables to have higher bandwidth between your initiator and gateway.

To optimize gateway performance, 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 to measure the total data throughput.

For your application, compare the measured throughput with the desired throughput. 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 your application environment

If your application can use additional CPU resources, then adding more CPUs can help your application to scale its I/O load.

Security in AWS Storage Gateway

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the Amazon Web Services Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS Compliance Programs](#). To learn about the compliance programs that apply to AWS Storage Gateway, see [AWS Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Storage Gateway. The following topics show you how to configure Storage Gateway to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Storage Gateway resources.

Topics

- [Data protection in AWS Storage Gateway](#)
- [Identity and Access Management for AWS Storage Gateway](#)
- [Compliance validation for AWS Storage Gateway](#)
- [Resilience in AWS Storage Gateway](#)
- [Infrastructure Security in AWS Storage Gateway](#)
- [AWS Security Best Practices](#)
- [Logging and Monitoring in AWS Storage Gateway](#)

Data protection in AWS Storage Gateway

The AWS [shared responsibility model](#) applies to data protection in AWS Storage Gateway. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see [Working with CloudTrail trails](#) in the *AWS CloudTrail User Guide*.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-3](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Storage Gateway or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Data encryption using AWS KMS

Storage Gateway uses SSL/TLS (Secure Socket Layers/Transport Layer Security) to encrypt data that is transferred between your gateway appliance and AWS storage. By default, Storage Gateway uses Amazon S3-Managed Encryption Keys (SSE-S3) to server-side encrypt all data it stores in Amazon S3. You have an option to use the Storage Gateway API to configure your gateway to encrypt data stored in the cloud using server-side encryption with AWS Key Management Service (SSE-KMS) keys.

Important

When you use an AWS KMS key for server-side encryption, you must choose a symmetric key. Storage Gateway does not support asymmetric keys. For more information, see [Using symmetric and asymmetric keys](#) in the *AWS Key Management Service Developer Guide*.

Encrypting a file share

For a file share, you can configure your gateway to encrypt your objects with AWS KMS–managed keys by using SSE-KMS. For information on using the Storage Gateway API to encrypt data written to a file share, see [CreateNFSFileShare](#) in the *AWS Storage Gateway API Reference*.

Encrypting a volume

For cached and stored volumes, you can configure your gateway to encrypt volume data stored in the cloud with AWS KMS–managed keys by using the Storage Gateway API. You can specify one of the managed keys as the KMS key. The key that you use to encrypt your volume can't be changed after the volume is created. For information on using the Storage Gateway API to encrypt data written to a cached or stored volume, see [CreateCachediSCSIVolume](#) or [CreateStorediSCSIVolume](#) in the *AWS Storage Gateway API Reference*.

Encrypting a tape

For a virtual tape, you can configure your gateway to encrypt tape data stored in the cloud with AWS KMS–managed keys by using the Storage Gateway API. You can specify one of the managed keys as the KMS key. The key that you use to encrypt your tape data can't be changed after the tape is created. For information on using the Storage Gateway API to encrypt data written to a virtual tape, see [CreateTapes](#) in the *AWS Storage Gateway API Reference*.

When using AWS KMS to encrypt your data, keep the following in mind:

- Your data is encrypted at rest in the cloud. That is, the data is encrypted in Amazon S3.
- IAM users must have the required permissions to call the AWS KMS API operations. For more information, see [Using IAM policies with AWS KMS](#) in the *AWS Key Management Service Developer Guide*.
- If you delete or deactivate your AWS KMS key or revoke the grant token, you can't access the data on the volume or tape. For more information, see [Deleting KMS keys](#) in the *AWS Key Management Service Developer Guide*.
- If you create a snapshot from a volume that is KMS-encrypted, the snapshot is encrypted. The snapshot inherits the volume's KMS key.
- If you create a new volume from a snapshot that is KMS-encrypted, the volume is encrypted. You can specify a different KMS key for the new volume.

Note

Storage Gateway doesn't support creating an unencrypted volume from a recovery point of a KMS-encrypted volume or a KMS-encrypted snapshot.

For more information about AWS KMS, see [What is AWS Key Management Service?](#)

Configuring CHAP authentication for your volumes

In Storage Gateway, your iSCSI initiators connect to your volumes as iSCSI targets. Storage Gateway uses Challenge-Handshake Authentication Protocol (CHAP) to authenticate iSCSI and initiator connections. CHAP provides protection against playback attacks by requiring authentication to access storage volume targets. For each volume target, you can define one or more CHAP credentials. You can view and edit these credentials for the different initiators in the Configure CHAP credentials dialog box.

To configure CHAP credentials

1. In the Storage Gateway Console, choose **Volumes** and select the volume for which you want to configure CHAP credentials.
2. For **Actions**, choose **Configure CHAP authentication**.
3. For **Initiator name**, type the name of your initiator. The name must be at least 1 character and at most 255 characters long.

4. For **Initiator secret**, provide the secret phrase you want to use to authenticate your iSCSI initiator. The initiator secret phrase must be at least 12 characters and at most 16 characters long.
5. For **Target secret**, provide the secret phrase you want used to authenticate your target for mutual CHAP. The target secret phrase must be at least 12 characters and at most 16 characters long.
6. Choose **Save** to save your entries.

To view or update CHAP credentials, you must have the necessary IAM role permissions that allow you to perform that operation.

Viewing and editing CHAP credentials

You can add, remove or update CHAP credentials for each user. You must have the necessary IAM role permissions to view or edit CHAP credentials, and initiator target must be attached to a functioning gateway.

To add CHAP credentials

1. In the Storage Gateway Console, choose **Volumes** and select the volume for which you want to add CHAP credentials.
2. For **Actions**, choose **Configure CHAP authentication**.
3. In the Configure CHAPS page, provide the **Initiator name**, **Initiator secret**, and **Target secret** in the respective boxes and choose **Save**.

To remove CHAP credentials

1. In the Storage Gateway Console, choose **Volumes** and select the volume for which you want to remove CHAP credentials.
2. For **Actions**, choose **Configure CHAP authentication**.
3. Click the **X** next to the credentials you want to remove and choose **Save**.

To update CHAP credentials

1. In the Storage Gateway Console, choose **Volumes** and select the volume for which you want to update CHAP.

2. For **Actions**, choose **Configure CHAP authentication**.
3. In Configure CHAP credentials page, change the entries for the credentials you to update.
4. Choose **Save**.

Identity and Access Management for AWS Storage Gateway

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use AWS SGW resources. IAM is an AWS service that you can use with no additional charge.

Topics

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)
- [How AWS Storage Gateway works with IAM](#)
- [Identity-based policy examples for Storage Gateway](#)
- [Troubleshooting AWS Storage Gateway identity and access](#)

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in AWS SGW.

Service user – If you use the AWS SGW service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more AWS SGW features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in AWS SGW, see [Troubleshooting AWS Storage Gateway identity and access](#).

Service administrator – If you're in charge of AWS SGW resources at your company, you probably have full access to AWS SGW. It's your job to determine which AWS SGW features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand

the basic concepts of IAM. To learn more about how your company can use IAM with AWS SGW, see [How AWS Storage Gateway works with IAM](#).

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to AWS SGW. To view example AWS SGW identity-based policies that you can use in IAM, see [Identity-based policy examples for Storage Gateway](#).

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see [How to sign in to your AWS account](#) in the *AWS Sign-In User Guide*.

If you access AWS programmatically, AWS provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see [Signing AWS API requests](#) in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see [Multi-factor authentication](#) in the *AWS IAM Identity Center User Guide* and [Using multi-factor authentication \(MFA\) in AWS](#) in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account.

We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

Federated identity

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access AWS services by using temporary credentials.

A *federated identity* is a user from your enterprise user directory, a web identity provider, the AWS Directory Service, the Identity Center directory, or any user that accesses AWS services by using credentials provided through an identity source. When federated identities access AWS accounts, they assume roles, and the roles provide temporary credentials.

For centralized access management, we recommend that you use AWS IAM Identity Center. You can create users and groups in IAM Identity Center, or you can connect and synchronize to a set of users and groups in your own identity source for use across all your AWS accounts and applications. For information about IAM Identity Center, see [What is IAM Identity Center?](#) in the *AWS IAM Identity Center User Guide*.

IAM users and groups

An [IAM user](#) is an identity within your AWS account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see [Rotate access keys regularly for use cases that require long-term credentials](#) in the *IAM User Guide*.

An [IAM group](#) is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see [When to create an IAM user \(instead of a role\)](#) in the *IAM User Guide*.

IAM roles

An [IAM role](#) is an identity within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. You can temporarily assume an IAM role in the AWS Management Console by [switching roles](#). You can assume a role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see [Using IAM roles](#) in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see [Creating a role for a third-party Identity Provider](#) in the *IAM User Guide*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For information about permissions sets, see [Permission sets](#) in the *AWS IAM Identity Center User Guide*.
- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some AWS services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.
- **Cross-service access** – Some AWS services use features in other AWS services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
- **Forward access sessions (FAS)** – When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).

- **Service role** – A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Creating a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.
- **Service-linked role** – A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see [When to create an IAM role \(instead of a user\)](#) in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. For more information about the structure and contents of JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the `iam:GetRole` action. A

user with that policy can get role information from the AWS Management Console, the AWS CLI, or the AWS API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Creating IAM policies](#) in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see [Choosing between managed policies and inline policies](#) in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the `Principal` field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each AWS account root user. For more information about Organizations and SCPs, see [Service control policies](#) in the *AWS Organizations User Guide*.
- **Session policies** – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

How AWS Storage Gateway works with IAM

Before you use IAM to manage access to AWS SGW, learn what IAM features are available to use with AWS SGW.

IAM features you can use with AWS Storage Gateway

IAM feature	AWS SGW support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys (service-specific)	Yes
ACLs	No
ABAC (tags in policies)	Partial
Temporary credentials	Yes
Forward access sessions (FAS)	Yes
Service roles	Yes
Service-linked roles	Yes

To get a high-level view of how AWS SGW and other AWS services work with most IAM features, see [AWS services that work with IAM](#) in the *IAM User Guide*.

Identity-based policies for AWS SGW

Supports identity-based policies: Yes

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Creating IAM policies](#) in the *IAM User Guide*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. You can't specify the principal in an identity-based policy because it applies to the user or role to which it is attached. To learn about all

of the elements that you can use in a JSON policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

Identity-based policy examples for AWS SGW

To view examples of AWS SGW identity-based policies, see [Identity-based policy examples for Storage Gateway](#).

Resource-based policies within AWS SGW

Supports resource-based policies: No

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are *IAM role trust policies* and *Amazon S3 bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. Adding a cross-account principal to a resource-based policy is only half of establishing the trust relationship. When the principal and the resource are in different AWS accounts, an IAM administrator in the trusted account must also grant the principal entity (user or role) permission to access the resource. They grant permission by attaching an identity-based policy to the entity. However, if a resource-based policy grants access to a principal in the same account, no additional identity-based policy is required. For more information, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Policy actions for AWS SGW

Supports policy actions: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Action` element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated AWS API operation.

There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

To see a list of AWS SGW actions, see [Actions Defined by AWS Storage Gateway](#) in the *Service Authorization Reference*.

Policy actions in AWS SGW use the following prefix before the action:

```
sgw
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [  
  "sgw:action1",  
  "sgw:action2"  
]
```

To view examples of AWS SGW identity-based policies, see [Identity-based policy examples for Storage Gateway](#).

Policy resources for AWS SGW

Supports policy resources: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

To see a list of AWS SGW resource types and their ARNs, see [Resources Defined by AWS Storage Gateway](#) in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see [Actions Defined by AWS Storage Gateway](#).

To view examples of AWS SGW identity-based policies, see [Identity-based policy examples for Storage Gateway](#).

Policy condition keys for AWS SGW

Supports service-specific policy condition keys: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element (or *Condition block*) lets you specify conditions in which a statement is in effect. The `Condition` element is optional. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple `Condition` elements in a statement, or multiple keys in a single `Condition` element, AWS evaluates them using a logical AND operation. If you specify multiple values for a single condition key, AWS evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see [IAM policy elements: variables and tags](#) in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

To see a list of AWS SGW condition keys, see [Condition Keys for AWS Storage Gateway](#) in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see [Actions Defined by AWS Storage Gateway](#).

To view examples of AWS SGW identity-based policies, see [Identity-based policy examples for Storage Gateway](#).

ACLs in AWS SGW

Supports ACLs: No

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

ABAC with AWS SGW

Supports ABAC (tags in policies): Partial

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In AWS, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many AWS resources. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide tag information in the [condition element](#) of a policy using the `aws:ResourceTag/key-name`, `aws:RequestTag/key-name`, or `aws:TagKeys` condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see [What is ABAC?](#) in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see [Use attribute-based access control \(ABAC\)](#) in the *IAM User Guide*.

Using temporary credentials with AWS SGW

Supports temporary credentials: Yes

Some AWS services don't work when you sign in using temporary credentials. For additional information, including which AWS services work with temporary credentials, see [AWS services that work with IAM](#) in the *IAM User Guide*.

You are using temporary credentials if you sign in to the AWS Management Console using any method except a user name and password. For example, when you access AWS using your

company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then switch roles. For more information about switching roles, see [Switching to a role \(console\)](#) in the *IAM User Guide*.

You can manually create temporary credentials using the AWS CLI or AWS API. You can then use those temporary credentials to access AWS. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see [Temporary security credentials in IAM](#).

Forward access sessions for AWS SGW

Supports forward access sessions (FAS): Yes

When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).

Service roles for AWS SGW

Supports service roles: Yes

A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Creating a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.

Warning

Changing the permissions for a service role might break AWS SGW functionality. Edit service roles only when AWS SGW provides guidance to do so.

Service-linked roles for AWS SGW

Supports service-linked roles: Yes

A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing service-linked roles, see [AWS services that work with IAM](#). Find a service in the table that includes a Yes in the **Service-linked role** column. Choose the **Yes** link to view the service-linked role documentation for that service.

Identity-based policy examples for Storage Gateway

By default, users and roles don't have permission to create or modify AWS SGW resources. They also can't perform tasks by using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS API. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see [Creating IAM policies](#) in the *IAM User Guide*.

For details about actions and resource types defined by AWS SGW, including the format of the ARNs for each of the resource types, see [Actions, Resources, and Condition Keys for AWS Storage Gateway](#) in the *Service Authorization Reference*.

Topics

- [Policy best practices](#)
- [Using the AWS SGW console](#)
- [Allow users to view their own permissions](#)

Policy best practices

Identity-based policies determine whether someone can create, access, or delete AWS SGW resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We

recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see [AWS managed policies](#) or [AWS managed policies for job functions](#) in the *IAM User Guide*.

- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see [Policies and permissions in IAM](#) in the *IAM User Guide*.
- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as AWS CloudFormation. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.
- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see [IAM Access Analyzer policy validation](#) in the *IAM User Guide*.
- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your AWS account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see [Configuring MFA-protected API access](#) in the *IAM User Guide*.

For more information about best practices in IAM, see [Security best practices in IAM](#) in the *IAM User Guide*.

Using the AWS SGW console

To access the AWS Storage Gateway console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the AWS SGW resources in your AWS account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

To ensure that users and roles can still use the AWS SGW console, also attach the AWS SGW *ConsoleAccess* or *ReadOnly* AWS managed policy to the entities. For more information, see [Adding permissions to a user](#) in the *IAM User Guide*.

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewOwnUserInfo",
      "Effect": "Allow",
      "Action": [
        "iam:GetUserPolicy",
        "iam:ListGroupsWithUser",
        "iam:ListAttachedUserPolicies",
        "iam:ListUserPolicies",
        "iam:GetUser"
      ],
      "Resource": ["arn:aws:iam::*:user/${aws:username}"]
    },
    {
      "Sid": "NavigateInConsole",
      "Effect": "Allow",
      "Action": [
        "iam:GetGroupPolicy",
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedGroupPolicies",
        "iam:ListGroupPolicies",
        "iam:ListPolicyVersions",
        "iam:ListPolicies",
        "iam:ListUsers"
      ],
      "Resource": "*"
    }
  ]
}
```

Troubleshooting AWS Storage Gateway identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with AWS SGW and IAM.

Topics

- [I am not authorized to perform an action in AWS SGW](#)
- [I am not authorized to perform iam:PassRole](#)
- [I want to allow people outside of my AWS account to access my AWS SGW resources](#)

I am not authorized to perform an action in AWS SGW

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the mateojackson IAM user tries to use the console to view details about a fictional *my-example-widget* resource but doesn't have the fictional `sgw:GetWidget` permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
sgw:GetWidget on resource: my-example-widget
```

In this case, the policy for the mateojackson user must be updated to allow access to the *my-example-widget* resource by using the `sgw:GetWidget` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to AWS SGW.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in AWS SGW. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my AWS SGW resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether AWS SGW supports these features, see [How AWS Storage Gateway works with IAM](#).
- To learn how to provide access to your resources across AWS accounts that you own, see [Providing access to an IAM user in another AWS account that you own](#) in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see [Providing access to AWS accounts owned by third parties](#) in the *IAM User Guide*.
- To learn how to provide access through identity federation, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Compliance validation for AWS Storage Gateway

Third-party auditors assess the security and compliance of AWS Storage Gateway as part of multiple AWS compliance programs. These include SOC, PCI, ISO, FedRAMP, HIPAA, MTSC, C5, K-ISMS, ENS High, OSPAR, and HITRUST CSF.

For a list of AWS services in scope of specific compliance programs, see [AWS Services in Scope by Compliance Program](#). For general information, see [AWS Compliance Programs](#).

You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#).

Your compliance responsibility when using Storage Gateway is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- [Security and Compliance Quick Start Guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying security- and compliance-focused baseline environments on AWS.
- [Architecting for HIPAA Security and Compliance Whitepaper](#) – This whitepaper describes how companies can use AWS to create HIPAA-compliant applications.
- [AWS Compliance Resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [Evaluating resources with rules](#) in the *AWS Config Developer Guide* – The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- [AWS Security Hub](#) – This AWS service provides a comprehensive view of your security state within AWS that helps you check your compliance with security industry standards and best practices.

Resilience in AWS Storage Gateway

The AWS global infrastructure is built around AWS Regions and Availability Zones.

An AWS Region is a physical location around the world where data centers are clustered. Each group of logical data centers is called an Availability Zone (AZ). Each AWS Region consists of a minimum of three isolated and physically separate AZs within a geographic area. Unlike other cloud providers, who often define a region as a single data center, the multiple AZ design of every AWS Region offers distinct advantages. Each AZ has independent power, cooling, and physical security and is connected via redundant, ultra-low-latency networks. If your deployment requires a focus on high availability, you can configure services and resources to in multiple AZs to achieve greater fault-tolerance.

AWS Regions meet the highest levels of infrastructure security, compliance, and data protection. All traffic between AZs is encrypted. The network performance is sufficient to accomplish

synchronous replication between AZs. AZs make partitioning services and resources for high availability easy. If your deployment is partitioned across AZs, your resources are better isolated and protected from issues such as power outages, lightning strikes, tornadoes, earthquakes, and more. AZs are physically separated by a meaningful distance from any other AZ, although all are within 100 km (60 miles) of each other.

For more information about AWS Regions and Availability Zones, see [AWS Global Infrastructure](#).

In addition to the AWS global infrastructure, Storage Gateway offers several features to help support your data resiliency and backup needs:

- Use VMware vSphere High Availability (VMware HA) to help protect storage workloads against hardware, hypervisor, or network failures. For more information, see [Using VMware vSphere High Availability with Storage Gateway](#).
- Use AWS Backup to back up your volumes. For more information, see [Backing up your volumes](#).
- Clone your volume from a recovery point. For more information, see [Cloning a cached volume from a recovery point](#).

Infrastructure Security in AWS Storage Gateway

As a managed service, AWS Storage Gateway is protected by the AWS global network security procedures that are described in the [Amazon Web Services: Overview of Security Processes](#) whitepaper.

You use AWS published API calls to access Storage Gateway through the network. Clients must support Transport Layer Security (TLS) 1.2. Clients must also support cipher suites with perfect forward secrecy (PFS) such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the [AWS Security Token Service](#) (AWS STS) to generate temporary security credentials to sign requests.

Note

You should treat the AWS Storage Gateway appliance as a managed virtual machine, and should not attempt to access or modify its installation in any way. Attempting to install scanning software or update any software packages using methods other than the normal

gateway update mechanism, may cause the gateway to malfunction and could impact our ability to support or fix the gateway.

AWS reviews, analyzes, and remediates CVEs on a regular basis. We incorporate fixes for these issues into Storage Gateway as part of our normal software release cycle. These fixes are typically applied as part of the normal gateway update process during scheduled maintenance windows. For more information about gateway updates, see .

AWS Security Best Practices

AWS provides a number of security features to consider as you develop and implement your own security policies. These best practices are general guidelines and don't represent a complete security solution. Because these practices might not be appropriate or sufficient for your environment, treat them as helpful considerations rather than prescriptions. For more information, see [AWS Security Best Practices](#).

Logging and Monitoring in AWS Storage Gateway

Storage Gateway is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Storage Gateway. CloudTrail captures all API calls for Storage Gateway as events. The calls captured include calls from the Storage Gateway console and code calls to the Storage Gateway API operations. If you create a trail, you can activate continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Storage Gateway. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Storage Gateway, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the [AWS CloudTrail User Guide](#).

Storage Gateway Information in CloudTrail

CloudTrail is activated on your Amazon Web Services account when you create the account. When activity occurs in Storage Gateway, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your Amazon Web Services account. For more information, see [Viewing Events with CloudTrail Event History](#).

For an ongoing record of events in your Amazon Web Services account, including events for Storage Gateway, create a trail. A *trail* allows CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- [Overview for Creating a Trail](#)
- [CloudTrail Supported Services and Integrations](#)
- [Configuring Amazon SNS Notifications for CloudTrail](#)
- [Receiving CloudTrail Log Files from Multiple Regions](#) and [Receiving CloudTrail Log Files from Multiple Accounts](#)

All of the Storage Gateway actions are logged and are documented in the [Actions](#) topic. For example, calls to the `ActivateGateway`, `ListGateways`, and `ShutdownGateway` actions generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail userIdentity Element](#).

Understanding Storage Gateway Log File Entries

A trail is a configuration that allows delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the action.

```

{ "Records": [{
    "eventVersion": "1.02",
    "userIdentity": {
        "type": "IAMUser",
        "principalId": "AIDAI5AUPEBH2M7JTNVC",
        "arn": "arn:aws:iam::111122223333:user/StorageGateway-team/JohnDoe",
        "accountId": "111122223333",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "userName": "JohnDoe"
    },
    "eventTime": "2014-12-04T16:19:00Z",
    "eventSource": "storagegateway.amazonaws.com",
    "eventName": "ActivateGateway",
    "awsRegion": "us-east-2",
    "sourceIPAddress": "192.0.2.0",
    "userAgent": "aws-cli/1.6.2 Python/2.7.6 Linux/2.6.18-164.el5",
    "requestParameters": {
        "gatewayTimezone": "GMT-5:00",
        "gatewayName": "cloudtrailgatewayvtl",
        "gatewayRegion": "us-east-2",
        "activationKey": "EHFBX-1NDD0-P0IVU-PI259-
DHK88",
        "gatewayType": "VTL"
    },
    "responseElements": {
        "gatewayARN":
"arn:aws:storagegateway:us-east-2:111122223333:gateway/cloudtrailgatewayvtl"
    },
    "requestID":
"54BTFGNQI71987UJD2IHTCT8NF1Q8GLLE1QEU3KPGG6F0KSTAUU0",
    "eventID": "635f2ea2-7e42-45f0-
bed1-8b17d7b74265",
    "eventType": "AwsApiCall",
    "apiVersion": "20130630",
    "recipientAccountId": "444455556666"
    }
]}

```

The following example shows a CloudTrail log entry that demonstrates the ListGateways action.

```

{
  "Records": [{
    "eventVersion": "1.02",

```

```
    "userIdentity": {
      "type": "IAMUser",
      "principalId": "AIDAI15AUEPBH2M7JTNVC",
      "arn": "arn:aws:iam::111122223333:user/StorageGateway-
team/JohnDoe",
      "accountId": "111122223333", "accessKeyId": "
AKIAIOSFODNN7EXAMPLE",
      "userName": "JohnDoe "
    },
    "eventTime": "2014 - 12 - 03T19: 41: 53Z ",
    "eventSource": "storagegateway.amazonaws.com ",
    "eventName": "ListGateways ",
    "awsRegion": "us-east-2 ",
    "sourceIPAddress": "192.0.2.0 ",
    "userAgent": "aws - cli / 1.6.2 Python / 2.7.6
Linux / 2.6.18 - 164.el5 ",
    "requestParameters": null,
    "responseElements": null,
    "requestID": "
6U2N42CU37KA08BG6V1I23FRSJ1Q8GLLE1QEU3KPGG6F0KSTAUU0 ",
    "eventID": "f76e5919 - 9362 - 48ff - a7c4 -
d203a189ec8d ",
    "eventType": "AwsApiCall ",
    "apiVersion": "20130630 ",
    "recipientAccountId": "444455556666"
  ]]
}
```

Troubleshooting your gateway

Following, you can find information about best practices and troubleshooting issues related to gateways, host platforms, volumes, high availability, data recovery, and snapshots. The on-premises gateway troubleshooting information covers gateways deployed on supported virtualization platforms. The troubleshooting information for high availability issues covers gateways running on VMware vSphere High Availability (HA) platform.

Topics

- [Troubleshooting: gateway offline issues](#) - Learn how to diagnose problems that can cause your gateway to appear offline in the Storage Gateway console.
- [Troubleshooting: internal error during gateway activation](#) - Learn what to do if you receive an internal error message when attempting to activate your Storage Gateway.
- [Troubleshooting on-premises gateway issues](#) - Learn about typical issues that you might encounter working with your on-premises gateways, and how to allow AWS Support to connect to your gateway to assist with troubleshooting.
- [Troubleshooting Microsoft Hyper-V setup](#) - Learn about typical issues that you might encounter when deploying Storage Gateway on the Microsoft Hyper-V platform.
- [Troubleshooting Amazon EC2 gateway issues](#) - Find information about typical issues that you might encounter when working with gateways deployed on Amazon EC2.
- [Troubleshooting hardware appliance issues](#) - Learn how to resolve issues that you might encounter with the Storage Gateway Hardware Appliance.
- [Troubleshooting volume issues](#) - Find information about most typical issues you might encounter when working with volumes, and the actions we suggest that you take to fix them.
- [Troubleshooting high availability issues](#) - Learn what to do if you experience issues with gateways that are deployed in a VMware HA environment.

Troubleshooting: gateway offline issues

Use the following troubleshooting information to determine what to do if the AWS Storage Gateway console shows that your gateway is offline.

Your gateway might be showing as offline for one or more of the following reasons:

- The gateway can't reach the Storage Gateway service endpoints.
- The gateway shut down unexpectedly.
- A cache disk associated with the gateway has been disconnected or modified, or has failed.

To bring your gateway back online, identify and resolve the issue that caused your gateway to go offline.

Check the associated firewall or proxy

If you configured your gateway to use a proxy, or you placed your gateway behind a firewall, then review the access rules of the proxy or firewall. The proxy or firewall must allow traffic to and from the network ports and service endpoints required by Storage Gateway. For more information, see [Network and firewall requirements](#).

Check for an ongoing SSL or deep-packet inspection of your gateway's traffic

If an SSL or deep-packet inspection is currently being performed on the network traffic between your gateway and AWS, then your gateway might not be able to communicate with the required service endpoints. To bring your gateway back online, you must disable the inspection.

Check for a power outage or hardware failure on the hypervisor host

A power outage or hardware failure on the hypervisor host of your gateway can cause your gateway to shut down unexpectedly and become unreachable. After you restore the power and network connectivity, your gateway will become reachable again.

After your gateway is back online, be sure to take steps to recover your data. For more information, see [Best practices for recovering your data](#).

Check for issues with an associated cache disk

Your gateway can go offline if at least one of the cache disks associated with your gateway was removed, changed, or resized, or if it is corrupted.

If a working cache disk was removed from the hypervisor host:

1. Shut down the gateway.

2. Re-add the disk.

Note

Make sure you add the disk to the same disk node.

3. Restart the gateway.

If a cache disk is corrupted, was replaced, or was resized:

1. Shut down the gateway.
2. Reset the cache disk.
3. Reconfigure the disk for cache storage.
4. Restart the gateway.

Troubleshooting: internal error during gateway activation

Storage Gateway activation requests traverse two network paths. Incoming activation requests sent by a client connect to the gateway's virtual machine (VM) or Amazon Elastic Compute Cloud (Amazon EC2) instance over port 80. If the gateway successfully receives the activation request, then the gateway communicates with the Storage Gateway endpoints to receive an activation key. If the gateway can't reach the Storage Gateway endpoints, then the gateway responds to the client with an internal error message.

Use the following troubleshooting information to determine what to do if you receive an internal error message when attempting to activate your AWS Storage Gateway.

Note

- Make sure you deploy new gateways using the latest virtual machine image file or Amazon Machine Image (AMI) version. You will receive an internal error if you attempt to activate a gateway that uses an outdated AMI.
- Make sure that you select the correct gateway type that you intend to deploy before you download the AMI. The .ova files and AMIs for each gateway type are different, and they are not interchangeable.

Resolve errors when activating your gateway using a public endpoint

To resolve activation errors when activating your gateway using a public endpoint, perform the following checks and configurations.

Check the required ports

For gateways deployed on-premises, check that the ports are open on your local firewall. For gateways deployed on an Amazon EC2 instance, check that the ports are open on the instance's security group. To confirm that the ports are open, run a telnet command on the public endpoint from a server. This server must be in the same subnet as the gateway. For example, the following telnet commands test the connection to port 443:

```
telnet d4kdq0yaxexbo.cloudfront.net 443
telnet storagegateway.region.amazonaws.com 443
telnet dp-1.storagegateway.region.amazonaws.com 443
telnet proxy-app.storagegateway.region.amazonaws.com 443
telnet client-cp.storagegateway.region.amazonaws.com 443
telnet anon-cp.storagegateway.region.amazonaws.com 443
```

To confirm that the gateway itself can reach the endpoint, access the gateway's local VM console (for gateways deployed on-premises). Or, you can SSH to the gateway's instance (for gateways deployed on Amazon EC2). Then, run a network connectivity test. Confirm that the test returns [PASSED]. For more information, see [Testing Your Gateway Connection to the Internet](#).

Note

The default login user name for the gateway console is `admin`, and the default password is `password`.

Make sure firewall security does not modify packets sent from the gateway to the public endpoints

SSL inspections, deep packet inspections, or other forms of firewall security can interfere with packets sent from the gateway. The SSL handshake fails if the SSL certificate is modified from what the activation endpoint expects. To confirm that there's no SSL inspection in progress, run an `OpenSSL` command on the main activation endpoint (`anon-`

cp.storagegateway.region.amazonaws.com) on port 443. You must run this command from a machine that's in the same subnet as the gateway:

```
$ openssl s_client -connect anon-cp.storagegateway.region.amazonaws.com:443 -
servername anon-cp.storagegateway.region.amazonaws.com
```

Note

Replace *region* with your AWS Region.

If there's no SSL inspection in progress, then the command returns a response similar to the following:

```
$ openssl s_client -connect anon-cp.storagegateway.us-east-2.amazonaws.com:443 -
servername anon-cp.storagegateway.us-east-2.amazonaws.com
CONNECTED(00000003)
depth=2 C = US, O = Amazon, CN = Amazon Root CA 1
verify return:1
depth=1 C = US, O = Amazon, OU = Server CA 1B, CN = Amazon
verify return:1
depth=0 CN = anon-cp.storagegateway.us-east-2.amazonaws.com
verify return:1
---
Certificate chain
 0 s:/CN=anon-cp.storagegateway.us-east-2.amazonaws.com
  i:/C=US/O=Amazon/OU=Server CA 1B/CN=Amazon
 1 s:/C=US/O=Amazon/OU=Server CA 1B/CN=Amazon
  i:/C=US/O=Amazon/CN=Amazon Root CA 1
 2 s:/C=US/O=Amazon/CN=Amazon Root CA 1
  i:/C=US/ST=Arizona/L=Scottsdale/O=Starfield Technologies, Inc./CN=Starfield Services
  Root Certificate Authority - G2
 3 s:/C=US/ST=Arizona/L=Scottsdale/O=Starfield Technologies, Inc./CN=Starfield Services
  Root Certificate Authority - G2
  i:/C=US/O=Starfield Technologies, Inc./OU=Starfield Class 2 Certification Authority
---
```

If there is an ongoing SSL inspection, then the response shows an altered certificate chain, similar to the following:


```
$ openssl s_client -connect anon-cp.storagegateway.ap-southeast-1.amazonaws.com:443 -
servername anon-cp.storagegateway.ap-southeast-1.amazonaws.com
CONNECTED(00000003)
depth=0 DC = com, DC = amazonaws, OU = AWS, CN = anon-cp.storagegateway.ap-
southeast-1.amazonaws.com
verify error:num=20:unable to get local issuer certificate
verify return:1
depth=0 DC = com, DC = amazonaws, OU = AWS, CN = anon-cp.storagegateway.ap-
southeast-1.amazonaws.com
verify error:num=21:unable to verify the first certificate
verify return:1
---
Certificate chain
 0 s:/DC=com/DC=amazonaws/OU=AWS/CN=anon-cp.storagegateway.ap-southeast-1.amazonaws.com
  i:/C=IN/O=Company/CN=Admin/ST=KA/L=New town/OU=SGW/emailAddress=admin@company.com
---
```

The activation endpoint accepts SSL handshakes only if it recognizes the SSL certificate. This means that the gateway's outbound traffic to the endpoints must be exempt from inspections performed by firewalls in your network. These inspections might be an SSL inspection or a deep packet inspection.

Check gateway time synchronization

Excessive time skews can cause SSL handshake errors. For on-premises gateways, you can use the gateway's local VM console to check your gateway's time synchronization. The time skew should be no larger than 60 seconds. For more information, see [Synchronizing Your Gateway VM Time](#).

The **System Time Management** option isn't available on gateways that are hosted on Amazon EC2 instances. To make sure Amazon EC2 gateways can properly synchronize time, confirm that the Amazon EC2 instance can connect to the following NTP server pool list over ports UDP and TCP 123:

- 0.amazon.pool.ntp.org
- 1.amazon.pool.ntp.org
- 2.amazon.pool.ntp.org
- 3.amazon.pool.ntp.org

Resolve errors when activating your gateway using an Amazon VPC endpoint

To resolve activation errors when activating your gateway using an Amazon Virtual Private Cloud (Amazon VPC) endpoint, perform the following checks and configurations.

Check the required ports

Make sure the required ports within your local firewall (for gateways deployed on-premises) or security group (for gateways deployed in Amazon EC2) are open. The ports required for connecting a gateway to a Storage Gateway VPC endpoint differ from those required when connecting a gateway to public endpoints. The following ports are required for connecting to a Storage Gateway VPC endpoint:

- TCP 443
- TCP 1026
- TCP 1027
- TCP 1028
- TCP 1031
- TCP 2222

For more information, see [Creating a VPC endpoint for Storage Gateway](#).

Additionally, check the security group that's attached to your Storage Gateway VPC endpoint. The default security group attached to the endpoint might not allow the required ports. Create a new security group that allows traffic from your gateway's IP address range over the required ports. Then, attach that security group to the VPC endpoint.

Note

Use the [Amazon VPC console](#) to verify the security group that's attached to the VPC endpoint. View your Storage Gateway VPC endpoint from the console, and then choose the **Security Groups** tab.

To confirm that the required ports are open, you can run telnet commands on the Storage Gateway VPC Endpoint. You must run these commands from a server that's in the same subnet as the

gateway. You can run the tests on the first DNS name that doesn't specify an Availability Zone. For example, the following telnet commands test the required port connections using the DNS name `vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com`:

```
telnet vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com 443
telnet vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com 1026
telnet vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com 1027
telnet vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com 1028
telnet vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com 1031
telnet vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com 2222
```

Make sure firewall security does not modify packets sent from the gateway to your Storage Gateway Amazon VPC endpoint

SSL inspections, deep packet inspections, or other forms of firewall security can interfere with packets sent from the gateway. The SSL handshake fails if the SSL certificate is modified from what the activation endpoint expects. To confirm that there's no SSL inspection in progress, run an OpenSSL command on your Storage Gateway VPC endpoint. You must run this command from a machine that's in the same subnet as the gateway. Run the command for each required port:

```
$ openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com:443 -servername
vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com

$ openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com:1026 -servername
vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com

$ openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com:1027 -servername
vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com

$ openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com:1028 -servername
vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com

$ openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com:1031 -servername
vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com
```

```
$ openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-
east-1.vpce.amazonaws.com:2222 -servername
vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com
```

If there's no SSL inspection in progress, then the command returns a response similar to the following:

```
openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-
east-1.vpce.amazonaws.com:1027 -servername
vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com
CONNECTED(00000005)
depth=2 C = US, O = Amazon, CN = Amazon Root CA 1
verify return:1
depth=1 C = US, O = Amazon, OU = Server CA 1B, CN = Amazon
verify return:1
depth=0 CN = anon-cp.storagegateway.us-east-1.amazonaws.com
verify return:1
---
Certificate chain
 0 s:CN = anon-cp.storagegateway.us-east-1.amazonaws.com
  i:C = US, O = Amazon, OU = Server CA 1B, CN = Amazon
 1 s:C = US, O = Amazon, OU = Server CA 1B, CN = Amazon
  i:C = US, O = Amazon, CN = Amazon Root CA 1
 2 s:C = US, O = Amazon, CN = Amazon Root CA 1
  i:C = US, ST = Arizona, L = Scottsdale, O = "Starfield Technologies, Inc.", CN =
Starfield Services Root Certificate Authority - G2
 3 s:C = US, ST = Arizona, L = Scottsdale, O = "Starfield Technologies, Inc.", CN =
Starfield Services Root Certificate Authority - G2
  i:C = US, O = "Starfield Technologies, Inc.", OU = Starfield Class 2 Certification
Authority
---
```

If there is an ongoing SSL inspection, then the response shows an altered certificate chain, similar to the following:

```
openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-
east-1.vpce.amazonaws.com:1027 -servername
vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com
CONNECTED(00000005)
depth=2 C = US, O = Amazon, CN = Amazon Root CA 1
verify return:1
depth=1 C = US, O = Amazon, OU = Server CA 1B, CN = Amazon
```

```
verify return:1
depth=0 DC = com, DC = amazonaws, OU = AWS, CN = anon-cp.storagegateway.us-
east-1.amazonaws.com
verify error:num=21:unable to verify the first certificate
verify return:1
---
Certificate chain
 0 s:/DC=com/DC=amazonaws/OU=AWS/CN=anon-cp.storagegateway.us-east-1.amazonaws.com
  i:/C=IN/O=Company/CN=Admin/ST=KA/L=New town/OU=SGW/emailAddress=admin@company.com
---
```

The activation endpoint accepts SSL handshakes only if it recognizes the SSL certificate. This means that the gateway's outbound traffic to your VPC endpoint over required ports is exempt from inspections performed by your network firewalls. These inspections might be SSL inspections or deep packet inspections.

Check gateway time synchronization

Excessive time skews can cause SSL handshake errors. For on-premises gateways, you can use the gateway's local VM console to check your gateway's time synchronization. The time skew should be no larger than 60 seconds. For more information, see [Synchronizing Your Gateway VM Time](#).

The **System Time Management** option isn't available on gateways that are hosted on Amazon EC2 instances. To make sure Amazon EC2 gateways can properly synchronize time, confirm that the Amazon EC2 instance can connect to the following NTP server pool list over ports UDP and TCP 123:

- 0.amazon.pool.ntp.org
- 1.amazon.pool.ntp.org
- 2.amazon.pool.ntp.org
- 3.amazon.pool.ntp.org

Check for an HTTP proxy and confirm associated security group settings

Before activation, check if you have an HTTP proxy on Amazon EC2 configured on the on-premises gateway VM as a Squid proxy on port 3128. In this case, confirm the following:

- The security group attached to the HTTP proxy on Amazon EC2 must have an inbound rule. This inbound rule must allow Squid proxy traffic on port 3128 from the gateway VM's IP address.

- The security group attached to the Amazon EC2 VPC endpoint must have inbound rules. These inbound rules must allow traffic on ports 1026-1028, 1031, 2222, and 443 from the IP address of the HTTP proxy on Amazon EC2.

Resolve errors when activating your gateway using a public endpoint and there is a Storage Gateway VPC endpoint in the same VPC

To resolve errors when activating your gateway using a public endpoint when there is a Amazon Virtual Private Cloud (Amazon VPC) endpoint in the same VPC, perform the following checks and configurations.

Confirm that the Enable Private DNS Name setting isn't enabled on your Storage Gateway VPC endpoint

If **Enable Private DNS Name** is enabled, you can't activate any gateways from that VPC to the public endpoint.

To disable the private DNS name option:

1. Open the [Amazon VPC console](#).
2. In the navigation pane, choose **Endpoints**.
3. Choose your Storage Gateway VPC endpoint.
4. Choose **Actions**.
5. Choose **Manage Private DNS Names**.
6. For **Enable Private DNS Name**, clear **Enable for this Endpoint**.
7. Choose **Modify Private DNS Names** to save the setting.

Troubleshooting on-premises gateway issues

You can find information following about typical issues that you might encounter working with your on-premises gateways, and how to activate AWS Support to help troubleshoot your gateway.

The following table lists typical issues that you might encounter working with your on-premises gateways.

Issue	Action to Take
You cannot find the IP address of your gateway.	<p>Use the hypervisor client to connect to your host to find the gateway IP address.</p> <ul style="list-style-type: none">• For VMware ESXi, the VM's IP address can be found in the vSphere client on the Summary tab.• For Microsoft Hyper-V, the VM's IP address can be found by logging into the local console. <p>If you are still having trouble finding the gateway IP address:</p> <ul style="list-style-type: none">• Check that the VM is turned on. Only when the VM is turned on does an IP address get assigned to your gateway.• Wait for the VM to finish startup. If you just turned on your VM, then it might take several minutes for the gateway to finish its boot sequence.
You're having network or firewall problems.	<ul style="list-style-type: none">• Allow the appropriate ports for your gateway.• SSL cert validation/inspection should not be activated. Storage Gateway utilizes mutual TLS authentication which would fail if any 3rd party application tries to intercept/sign either certificate.• If you use a firewall or router to filter or limit network traffic, you must configure your firewall and router to allow these service endpoints for outbound communication to AWS. For more information about network and firewall requirements, see Network and firewall requirements.
Your gateway's activation fails when you click the Proceed to Activation button in the Storage Gateway Management Console.	<ul style="list-style-type: none">• Check that the gateway VM can be accessed by pinging the VM from your client.• Check that your VM has network connectivity to the internet. Otherwise, you'll need to configure a SOCKS proxy. For more information on doing so, see Configuring a SOCKS5 proxy for your on-premises gateway.

Issue	Action to Take
	<ul style="list-style-type: none">• Check that the host has the correct time, that the host is configured to synchronize its time automatically to a Network Time Protocol (NTP) server, and that the gateway VM has the correct time. For information about synchronizing the time of hypervisor hosts and VMs, see Synchronize VM time with Hyper-V or Linux KVM host time.• After performing these steps, you can retry the gateway deployment using the Storage Gateway console and the Setup and Activate Gateway wizard.• SSL cert validation/inspection should not be activated. Storage Gateway utilizes mutual TLS authentication which would fail if any 3rd party application tries to intercept/sign either certificate.• Check that your VM has at least 7.5 GB of RAM. Gateway allocation fails if there is less than 7.5 GB of RAM. For more information, see Requirements for setting up Volume Gateway.
<p>You need to remove a disk allocated as upload buffer space. For example, you might want to reduce the amount of upload buffer space for a gateway, or you might need to replace a disk used as an upload buffer that has failed.</p>	<p>For instructions about removing a disk allocated as upload buffer space, see Removing Disks from Your Gateway.</p>

Issue	Action to Take
You need to improve bandwidth between your gateway and AWS.	<p>You can improve the bandwidth from your gateway to AWS by setting up your internet connection to AWS on a network adapter (NIC) separate from that connecting your applications and the gateway VM. Taking this approach 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. For more on this subject, see Measuring Performance Between Your Gateway and AWS. Improving your internet connectivity helps to ensure that your upload buffer does not fill up.</p>

Issue	Action to Take
<p>Throughput to or from your gateway drops to zero.</p>	<ul style="list-style-type: none"> • On the Gateway tab of the Storage Gateway console, verify that the IP addresses for your gateway VM are the same that you see using your hypervisor client software (that is, the VMware vSphere client or Microsoft Hyper-V Manager). If you find a mismatch, restart your gateway from the Storage Gateway console, as shown in Shutting Down Your Gateway VM. After the restart, the addresses in the IP Addresses list in the Storage Gateway console's Gateway tab should match the IP addresses for your gateway, which you determine from the hypervisor client. <ul style="list-style-type: none"> • For VMware ESXi, the VM's IP address can be found in the vSphere client on the Summary tab. • For Microsoft Hyper-V, the VM's IP address can be found by logging into the local console. • Check your gateway's connectivity to AWS as described in Testing your gateway connection to the internet. • Check your gateway's network adapter configuration, and ensure that all the interfaces you intended to be activated for the gateway are activated. To view the network adapter configuration for your gateway, follow the instructions in Configuring Your Gateway Network 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 and AWS, see Measuring Performance Between Your Gateway and AWS.</p>
<p>You are having trouble importing (deploying) Storage Gateway on Microsoft Hyper-V.</p>	<p>See Troubleshooting Microsoft Hyper-V setup, which discusses some of the common issues of deploying a gateway on Microsoft Hyper-V.</p>

Issue	Action to Take
You receive a message that says: "The data that has been written to the volume in your gateway isn't securely stored at AWS".	You receive this message if your gateway VM was created from a clone or snapshot of another gateway VM. If this isn't the case, contact AWS Support.

Allowing AWS Support to help troubleshoot your gateway hosted on-premises

Storage Gateway provides a local console you can use to perform several maintenance tasks, including activating AWS Support to access your gateway to assist you with troubleshooting gateway issues. By default, AWS Support access to your gateway is deactivated. You provide this access through the host's local console. To give AWS Support access to your gateway, you first log in to the local console for the host, navigate to the Storage Gateway's console, and then connect to the support server.

To allow AWS Support access to your gateway

1. Log in to your host's local console.
 - VMware ESXi – for more information, see [Accessing the Gateway Local Console with VMware ESXi](#).
 - Microsoft Hyper-V – for more information, see [Access the Gateway Local Console with Microsoft Hyper-V](#).
2. At the prompt, enter the corresponding numeral to select **Gateway Console**.
3. Enter **h** to open the list of available commands.
4. Do one of the following:
 - If your gateway is using a public endpoint, in the **AVAILABLE COMMANDS** window, enter **open-support-channel** to connect to customer support for Storage Gateway. Allow TCP port 22 so you can open a support channel to AWS. When you connect to customer support, Storage Gateway assigns you a support number. Make a note of your support number.

- If your gateway is using a VPC endpoint, in the **AVAILABLE COMMANDS** window, enter **open-support-channel**. If your gateway is not activated, provide the VPC endpoint or IP address to connect to customer support for Storage Gateway. Allow TCP port 22 so you can open a support channel to AWS. When you connect to customer support, Storage Gateway assigns you a support number. Make a note of your support number.

Note

The channel number is not a Transmission Control Protocol/User Datagram Protocol (TCP/UDP) port number. Instead, the gateway makes a Secure Shell (SSH) (TCP 22) connection to Storage Gateway servers and provides the support channel for the connection.

5. After the support channel is established, provide your support service number to AWS Support so AWS Support can provide troubleshooting assistance.
6. When the support session is completed, enter **q** to end it. Don't close the session until Amazon Web Services Support notifies you that the support session is complete.
7. Enter **exit** to log out of the gateway console console.
8. Follow the prompts to exit the local console.

Troubleshooting Microsoft Hyper-V setup

The following table lists typical issues that you might encounter when deploying Storage Gateway on the Microsoft Hyper-V platform.

Issue	Action to Take
<p>You try to import a gateway and receive the following error message:</p> <p>"A server error occurred while attempting to import the virtual machine. Import failed.</p>	<p>This error can occur for the following reasons:</p> <ul style="list-style-type: none"> • If you are not pointing to the root of the unzipped gateway source files. The last part of the location you specify in the Import Virtual Machine dialog box should be AWS-Storage-Gateway . For example: <p>C:\prod-gateway\unzippedSourceVM\AWS-Storage-Gateway\ .</p>

Issue	Action to Take
<p>Unable to find virtual machine import files under location [...]. You can import a virtual machine only if you used Hyper-V to create and export it."</p>	<ul style="list-style-type: none">• If you have already deployed a gateway and you did not select the Copy the virtual machine option and check the Duplicate all files option in the Import Virtual Machine dialog box, then the VM was created in the location where you have the unzipped gateway files and you cannot import from this location again. To fix this problem, get a fresh copy of the unzipped gateway source files and copy to a new location. Use the new location as the source of the import. <p>If you plan on creating multiple gateways from one unzipped source files location, you must select Copy the virtual machine and check the Duplicate all files box in the Import Virtual Machine dialog box.</p>
<p>You try to import a gateway and receive the following error message:</p> <p>"A server error occurred while attempting to import the virtual machine. Import failed. Import task failed to copy file from [...]: The file exists. (0x80070050)"</p>	<p>If you have already deployed a gateway and you try to reuse the default folders that store the virtual hard disk files and virtual machine configuration files, then this error will occur. To fix this problem, specify new locations under Server in the panel on the left side of the Hyper-V Settings dialog box.</p>

Issue	Action to Take
<p>You try to import a gateway and receive the following error message:</p> <p>"A server error occurred while attempting to import the virtual machine. Import failed. Import failed because the virtual machine must have a new identifier. Select a new identifier and try the import again."</p>	<p>When you import the gateway make sure you select Copy the virtual machine and check the Duplicate all files box in the Import Virtual Machine dialog box to create a new unique ID for the VM.</p>
<p>You try to start a gateway VM and receive the following error message:</p> <p>"An error occurred while attempting to start the selected virtual machine(s). The child partition processor setting is incompatible with parent partition. 'AWS-Storage-Gateway' could not initialize. (Virtual machine ID [...])"</p>	<p>This error is likely caused by a CPU discrepancy between the required CPUs for the gateway and the available CPUs on the host. Ensure that the VM CPU count is supported by the underlying hypervisor.</p> <p>For more information about the requirements for Storage Gateway, see Requirements for setting up Volume Gateway.</p>

Issue	Action to Take
<p>You try to start a gateway VM and receive the following error message:</p> <p>"An error occurred while attempting to start the selected virtual machine(s). 'AWS-Storage-Gateway' could not initialize. (Virtual machine ID [...]) Failed to create partition: Insufficient system resources exist to complete the requested service. (0x800705AA)"</p>	<p>This error is likely caused by a RAM discrepancy between the required RAM for the gateway and the available RAM on the host.</p> <p>For more information about the requirements for Storage Gateway, see Requirements for setting up Volume Gateway.</p>
<p>Your snapshots and gateway software updates are occurring at slightly different times than expected.</p>	<p>The gateway VM's clock might be offset from the actual time, known as clock drift. Check and correct the VM's time using local gateway console's time synchronization option. For more information, see Synchronize VM time with Hyper-V or Linux KVM host time.</p>
<p>You need to put the unzipped Microsoft Hyper-V Storage Gateway files on the host file system.</p>	<p>Access the host as you do a typical Microsoft Windows server. For example, if the hypervisor host is name <code>hyperv-server</code>, then you can use the following UNC path <code>\\hyperv-server\c\$</code>, which assumes that the name <code>hyperv-server</code> can be resolved or is defined in your local hosts file.</p>
<p>You are prompted for credentials when connecting to hypervisor.</p>	<p>Add your user credentials as a local administrator for the hypervisor host by using the <code>Sconfig.cmd</code> tool.</p>

Issue	Action to Take
You may notice poor network performance if you turn on virtual machine queue (VMQ) for a Hyper-V host that's using a Broadcom network adapter.	For information about a workaround, see the Microsoft documentation, see Poor network performance on virtual machines on a Windows Server 2012 Hyper-V host if VMQ is turned on.

Troubleshooting Amazon EC2 gateway issues

In the following sections, you can find typical issues that you might encounter working with your gateway deployed on Amazon EC2. For more information about the difference between an on-premises gateway and a gateway deployed in Amazon EC2, see [Deploy a customized Amazon EC2 instance for Volume Gateway](#).

Topics

- [Your gateway activation hasn't occurred after a few moments](#)
- [You can't find your EC2 gateway instance in the instance list](#)
- [You created an Amazon EBS volume but can't attach it to your EC2 gateway instance](#)
- [You can't attach an initiator to a volume target of your EC2 gateway](#)
- [You get a message that you have no disks available when you try to add storage volumes](#)
- [You want to remove a disk allocated as upload buffer space to reduce upload buffer space](#)
- [Throughput to or from your EC2 gateway drops to zero](#)
- [You want AWS Support to help troubleshoot your EC2 gateway](#)
- [You want to connect to your gateway instance using the Amazon EC2 serial console](#)

Your gateway activation hasn't occurred after a few moments

Check the following in the Amazon EC2 console:

- Port 80 is activated in the security group that you associated with the instance. For more information about adding a security group rule, see [Adding a security group rule](#) in the *Amazon EC2 User Guide*.
- The gateway instance is marked as running. In the Amazon EC2 console, the **State** value for the instance should be RUNNING.
- Make sure that your Amazon EC2 instance type meets the minimum requirements, as described in [Storage requirements](#).

After correcting the problem, try activating the gateway again. To do this, open the Storage Gateway console, choose **Deploy a new Gateway on Amazon EC2**, and re-enter the IP address of the instance.

You can't find your EC2 gateway instance in the instance list

If you didn't give your instance a resource tag and you have many instances running, it can be hard to tell which instance you launched. In this case, you can take the following actions to find the gateway instance:

- Check the name of the Amazon Machine Image (AMI) on the **Description** tab of the instance. An instance based on the Storage Gateway AMI should start with the text **aws-storage-gateway-ami**.
- If you have several instances based on the Storage Gateway AMI, check the instance launch time to find the correct instance.

You created an Amazon EBS volume but can't attach it to your EC2 gateway instance

Check that the Amazon EBS volume in question is in the same Availability Zone as the gateway instance. If there is a discrepancy in Availability Zones, create a new Amazon EBS volume in the same Availability Zone as your instance.

You can't attach an initiator to a volume target of your EC2 gateway

Check that the security group that you launched the instance with includes a rule that allows the port that you are using for iSCSI access. The port is usually set as 3260. For more information on connecting to volumes, see [Connecting to your volumes from a Windows client](#).

You get a message that you have no disks available when you try to add storage volumes

For a newly activated gateway, no volume storage is defined. Before you can define volume storage, you must allocate local disks to the gateway to use as an upload buffer and cache storage. For a gateway deployed to Amazon EC2, the local disks are Amazon EBS volumes attached to the instance. This error message likely occurs because no Amazon EBS volumes are defined for the instance.

Check block devices defined for the instance that is running the gateway. If there are only two block devices (the default devices that come with the AMI), then you should add storage. For more information on doing so, see [Deploy a customized Amazon EC2 instance for Volume Gateway](#). After attaching two or more Amazon EBS volumes, try creating volume storage on the gateway.

You want to remove a disk allocated as upload buffer space to reduce upload buffer space

Follow the steps in [Determining the size of upload buffer to allocate](#).

Throughput to or from your EC2 gateway drops to zero

Verify that the gateway instance is running. If the instance is starting due to a reboot, for example, wait for the instance to restart.

Also, verify that the gateway IP has not changed. If the instance was stopped and then restarted, the IP address of the instance might have changed. In this case, you need to activate a new gateway.

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 and AWS, see [Measuring Performance Between Your Gateway and AWS](#).

You want AWS Support to help troubleshoot your EC2 gateway

Storage Gateway provides a local console you can use to perform several maintenance tasks, including activating AWS Support to access your gateway to assist you with troubleshooting gateway issues. By default, AWS Support access to your gateway is deactivated. You provide this access through the Amazon EC2 local console. You log in to the Amazon EC2 local console through

a Secure Shell (SSH). To successfully log in through SSH, your instance's security group must have a rule that opens TCP port 22.

Note

If you add a new rule to an existing security group, the new rule applies to all instances that use that security group. For more information about security groups and how to add a security group rule, see [Amazon EC2 security groups](#) in the *Amazon EC2 User Guide*.

To let AWS Support connect to your gateway, you first log in to the local console for the Amazon EC2 instance, navigate to the Storage Gateway's console, and then provide the access.

To activate AWS Support access to a gateway deployed on an Amazon EC2 instance

1. Log in to the local console for your Amazon EC2 instance. For instructions, go to [Connect to your instance](#) in the *Amazon EC2 User Guide*.

You can use the following command to log in to the EC2 instance's local console.

```
ssh -i PRIVATE-KEY admin@INSTANCE-PUBLIC-DNS-NAME
```

Note

The *PRIVATE-KEY* is the `.pem` file containing the private certificate of the EC2 key pair that you used to launch the Amazon EC2 instance. For more information, see [Retrieving the public key for your key pair](#) in the *Amazon EC2 User Guide*.

The *INSTANCE-PUBLIC-DNS-NAME* is the public Domain Name System (DNS) name of your Amazon EC2 instance that your gateway is running on. You obtain this public DNS name by selecting the Amazon EC2 instance in the EC2 console and clicking the **Description** tab.

2. At the prompt, enter **6 - Command Prompt** to open the AWS Support Channel console.
3. Enter **h** to open the **AVAILABLE COMMANDS** window.
4. Do one of the following:
 - If your gateway is using a public endpoint, in the **AVAILABLE COMMANDS** window, enter **open-support-channel** to connect to customer support for Storage Gateway. Allow TCP

port 22 so you can open a support channel to AWS. When you connect to customer support, Storage Gateway assigns you a support number. Make a note of your support number.

- If your gateway is using a VPC endpoint, in the **AVAILABLE COMMANDS** window, enter **open-support-channel1**. If your gateway is not activated, provide the VPC endpoint or IP address to connect to customer support for Storage Gateway. Allow TCP port 22 so you can open a support channel to AWS. When you connect to customer support, Storage Gateway assigns you a support number. Make a note of your support number.

Note

The channel number is not a Transmission Control Protocol/User Datagram Protocol (TCP/UDP) port number. Instead, the gateway makes a Secure Shell (SSH) (TCP 22) connection to Storage Gateway servers and provides the support channel for the connection.

5. After the support channel is established, provide your support service number to AWS Support so AWS Support can provide troubleshooting assistance.
6. When the support session is completed, enter **q** to end it. Don't close the session until AWS Support notifies you that the support session is complete.
7. Enter **exit** to exit the Storage Gateway console.
8. Follow the console menus to log out of the Storage Gateway instance.

You want to connect to your gateway instance using the Amazon EC2 serial console

You can use the Amazon EC2 serial console to troubleshoot boot, network configuration, and other issues. For instructions and troubleshooting tips, see [Amazon EC2 Serial Console](#) in the *Amazon Elastic Compute Cloud User Guide*.

Troubleshooting hardware appliance issues

The following topics discuss issues that you might encounter with the Storage Gateway Hardware Appliance, and suggestions on troubleshooting these.

You can't determine the service IP address

When attempting to connect to your service, make sure that you are using the service's IP address and not the host IP address. Configure the service IP address in the service console, and the host IP address in the hardware console. You see the hardware console when you start the hardware appliance. To go to the service console from the hardware console, choose **Open Service Console**.

How do you perform a factory reset?

If you need to perform a factory reset on your appliance, contact the Storage Gateway Hardware Appliance team for support, as described in the Support section following.

How do you perform a remote restart?

If you need to perform a remote restart of your appliance, you can do so using the Dell iDRAC management interface. For more information, see [iDRAC9 Virtual Power Cycle: Remotely power cycle Dell EMC PowerEdge Servers](#) on the Dell Technologies InfoHub website.

Where do you obtain Dell iDRAC support?

The Dell PowerEdge R640 server comes with the Dell iDRAC management interface. We recommend the following:

- If you use the iDRAC management interface, you should change the default password. For more information about the iDRAC credentials, see [Dell PowerEdge - What is the default sign-in credentials for iDRAC?](#)
- Make sure that the firmware is up-to-date to prevent security breaches.
- Moving the iDRAC network interface to a normal (em) port can cause performance issues or prevent the normal functioning of the appliance.

You can't find the hardware appliance serial number

You can find the serial number for your Storage Gateway Hardware Appliance using the Storage Gateway console.

To find the hardware appliance serial number:

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.

2. Choose **Hardware** from the navigation menu on the left side of the page.
3. Select your hardware appliance from the list.
4. Locate the **Serial Number** field on the **Details** tab for your appliance.

Where to obtain hardware appliance support

To contact AWS about technical support for your hardware appliance, see [AWS Support](#).

The AWS Support team might ask you to activate the support channel to troubleshoot your gateway issues remotely. You don't need this port to be open for the normal operation of your gateway, but it is required for troubleshooting. You can activate the support channel from the hardware console as shown in the procedure following.

To open a support channel for AWS

1. Open the hardware console.
2. Choose **Open Support Channel** at the bottom of the main page of the hardware console, and then press Enter.

The assigned port number should appear within 30 seconds if there are no network connectivity or firewall issues. For example:

Status: Open on port 19599

3. Note the port number and provide it to AWS Support.

Troubleshooting volume issues

You can find information about the most typical issues you might encounter when working with volumes, and actions that we suggest that you take to fix them.

Topics

- [The Console Says That Your Volume Is Not Configured](#)
- [The Console Says That Your Volume Is Irrecoverable](#)
- [Your Cached Gateway is Unreachable And You Want to Recover Your Data](#)
- [The Console Says That Your Volume Has PASS THROUGH Status](#)
- [You Want to Verify Volume Integrity and Fix Possible Errors](#)

- [Your Volume's iSCSI Target Doesn't Appear in Windows Disk Management Console](#)
- [You Want to Change Your Volume's iSCSI Target Name](#)
- [Your Scheduled Volume Snapshot Did Not Occur](#)
- [You Need to Remove or Replace a Disk That Has Failed](#)
- [Throughput from Your Application to a Volume Has Dropped to Zero](#)
- [A Cache Disk in Your Gateway Encounters a Failure](#)
- [A Volume Snapshot Has PENDING Status Longer Than Expected](#)
- [High Availability Health Notifications](#)

The Console Says That Your Volume Is Not Configured

If the Storage Gateway console indicates that your volume has a status of **UPLOAD BUFFER NOT CONFIGURED**, add upload buffer capacity to your gateway. You cannot use a gateway to store your application data if the upload buffer for the gateway is not configured. For more information, see [To configure additional upload buffer or cache storage for your gateway](#).

The Console Says That Your Volume Is Irrecoverable

For stored volumes, if the Storage Gateway console indicates that your volume has a status of **IRRECOVERABLE**, you can no longer use this volume. You can try to delete the volume in the Storage Gateway console. If there is data on the volume, then you can recover the data when you create a new volume based on the local disk of the VM that was initially used to create the volume. When you create the new volume, select **Preserve existing data**. Make sure to delete pending snapshots of the volume before deleting the volume. For more information, see [Deleting snapshots of your storage volumes](#). If deleting the volume in the Storage Gateway console does not work, then the disk allocated for the volume might have been improperly removed from the VM and cannot be removed from the appliance.

For cached volumes, if the Storage Gateway console indicates that your volume has a status of **IRRECOVERABLE**, you can no longer use this volume. If there is data on the volume, you can create a snapshot of the volume and then recover your data from the snapshot or you can clone the volume from the last recovery point. You can delete the volume after you have recovered your data. For more information, see [Your Cached Gateway is Unreachable And You Want to Recover Your Data](#).

For stored volumes, you can create a new volume from the disk that was used to create the irrecoverable volume. For more information, see [Creating a storage volume](#). For information about volume status, see [Understanding Volume Statuses and Transitions](#).

Your Cached Gateway is Unreachable And You Want to Recover Your Data

When your gateway becomes unreachable (such as when you shut it down), you have the option of either creating a snapshot from a volume recovery point and using that snapshot, or cloning a new volume from the last recovery point for an existing volume. Cloning from a volume recovery point is faster and more cost effective than creating a snapshot. For more information about cloning a volume, see [Cloning a cached volume from a recovery point](#).

Storage Gateway provides recovery points for each volume in a cached Volume Gateway architecture. A *volume recovery point* is a point in time at which all data of the volume is consistent and from which you can create a snapshot or clone a volume.

The Console Says That Your Volume Has PASS THROUGH Status

In some cases, the Storage Gateway console might indicate that your volume has a status of PASSTHROUGH. A volume can have PASSTHROUGH status for several reasons. Some reasons require action, and some do not.

An example of when you should take action if your volume has the PASS THROUGH status is when your gateway has run out of upload buffer space. To verify if your upload buffer was exceeded in the past, you can view the `UploadBufferPercentUsed` metric in the Amazon CloudWatch console; for more information, see [Monitoring the upload buffer](#). If your gateway has the PASS THROUGH status because it has run out of upload buffer space, you should allocate more upload buffer space to your gateway. Adding more buffer space will cause your volume to transition from PASS THROUGH to BOOTSTRAPPING to AVAILABLE automatically. While the volume has the BOOTSTRAPPING status, the gateway reads data off the volume's disk, uploads this data to Amazon S3, and catches up as needed. When the gateway has caught up and saved the volume data to Amazon S3, the volume status becomes AVAILABLE and snapshots can be started again. Note that when your volume has the PASS THROUGH or BOOTSTRAPPING status, you can continue to read and write data from the volume disk. For more information about adding more upload buffer space, see [Determining the size of upload buffer to allocate](#).

To take action before the upload buffer is exceeded, you can set a threshold alarm on a gateway's upload buffer. For more information, see [To set an upper threshold alarm for a gateway's upload buffer](#).

In contrast, an example of not needing to take action when a volume has the PASS THROUGH status is when the volume is waiting to be bootstrapped because another volume is currently being bootstrapped. The gateway bootstraps volumes one at a time.

Infrequently, the PASS THROUGH status can indicate that a disk allocated for an upload buffer has failed. In this is the case, you should remove the disk. For more information, see [Working with Volume Gateway storage resources](#). For information about volume status, see [Understanding Volume Statuses and Transitions](#).

You Want to Verify Volume Integrity and Fix Possible Errors

If you want to verify volume integrity and fix possible errors, and your gateway uses Microsoft Windows initiators to connect to its volumes, you can use the Windows CHKDSK utility to verify the integrity of your volumes and fix any errors on the volumes. Windows can automatically run the CHKDSK tool when volume corruption is detected, or you can run it yourself.

Your Volume's iSCSI Target Doesn't Appear in Windows Disk Management Console

If your volume's iSCSI target does not show up in the Disk Management Console in Windows, check that you have configured the upload buffer for the gateway. For more information, see [To configure additional upload buffer or cache storage for your gateway](#).

You Want to Change Your Volume's iSCSI Target Name

If you want to change the iSCSI target name of your volume, you must delete the volume and add it again with a new target name. If you do so, you can preserve the data on the volume.

Your Scheduled Volume Snapshot Did Not Occur

If your scheduled snapshot of a volume did not occur, check whether your volume has the PASSTHROUGH status, or if the gateway's upload buffer was filled just prior to the scheduled snapshot time. You can check the `UploadBufferPercentUsed` metric for the gateway in the Amazon CloudWatch console and create an alarm for this metric. For more information, see [Monitoring the upload buffer](#) and [To set an upper threshold alarm for a gateway's upload buffer](#).

You Need to Remove or Replace a Disk That Has Failed

If you need to replace a volume disk that has failed or replace a volume because it isn't needed, you should remove the volume first using the Storage Gateway console. For more information, see [To delete a volume](#). You then use the hypervisor client to remove the backing storage:

- For VMware ESXi, remove the backing storage as described in [Deleting storage volumes](#).
- For Microsoft Hyper-V, remove the backing storage.

Throughput from Your Application to a Volume Has Dropped to Zero

If throughput from your application to a volume has dropped to zero, try the following:

- If you are using the VMware vSphere client, check that your 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** address for a storage volume in the Storage Gateway console in the **Details** tab for the 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 Storage Gateway console as shown in [Shutting Down Your Gateway VM](#). 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.
- If there is no IP address in the **Host IP** box for the volume and the gateway is online. For example, this could occur if you create a volume associated with an IP address of a network adapter of a gateway that has two or more network adapters. When you remove or deactivate the network adapter that the volume is associated with, the IP address might not appear in the **Host IP** box. To address this issue, delete the volume and then re-create it preserving its existing data.
- 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 [Connecting to your volumes from a Windows client](#).

You can view the throughput for volumes and create alarms from the Amazon CloudWatch console. For more information about measuring throughput from your application to a volume, see [Measuring Performance Between Your Application and Gateway](#).

A Cache Disk in Your Gateway Encounters a Failure

If one or more cache disks in your gateway encounters a failure, the gateway prevents read and write operations to your virtual tapes and volumes. To resume normal functionality, reconfigure your gateway as described following:

- If the cache disk is inaccessible or unusable, delete the disk from your gateway configuration.
- If the cache disk is still accessible and useable, reconnect it to your gateway.

Note

If you delete a cache disk, tapes or volumes that have clean data (that is, for which data in the cache disk and Amazon S3 are synchronized) will continue to be available when the gateway resumes normal functionality. For example, if your gateway has three cache disks and you delete two, tapes or volumes that are clean will have AVAILABLE status. Other tapes and volumes will have IRRECOVERABLE status.

If you use ephemeral disks as cache disks for your gateway or mount your cache disks on an ephemeral drive, your cache disks will be lost when you shut down the gateway. Shutting down the gateway when your cache disk and Amazon S3 are not synchronized can result in data loss. As a result, we don't recommend using ephemeral drives or disks.

A Volume Snapshot Has PENDING Status Longer Than Expected

If a volume snapshot remains in PENDING state longer than expected, the gateway VM might have crashed unexpectedly or the status of a volume might have changed to PASS THROUGH or IRRECOVERABLE. If any of these are the case, the snapshot remains in PENDING status and the snapshot does not successfully complete. In these cases, we recommend that you delete the snapshot. For more information, see [Deleting snapshots of your storage volumes](#).

When the volume returns to AVAILABLE status, create a new snapshot of the volume. For information about volume status, see [Understanding Volume Statuses and Transitions](#).

High Availability Health Notifications

When running your gateway on the VMware vSphere High Availability (HA) platform, you may receive health notifications. For more information about health notifications, see [Troubleshooting high availability issues](#).

Troubleshooting high availability issues

You can find information following about actions to take if you experience availability issues.

Topics

- [Health notifications](#)
- [Metrics](#)

Health notifications

When you run your gateway on VMware vSphere HA, all gateways produce the following health notifications to your configured Amazon CloudWatch log group. These notifications go into a log stream called `AvailabilityMonitor`.

Topics

- [Notification: Reboot](#)
- [Notification: HardReboot](#)
- [Notification: HealthCheckFailure](#)
- [Notification: AvailabilityMonitorTest](#)

Notification: Reboot

You can get a reboot notification when the gateway VM is restarted. You can restart a gateway VM by using the VM Hypervisor Management console or the Storage Gateway console. You can also restart by using the gateway software during the gateway's maintenance cycle.

Action to Take

If the time of the reboot is within 10 minutes of the gateway's configured [maintenance start time](#), this is probably a normal occurrence and not a sign of any problem. If the reboot occurred significantly outside the maintenance window, check whether the gateway was restarted manually.

Notification: HardReboot

You can get a `HardReboot` notification when the gateway VM is restarted unexpectedly. Such a restart can be due to loss of power, a hardware failure, or another event. For VMware gateways, a reset by vSphere High Availability Application Monitoring can launch this event.

Action to Take

When your gateway runs in such an environment, check for the presence of the `HealthCheckFailure` notification and consult the VMware events log for the VM.

Notification: HealthCheckFailure

For a gateway on VMware vSphere HA, you can get a `HealthCheckFailure` notification when a health check fails and a VM restart is requested. This event also occurs during a test to monitor availability, indicated by an `AvailabilityMonitorTest` notification. In this case, the `HealthCheckFailure` notification is expected.

Note

This notification is for VMware gateways only.

Action to Take

If this event repeatedly occurs without an `AvailabilityMonitorTest` notification, check your VM infrastructure for issues (storage, memory, and so on). If you need additional assistance, contact AWS Support.

Notification: AvailabilityMonitorTest

For a gateway on VMware vSphere HA, you can get an `AvailabilityMonitorTest` notification when you [run a test](#) of the [Availability and application monitoring](#) system in VMware.

Metrics

The `AvailabilityNotifications` metric is available on all gateways. This metric is a count of the number of availability-related health notifications generated by the gateway. Use the Sum statistic to observe whether the gateway is experiencing any availability-related events. Consult with your configured CloudWatch log group for details about the events.

Best practices for Volume Gateway

This section contains the following topics, which provide information about the best practices for working with gateways, local disks, snapshots, and data. We recommend that you familiarize yourself with the information outlined in this section, and attempt to follow these guidelines in order to avoid problems with your AWS Storage Gateway. For additional guidance on diagnosing and solving common issues you might encounter with your deployment, see [Troubleshooting your gateway](#).

Topics

- [Best practices: recovering your data](#)
- [Cleaning up unnecessary resources](#)
- [Reducing the amount of billed storage on a volume](#)

Best practices: recovering your data

Although it is rare, your gateway might encounter an unrecoverable failure. Such a failure can occur in your virtual machine (VM), the gateway itself, the local storage, or elsewhere. If a failure occurs, we recommend that you follow the instructions in the appropriate section following to recover your data.

Important

Storage Gateway doesn't support recovering a gateway VM from a snapshot that is created by your hypervisor or from your Amazon EC2 Amazon Machine Image (AMI). If your gateway VM malfunctions, activate a new gateway and recover your data to that gateway using the instructions following.

Topics

- [Recovering from an unexpected virtual machine shutdown](#)
- [Recovering your data from a malfunctioning gateway or VM](#)
- [Recovering your data from an irrecoverable volume](#)
- [Recovering your data from a malfunctioning cache disk](#)
- [Recovering your data from a corrupted file system](#)

- [Recovering your data from an inaccessible data center](#)

Recovering from an unexpected virtual machine shutdown

If your VM shuts down unexpectedly, for example during a power outage, your gateway becomes unreachable. When power and network connectivity are restored, your gateway becomes reachable and starts to function normally. Following are some steps you can take at that point to help recover your data:

- If an outage causes network connectivity issues, you can troubleshoot the issue. For information about how to test network connectivity, see [Testing your gateway connection to the internet](#).
- For cached volumes setups, when your gateway becomes reachable, your volumes go into BOOTSTRAPPING status. This functionality ensures that your locally stored data continues to be synchronized with AWS. For more information on this status, see [Understanding Volume Statuses and Transitions](#).
- If your gateway malfunctions and issues occur with your volumes or tapes as a result of an unexpected shutdown, you can recover your data. For information about how to recover your data, see the sections following that apply to your scenario.

Recovering your data from a malfunctioning gateway or VM

If your gateway or virtual machine malfunctions, you can recover data that has been uploaded to AWS and stored on a volume in Amazon S3. For cached volumes gateways, you recover data from a recovery snapshot. For stored volumes gateways, you can recover data from your most recent Amazon EBS snapshot of the volume. For Tape Gateways, you recover one or more tapes from a recovery point to a new Tape Gateway.

If your cached volumes gateway becomes unreachable, you can use the following steps to recover your data from a recovery snapshot:

1. In the AWS Management Console, choose the malfunctioning gateway, choose the volume you want to recover, and then create a recovery snapshot from it.
2. Deploy and activate a new Volume Gateway. Or, if you have an existing functioning Volume Gateway, you can use that gateway to recover your volume data.
3. Find the snapshot you created and restore it to a new volume on the functioning gateway.
4. Mount the new volume as an iSCSI device on your on-premises application server.

For detailed information on how to recover cached volumes data from a recovery snapshot, see [Your Cached Gateway is Unreachable And You Want to Recover Your Data](#).

Recovering your data from an irrecoverable volume

If the status of your volume is IRRECOVERABLE, you can no longer use this volume.

For stored volumes, you can retrieve your data from the irrecoverable volume to a new volume by using the following steps:

1. Create a new volume from the disk that was used to create the irrecoverable volume.
2. Preserve existing data when you are creating the new volume.
3. Delete all pending snapshot jobs for the irrecoverable volume.
4. Delete the irrecoverable volume from the gateway.

For cached volumes, we recommend using the last recovery point to clone a new volume.

For detailed information about how to retrieve your data from an irrecoverable volume to a new volume, see [The Console Says That Your Volume Is Irrecoverable](#).

Recovering your data from a malfunctioning cache disk

If your cache disk encounters a failure, we recommend you use the following steps to recover your data depending on your situation:

- If the malfunction occurred because a cache disk was removed from your host, shut down the gateway, re-add the disk, and restart the gateway.
- If the cache disk is corrupted or not accessible, shut down the gateway, reset the cache disk, reconfigure the disk for cache storage, and restart the gateway.

Recovering your data from a corrupted file system

If your file system gets corrupted, you can use the **fsck** command to check your file system for errors and repair it. If you can repair the file system, you can then recover your data from the volumes on the file system, as described following:

1. Shut down your virtual machine and use the Storage Gateway Management Console to create a recovery snapshot. This snapshot represents the most current data stored in AWS.

Note

You use this snapshot as a fallback if your file system can't be repaired or the snapshot creation process can't be completed successfully.

For information about how to create a recovery snapshot, see [Your Cached Gateway is Unreachable And You Want to Recover Your Data](#).

2. Use the **fsck** command to check your file system for errors and attempt a repair.
3. Restart your gateway VM.
4. When your hypervisor host starts to boot up, press and hold down shift key to enter the grub boot menu.
5. From the menu, press **e** to edit.
6. Choose the kernel line (the second line), and then press **e** to edit.
7. Append the following option to the kernel command line: **init=/bin/bash**. Use a space to separate the previous option from the option you just appended.
8. Delete both `console=` lines, making sure to delete all values following the = symbol, including those separated by commas.
9. Press **Return** to save the changes.
10. Press **b** to boot your computer with the modified kernel option. Your computer will boot to a `bash#` prompt.
11. Enter `/sbin/fsck -f /dev/sda1` to run this command manually from the prompt, to check and repair your file system. If the command does not work with the `/dev/sda1` path, you can use `lsblk` to determine the root filesystem device for `/` and use that path instead.
12. When the file system check and repair is complete, reboot the instance. The grub settings will revert to the original values, and the gateway will boot up normally.
13. Wait for snapshots that are in-progress from the original gateway to complete, and then validate the snapshot data.

You can continue to use the original volume as-is, or you can create a new gateway with a new volume based on either the recovery snapshot or the completed snapshot. Alternatively, you can create a new volume from any of your completed snapshots from this volume.

Recovering your data from an inaccessible data center

If your gateway or data center becomes inaccessible for some reason, you can recover your data to another gateway in a different data center or recover to a gateway hosted on an Amazon EC2 instance. If you don't have access to another data center, we recommend creating the gateway on an Amazon EC2 instance. The steps you follow depends on the gateway type you are covering the data from.

To recover data from a Volume Gateway in an inaccessible data center

1. Create and activate a new Volume Gateway on an Amazon EC2 host. For more information, see [Deploy a customized Amazon EC2 instance for Volume Gateway](#).

Note

Gateway stored volumes can't be hosted on Amazon EC2 instance.

2. Create a new volume and choose the EC2 gateway as the target gateway. For more information, see [Creating a storage volume](#).

Create the new volume based on an Amazon EBS snapshot or clone from last recovery point of the volume you want to recover.

If your volume is based on a snapshot, provide the snapshot id.

If you are cloning a volume from a recovery point, choose the source volume.

Cleaning up unnecessary resources

If you created your gateway as an example exercise or a test, consider cleaning up to avoid incurring unexpected or unnecessary charges.

To clean up resources you don't need

1. Delete any snapshots. For instructions, see [Deleting snapshots of your storage volumes](#).
2. Unless you plan to continue using the gateway, delete it. For more information, see [Deleting your gateway and removing associated resources](#).
3. Delete the Storage Gateway VM from your on-premises host. If you created your gateway on an Amazon EC2 instance, terminate the instance.

Reducing the amount of billed storage on a volume

Deleting files from your file system doesn't necessarily delete data from the underlying block device or reduce the amount of data stored on your volume. If you want to reduce the amount of billed storage on your volume, we recommend overwriting your files with zeros to compress the storage to a negligible amount of actual storage. Storage Gateway charges for volume usage based on compressed storage.

Note

If you use a delete tool that overwrites the data on your volume with random data, your usage will not be reduced. This is because the random data is not compressible.

Additional Storage Gateway Resources

This section describes AWS and third-party software, tools, and resources that can help you set up or manage your gateway, and also Storage Gateway quotas.

Topics

- [Deploying and configuring the gateway VM host](#) - Learn how to deploy and configure a virtual machine host for your gateway.
- [Working with Volume Gateway storage resources](#) - Learn about procedures related to Volume Gateway storage resources, such as removing local disks and managing Amazon EBS volumes on gateway Amazon EC2 instances.
- [Getting an activation key for your gateway](#) - Learn where to find the activation key that you need to provide when you deploy a new gateway.
- [Connecting iSCSI Initiators](#) - Learn how to work with volumes or virtual tape library (VTL) devices that are exposed as Internet Small Computer System Interface (iSCSI) targets.
- [Using AWS Direct Connect with Storage Gateway](#) - Learn how to create a dedicated network connection between your on-premises gateway and the AWS cloud.
- [Port requirements for Volume Gateway](#) - Find specific information about the network ports required by Volume Gateway.
- [Getting the IP address for your gateway appliance](#) - Learn where to find the gateway's virtual machine host IP address, which you need to provide when you deploy a new gateway.
- [Understanding Storage Gateway Resources and Resource IDs](#) - Learn how AWS identifies the resources and subresources that are created by Storage Gateway.
- [Tagging Storage Gateway Resources](#) - Learn how to use metadata tags to categorize your resources and make them easier to manage.
- [Working with open-source components for Storage Gateway](#) - Learn about the third-party tools and licenses that are used to deliver Storage Gateway functionality.
- [AWS Storage Gateway quotas](#) - Learn about limits and quotas for Volume Gateway, including maximum limitations for volume size and quantity, and local disk size recommendations.

Deploying and configuring the gateway VM host

The topics in this section describe how to set up and manage the virtual machine host for your Storage Gateway appliance, including on-premises appliances running on VMware, Hyper-V, or Linux KVM, and appliances running on Amazon EC2 instances in the AWS cloud.


Topics

- [Deploy a default Amazon EC2 host for Volume Gateway](#) - Learn about how to deploy and activate a Volume Gateway on an Amazon Elastic Compute Cloud (Amazon EC2) instance using the default specifications.
- [Deploy a customized Amazon EC2 instance for Volume Gateway](#) - Learn about how to deploy and activate a Volume Gateway on an Amazon Elastic Compute Cloud (Amazon EC2) instance using customized settings.
- [Modify Amazon EC2 instance metadata options](#) - Learn about how to configure your Amazon EC2 gateway instance to accept incoming metadata requests that use IMDS Version 1 (IMDSv1) or require that all metadata requests use IMDS Version 2 (IMDSv2).
- [Synchronize VM time with Hyper-V or Linux KVM host time](#) - Learn about how to view and synchronize the time of an on-premises Hyper-V or Linux KVM gateway virtual machine to a Network Time Protocol (NTP) server.
- [Synchronize VM time with VMware host time](#) - Learn about how to check the host time for a VMware gateway virtual machine and, if needed, set the time and configure the host to synchronize its time automatically to a Network Time Protocol (NTP) server.
- [Configuring paravirtualization on a VMware host](#) - Learn about how you can configure the VMware host platform for your Storage Gateway appliance to use paravirtual Internet Small Computer System Interface Protocol (iSCSI) controllers.
- [Configuring network adapters for your gateway](#) - Learn about how you can reconfigure your gateway to use the VMXNET3 (10 GbE) network adapter, or to use more than one network adapter so that it can be accessed from multiple IP addresses.
- [Using VMware vSphere High Availability with Storage Gateway](#) - Learn about how to protect your storage workloads against hardware, hypervisor, or network failures by configuring Storage Gateway to work with VMware vSphere High Availability.

Deploy a default Amazon EC2 host for Volume Gateway

This topic lists the steps to deploy an Amazon EC2 host using the default specifications.

You can deploy and activate a Volume Gateway on an Amazon Elastic Compute Cloud (Amazon EC2) instance. The AWS Storage Gateway Amazon Machine Image (AMI) is available as a community AMI.

 **Note**

Storage Gateway community AMIs are published and fully supported by AWS. You can see that the publisher is AWS, a verified provider.

1. To set up the Amazon EC2 instance, choose **Amazon EC2** as the **Host platform** in the **Platform options** section of the workflow. For instructions on configuring the Amazon EC2 instance, see [Deploying an Amazon EC2 instance to host your Volume Gateway](#).
2. Select **Launch instance** to open the AWS Storage Gateway AMI template in the Amazon EC2 console and customize additional settings such as **Instance types**, **Network settings** and **Configure storage**.
3. Optionally, you can select **Use default settings** in the Storage Gateway console to deploy an Amazon EC2 instance with the default configuration.

The Amazon EC2 instance that **Use default settings** creates has the following default specifications:

- **Instance type** — *m5.xlarge*
- **Network Settings**
 - For **VPC**, select the VPC that you want your EC2 instance to run in.
 - For **Subnet**, specify the subnet that your EC2 instance should be launched in.

 **Note**

VPC subnets will appear in the drop down only if they have the auto-assign public IPv4 address setting activated from the VPC management console.

- **Auto-assign Public IP** — *Activated*

An EC2 security group is created and associated with the EC2 instance. The security group has the following inbound port rules:

Note

You will need Port 80 open during gateway activation. The port is closed immediately following activation. Thereafter, your EC2 instance can only be accessed over the other ports from the selected VPC.

The iSCSI targets on your gateway are only accessible from the hosts in the same VPC as the gateway. If the iSCSI targets need to be accessed from hosts outside of the VPC, you should update the appropriate security group rules.

You can edit security groups at any time by navigating to the Amazon EC2 instance details page, selecting **Security**, navigating to **Security group details**, and choosing the security group ID.

Port	Protocol	File System Protocol				
80	TCP	HTTP access for activation				
3260	TCP	iSCSI				

- **Configure storage**

Default Settings	AMI Root Volume	Volume 2 Cache	Volume 3 Cache			
Device Name		'/dev/sdb'	'/dev/sdc'			
Size	80 Gib	165 GiB	150 GiB			
Volume Type	gp3	gp3	gp3			
IOPS	3000	3000	3000			

Default Settings	AMI Root Volume	Volume 2 Cache	Volume 3 Cache			
Delete on termination	Yes	Yes	Yes			
Encrypted	No	No	No			
Throughput	125	125	125			

Deploy a customized Amazon EC2 instance for Volume Gateway

You can deploy and activate a Volume Gateway on an Amazon Elastic Compute Cloud (Amazon EC2) instance. The AWS Storage Gateway Amazon Machine Image (AMI) is available as a community AMI.

Note

Storage Gateway community AMIs are published and fully supported by AWS. You can see that the publisher is AWS, a verified provider.

Volume Gateway AMIs use the following naming convention. The version number appended to the AMI name changes with each version release.

`aws-storage-gateway-CLASSIC-2.9.0`

To deploy an Amazon EC2 instance to host your Volume Gateway

1. Start setting up a new gateway using the Storage Gateway console. For instructions, see [Set up a Volume Gateway](#). When you reach the **Platform options** section, choose **Amazon EC2** as the **Host platform**, then use the following steps to launch the Amazon EC2 instance that will host your Volume Gateway.

Note

The Amazon EC2 host platform supports **Cached volumes** only. Stored volume gateways cannot be deployed on EC2 instances.

2. Choose **Launch instance** to open the AWS Storage Gateway AMI template in the Amazon EC2 console, where you can configure additional settings.

Use **Quicklaunch** to launch the Amazon EC2 instance with default settings. For more information on Amazon EC2 Quicklaunch default specifications, see [Quicklaunch Configuration Specifications for Amazon EC2](#).

3. For **Name**, enter a name for the Amazon EC2 instance. After the instance is deployed, you can search for this name to find your instance on list pages in the Amazon EC2 console.
4. In the **Instance type** section, for **Instance type**, choose the hardware configuration for your instance. The hardware configuration must meet certain minimum requirements to support your gateway. We recommend starting with the **m5.xlarge** instance type, which meets the minimum hardware requirements for your gateway to function properly. For more information, see [Requirements for Amazon EC2 instance types](#).


You can resize your instance after you launch, if necessary. For more information, see [Resizing your instance](#) in the *Amazon EC2 User Guide*.

Note

Certain instance types, particularly i3 EC2, use NVMe SSD disks. These can cause problems when you start or stop Volume Gateway; for example, you can lose data from the cache. Monitor the `CachePercentDirty` Amazon CloudWatch metric, and only start or stop your system when that parameter is 0. To learn more about monitoring metrics for your gateway, see [Storage Gateway metrics and dimensions](#) in the CloudWatch documentation.

5. In the **Key pair (login)** section, for **Key pair name - required**, select the key pair you want to use to securely connect to your instance. You can create a new key pair if necessary. For more information, see [Create a key pair](#) in the *Amazon Elastic Compute Cloud User Guide for Linux Instances*.

6. In the **Network settings** section, review the preconfigured settings and choose **Edit** to make changes to the following fields:
 - a. For **VPC - required**, choose the VPC where you want to launch your Amazon EC2 instance. For more information, see [How Amazon VPC works](#) in the *Amazon Virtual Private Cloud User Guide*.
 - b. (Optional) For **Subnet**, choose the subnet where you want to launch your Amazon EC2 instance.
 - c. For **Auto-assign Public IP**, choose **Enable**.
7. In the **Firewall (security groups)** subsection, review the preconfigured settings. You can change the default name and description of the new security group to be created for your Amazon EC2 instance if you want, or choose to apply firewall rules from an existing security group instead.
8. In the **Inbound security groups rules** subsection, add firewall rules to open the ports that clients will use to connect to your instance. For more information on the ports required for Volume Gateway, see [Port requirements](#). For more information on adding firewall rules, see [Security group rules](#) in the *Amazon Elastic Compute Cloud User Guide for Linux Instances*.

 **Note**

Volume Gateway requires TCP port 80 to be open for inbound traffic and for one-time HTTP access during gateway activation. After activation, you can close this port. Additionally, you must open TCP port 3260 for iSCSI access.

9. In the **Advanced network configuration** subsection, review the preconfigured settings and make changes if necessary.
10. In the **Configure storage** section, choose **Add new volume** to add storage to your gateway instance.

 **Important**

You must add at least one Amazon EBS volume with at least **165 GiB** capacity for cache storage, and at least one Amazon EBS volume with at least **150 GiB** capacity for upload buffer, in addition to the preconfigured **Root volume**. For increased performance, we recommend allocating multiple EBS volumes for cache storage with at least 150 GiB each.

11. In the **Advanced details** section, review the preconfigured settings and make changes if necessary.
12. Choose **Launch instance** to launch your new Amazon EC2 gateway instance with the configured settings.
13. To verify that your new instance launched successfully, navigate to the **Instances** page in the Amazon EC2 console and search for your new instance by name. Ensure that that **Instance state** displays **Running** with a green check mark, and that the **Status check** is complete, and shows a green check mark.
14. Select your instance from the details page. Copy the **Public IPv4 address** from the **Instance summary** section, then return to the **Set up gateway** page in the Storage Gateway console to resume setting up your Volume Gateway.

You can determine the AMI ID to use for launching a Volume Gateway by using the Storage Gateway console or by querying the AWS Systems Manager parameter store.

To determine the AMI ID, do one of the following:

- Start setting up a new gateway using the Storage Gateway console. For instructions, see [Set up a Volume Gateway](#). When you reach the **Platform options** section, choose **Amazon EC2** as the **Host platform**, then choose **Launch instance** to open the AWS Storage Gateway AMI template in the Amazon EC2 console.

You are redirected to the EC2 community AMI page, where you can see the AMI ID for your AWS Region in the URL.

- Query the Systems Manager parameter store. You can use the AWS CLI or Storage Gateway API to query the Systems Manager public parameter under the namespace `/aws/service/storagegateway/ami/CACHED/latest` for Cached Volume Gateways or `/aws/service/storagegateway/ami/STORED/latest` for Stored Volume Gateways. For example, using the following CLI command returns the ID of the current AMI in the AWS Region you specify.

```
aws --region us-east-2 ssm get-parameter --name /aws/service/storagegateway/ami/STORED/latest
```

The CLI command returns output similar to the following.

```
{  
  "Parameter": {
```

```
    "Type": "String",
    "LastModifiedDate": 1561054105.083,
    "Version": 4,
    "ARN": "arn:aws:ssm:us-east-2::parameter/aws/service/storagegateway/ami/
STORED/latest",
    "Name": "/aws/service/storagegateway/ami/STORED/latest",
    "Value": "ami-123c45dd67d891000"
  }
}
```

Modify Amazon EC2 instance metadata options

The instance metadata service (IMDS) is an on-instance component that provides secure access to Amazon EC2 instance metadata. An instance can be configured to accept incoming metadata requests that use IMDS Version 1 (IMDSv1) or require that all metadata requests use IMDS Version 2 (IMDSv2). IMDSv2 uses session-oriented requests and mitigates several types of vulnerabilities that could be used to try to access the IMDS. For information about IMDSv2, see [How Instance Metadata Service Version 2 works](#) in the *Amazon Elastic Compute Cloud User Guide*.

We recommend that you require IMDSv2 for all Amazon EC2 instances that host Storage Gateway. IMDSv2 is required by default on all newly launched gateway instances. If you have existing instances that are still configured to accept IMDSv1 metadata requests, see [Require the use of IMDSv2](#) in the *Amazon Elastic Compute Cloud User Guide* for instructions to modify your instance metadata options to require the use of IMDSv2. Applying this change does not require an instance reboot.

Synchronize VM time with Hyper-V or Linux KVM host time

For a gateway deployed on VMware ESXi, setting the hypervisor host time and synchronizing the virtual machine time to the host is sufficient to avoid time drift. For more information, see [Synchronize VM time with VMware host time](#). For a gateway deployed on Microsoft Hyper-V or Linux KVM, we recommend that you periodically check the virtual machine time using the procedure described following.

To view and synchronize the time of a hypervisor gateway virtual machine to a Network Time Protocol (NTP) server

1. Log in to your gateway's local console:

- For more information on logging in to the Microsoft Hyper-V local console, see [Access the Gateway Local Console with Microsoft Hyper-V](#).
 - For more information on logging in to the local console for Linux Kernel-based Virtual Machine (KVM), see [Accessing the Gateway Local Console with Linux KVM](#).
2. On the **Storage Gateway Configuration** main menu screen, enter the corresponding numeral to select **System Time Management**.
 3. On the **System Time Management** menu screen, enter the corresponding numeral to select **View and Synchronize System Time**.

The gateway local console displays the current system time and compares it with the time reported by the NTP server, then reports the exact discrepancy between the two times in seconds.

4. If the time discrepancy is greater than 60 seconds, enter **y** to synchronize the system time with NTP time. Otherwise, enter **n**.

Time synchronization might take a few moments.

Synchronize VM time with VMware host time

To successfully activate your gateway, you must ensure that your VM time is synchronized to the host time, and that the host time is correctly set. In this section, 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.

Important

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

To synchronize VM time with host time

1. Configure your VM time.
 - a. In the vSphere client, right-click on the name of your gateway VM in panel on the left side of the application window to open the context menu for the VM, and then choose **Edit Settings**.

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

- b. Choose the **Options** tab, and then choose **VMware Tools** from the options list.
- c. Check the **Synchronize guest time with host** option in the **Advanced** section on the right side of the **Virtual Machine Properties** dialog box, and then choose **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, perform the following steps to set and synchronize it with an NTP server.

- a. In the VMware vSphere client, select the vSphere host node in the left panel, and then choose the **Configuration** tab.
- b. Select **Time Configuration** in the **Software** panel, and then choose the **Properties** link.

The **Time Configuration** dialog box appears.

- c. Under **Date and Time**, set the date and time for your vSphere host.
- d. Configure the host to synchronize its time automatically to an NTP server.
 - i. Choose **Options** in the **Time Configuration** dialog box, and then in the **NTP Daemon (ntpd) Options** dialog box, choose **NTP Settings** in the left panel.
 - ii. Choose **Add** to add a new NTP server.
 - iii. In the **Add NTP Server** dialog box, type the IP address or the fully qualified domain name of an NTP server, and then choose **OK**.

You can use `pool.ntp.org` as the domain name.

- iv. In the **NTP Daemon (ntpd) Options** dialog box, choose **General** in the left panel.
- v. Under **Service Commands**, choose **Start** to start the service.

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

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

Configuring paravirtualization on a VMware host

The following procedure describes how to configure the VMware host platform for your Storage Gateway appliance to use paravirtual Internet Small Computer System Interface Protocol (iSCSI) controllers. Paravirtual iSCSI controllers are high performance storage controllers that can result in greater throughput and lower CPU use. These controllers are best suited for high performance storage environments. When you configure iSCSI controllers this way, the Storage Gateway virtual machine works with the host operating system to allow the gateway console to identify the virtual disks that you add to your virtual machine.

Note

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

To configure your VMware host platform to use paravirtualized controllers

1. In the VMware vSphere client, right-click on the name of your gateway virtual machine in the navigation pane on the left side of the application window to open the context menu, and then choose **Edit Settings**.
2. In the **Virtual Machine Properties** dialog box, choose the **Hardware** tab.
3. On the **Hardware** tab, select **SCSI controller 0**, and then choose **Change Type**.
4. In the **Change SCSI Controller Type** dialog box, select the **VMware Paravirtual** SCSI controller type, and then choose **OK** to save the configuration.

Configuring network adapters for your gateway

By default, Storage Gateway is configured to use the E1000 network adapter type, but you can reconfigure your gateway to use the VMXNET3 (10 GbE) network adapter. You can also configure Storage Gateway so it can be accessed by more than one IP address. You do this by configuring your gateway to use more than one network adapter.

Topics

- [Configuring Your Gateway to Use the VMXNET3 Network Adapter](#)
- [Configuring Your Gateway for Multiple NICs](#)

Configuring Your Gateway to Use the VMXNET3 Network Adapter

Storage Gateway supports the E1000 network adapter type in both VMware ESXi and Microsoft Hyper-V hypervisor hosts. However, the VMXNET3 (10 GbE) network adapter type is supported in VMware ESXi hypervisor only. If your gateway is hosted on a VMware ESXi hypervisor, you can reconfigure your gateway to use the VMXNET3 (10 GbE) adapter type. For more information on these adapters, see [Choosing a network adapter for your virtual machine](#) on the Broadcom (VMware) website.

Important

To select VMXNET3, your guest operating system type must be **Other Linux64**.

Following are the steps you take to configure your gateway to use the VMXNET3 adapter:

1. Remove the default E1000 adapter.
2. Add the VMXNET3 adapter.
3. Restart your gateway.
4. Configure the adapter for the network.

Details on how to perform each step follow.

To remove the default E1000 adapter and configure your gateway to use the VMXNET3 adapter

1. In VMware, open the context (right-click) menu for your gateway and choose **Edit Settings**.
2. In the **Virtual Machine Properties** window, choose the **Hardware** tab.
3. For **Hardware**, choose **Network adapter**. Notice that the current adapter is E1000 in the **Adapter Type** section. You will replace this adapter with the VMXNET3 adapter.
4. Choose the E1000 network adapter, and then choose **Remove**. In this example, the E1000 network adapter is **Network adapter 1**.

Note

Although you can run the E1000 and VMXNET3 network adapters in your gateway at the same time, we don't recommend doing so because it can cause network problems.

5. Choose **Add** to open the Add Hardware wizard.
6. Choose **Ethernet Adapter**, and then choose **Next**.
7. In the Network Type wizard, select **VMXNET3** for **Adapter Type**, and then choose **Next**.
8. In the Virtual Machine properties wizard, verify in the **Adapter Type** section that **Current Adapter** is set to **VMXNET3**, and then choose **OK**.
9. In the VMware vSphere client, shut down your gateway.
10. In the VMware vSphere client, restart your gateway.

After your gateway restarts, reconfigure the adapter you just added to make sure that network connectivity to the internet is established.

To configure the adapter for the network

1. In the vSphere client, choose the **Console** tab to start the local console. Use the default login credentials to log in to the gateway's local console for this configuration task. For information about how to log in using the default credentials, see [Logging in to the Local Console Using Default Credentials](#).
2. At the prompt, enter the corresponding numeral to select **Network Configuration**.
3. At the prompt, enter the corresponding numeral to select **Reset all to DHCP**, and then enter **y** (for yes) at the prompt to set all adapters to use Dynamic Host Configuration Protocol (DHCP). All available adapters are set to use DHCP.

If your gateway is already activated, you must shut it down and restart it from the Storage Gateway Management Console. After the gateway restarts, you must test network connectivity to the internet. For information about how to test network connectivity, see [Testing Your Gateway Connection to the Internet](#).

Configuring Your Gateway for Multiple NICs

If you configure your gateway to use multiple network adapters (NICs), it can be accessed by more than one IP address. You might want to do this in the following situations:

- **Maximizing throughput** – You might want to maximize throughput to a gateway when network adapters are a bottleneck.

- **Application separation** – You might need to separate how your applications write to a gateway's volumes. For example, you might choose to have a critical storage application exclusively use one particular adapter defined for your gateway.
- **Network constraints** – Your application environment might require that you keep your iSCSI targets and the initiators that connect to them in an isolated network that is different from the network by which the gateway communicates with AWS.

In a typical multiple-adapter use case, one adapter is configured as the route by which the gateway communicates with AWS (that is, as the default gateway). Except for this one adapter, initiators must be in the same subnet as the adapter that contains the iSCSI targets to which they connect. Otherwise, communication with the intended targets might not be possible. If a target is configured on the same adapter that is used for communication with AWS, then iSCSI traffic for that target and AWS traffic will flow through the same adapter.

When you configure one adapter to connect to the Storage Gateway console and then add a second adapter, Storage Gateway automatically configures the route table to use the second adapter as the preferred route. For instructions on how to configure multiple-adapters, see the following sections.

- [Configuring multiple network adapters on a VMware ESXi host](#)
- [Configuring multiple network adapters on Microsoft Hyper-V host](#)

Configuring multiple network adapters on a VMware ESXi host

The following procedure assumes that your gateway VM already has one network adapter defined, and describes how to add an adapter on VMware ESXi.

To configure your gateway to use an additional network adapter in VMware ESXi host

1. Shut down the gateway.
2. In the VMware vSphere client, select your gateway VM.


The VM can remain turned on for this procedure.

3. In the client, open the context (right-click) menu for your gateway VM, and choose **Edit Settings**.
4. On the **Hardware** tab of the **Virtual Machine Properties** dialog box, choose **Add** to add a device.

5. Follow the Add Hardware wizard to add a network adapter.
 - a. In the **Device Type** pane, choose **Ethernet Adapter** to add an adapter, and then choose **Next**.
 - b. In the **Network Type** pane, ensure that **Connect at power on** is selected for **Type**, and then choose **Next**.

We recommend that you use the VMXNET3 network adapter with Storage Gateway. For more information on the adapter types that might appear in the adapter list, see Network Adapter Types in the [ESXi and vCenter Server Documentation](#).

- c. In the **Ready to Complete** pane, review the information, and then choose **Finish**.
6. Choose the **Summary** tab for the VM, and choose **View All** next to the **IP Address** box. The **Virtual Machine IP Addresses** window displays all the IP addresses you can use to access the gateway. Confirm that a second IP address is listed for the gateway.

 **Note**

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

7. In the Storage Gateway console, turn on the gateway.
8. In the **Navigation** pane of the Storage Gateway console, choose **Gateways** and choose the gateway to which you added the adapter. Confirm that the second IP address is listed in the **Details** tab.

For information about local console tasks common to VMware, Hyper-V, and KVM hosts, see [Performing Tasks on the VM Local Console](#)

Configuring multiple network adapters on Microsoft Hyper-V host

The following procedure assumes that your gateway VM already has one network adapter defined and that you are adding a second adapter. This procedure shows how to add an adapter for a Microsoft Hyper-V host.

To configure your gateway to use an additional network adapter in a Microsoft Hyper-V Host

1. On the Storage Gateway console, turn off the gateway.
2. In the Microsoft Hyper-V Manager, select your gateway VM from the **Virtual Machines** panel.

3. If the gateway VM isn't turned off already, right-click the VM name to open the context menu, and then choose **Turn Off**.
4. Right-click the gateway VM name to open the context menu, and then choose **Settings**.
5. In the **Settings** dialog box, under **Hardware**, choose **Add Hardware**.
6. In the **Add Hardware** panel on the right side of the **Settings** dialog box, choose **Network Adapter**, and then choose **Add** to add a device.
7. Configure the network adapter, and then choose **Apply** to apply settings.
8. In the **Settings** dialog box, under **Hardware**, confirm that the new network adapter was added to the hardware list, and then choose **OK**.
9. Turn on the gateway using the Storage Gateway console.
10. In the **Navigation** panel of the Storage Gateway console, choose **Gateways**, then select the gateway to which you added the adapter. Confirm that a second IP address is listed in the **Details** tab.

For information about local console tasks common to VMware, Hyper-V, and KVM hosts, see [Performing Tasks on the VM Local Console](#)

Using VMware vSphere High Availability with Storage Gateway

Storage Gateway provides high availability on VMware through a set of application-level health checks integrated with VMware vSphere High Availability (VMware HA). This approach helps protect storage workloads against hardware, hypervisor, or network failures. It also helps protect against software errors, such as connection timeouts and file share or volume unavailability.

vSphere HA works by pooling virtual machines and the hosts they reside on into a cluster for redundancy. Hosts in the cluster are monitored and in the event of a failure, the virtual machines on a failed host are restarted on alternate hosts. Generally, this recovery happens quickly and without data loss. For more information about vSphere HA, see [How vSphere HA Works](#) in the VMware documentation.

Note

The time required to restart a failed virtual machine and re-establish the iSCSI connection on a new host depends on many factors, such as the host operating system and resource load, disk speed, network connection, and SAN/storage infrastructure. To minimize failover downtime, implement the recommendations outlined in [Optimizing Gateway Performance](#).

To use Storage Gateway with VMware HA, we recommend doing the following things:

- Deploy the VMware ESX .ova downloadable package that contains the Storage Gateway VM on only one host in a cluster.
- When deploying the .ova package, select a data store that is not local to one host. Instead, use a data store that is accessible to all hosts in the cluster. If you select a data store that is local to a host and the host fails, then the data source might not be accessible to other hosts in the cluster and failover to another host might not succeed.
- To prevent your initiator from disconnecting from storage volume targets during failover, follow the recommended iSCSI settings for your operating system. In a failover event, it can take from 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 both Windows and Linux clients are greater than the typical time it takes for failover to occur. For more information on customizing Windows clients' timeout settings, see [Customizing Your Windows iSCSI Settings](#). For more information on customizing Linux clients' timeout settings, see [Customizing Your Linux iSCSI Settings](#).
- With clustering, if you deploy the .ova package to the cluster, select a host when you are prompted to do so. Alternately, you can deploy directly to a host in a cluster.

The following topics describe how to deploy Storage Gateway in a VMware HA cluster:

Topics

- [Configure Your vSphere VMware HA Cluster](#)
- [Download the .ova Image from the Storage Gateway console](#)
- [Deploy the Gateway](#)
- [\(Optional\) Add Override Options for Other VMs on Your Cluster](#)
- [Activate Your Gateway](#)
- [Test Your VMware High Availability Configuration](#)

Configure Your vSphere VMware HA Cluster

First, if you haven't already created a VMware cluster, create one. For information about how to create a VMware cluster, see [Create a vSphere HA Cluster](#) in the VMware documentation.

Next, configure your VMware cluster to work with Storage Gateway.

To configure your VMware cluster

1. On the **Edit Cluster Settings** page in VMware vSphere, make sure that VM monitoring is configured for VM and application monitoring. To do so, set the following values for each option:
 - **Host Failure Response: Restart VMs**
 - **Response for Host Isolation: Shut down and restart VMs**
 - **Datastore with PDL: Disabled**
 - **Datastore with APD: Disabled**
 - **VM Monitoring: VM and Application Monitoring**
2. Fine-tune the sensitivity of the cluster by adjusting the following values:
 - **Failure interval** – After this interval, the VM is restarted if a VM heartbeat isn't received.
 - **Minimum uptime** – The cluster waits this long after a VM starts to begin monitoring for VM tools' heartbeats.
 - **Maximum per-VM resets** – The cluster restarts the VM a maximum of this many times within the maximum resets time window.
 - **Maximum resets time window** – The window of time in which to count the maximum resets per-VM resets.

If you aren't sure what values to set, use these example settings:

- **Failure interval: 30** seconds
- **Minimum uptime: 120** seconds
- **Maximum per-VM resets: 3**
- **Maximum resets time window: 1** hour

If you have other VMs running on the cluster, you might want to set these values specifically for your VM. You can't do this until you deploy the VM from the .ova. For more information on setting these values, see [\(Optional\) Add Override Options for Other VMs on Your Cluster](#).

Download the .ova Image from the Storage Gateway console

To download the .ova image for your gateway

- On the **Set up gateway** page in the Storage Gateway console, select your gateway type and host platform, then use the link provided in the console to download the .ova as outlined in [Set up a Volume Gateway](#).

Deploy the Gateway

In your configured cluster, deploy the .ova image to one of the cluster's hosts.

To deploy the gateway .ova image

1. Deploy the .ova image to one of the hosts in the cluster.
2. Make sure the data stores that you choose for the root disk and the cache are available to all hosts in the cluster. When deploying the Storage Gateway .ova file in a VMware or on-prem environment, the disks are described as paravirtualized SCSI disks. *Paravirtualization* is a mode where the gateway VM works with the host operating system so the console can identify the virtual disks that you add to your VM.

To configure your VM to use paravirtualized controllers

1. In the VMware vSphere client, open the context (right-click) menu for your gateway VM, and then choose **Edit Settings**.
2. In the **Virtual Machine Properties** dialog box, choose the **Hardware** tab, select the **SCSI controller 0**, and then choose **Change Type**.
3. In the **Change SCSI Controller Type** dialog box, select the **VMware Paravirtual SCSI** controller type, and then choose **OK**.

(Optional) Add Override Options for Other VMs on Your Cluster

If you have other VMs running on your cluster, you might want to set the cluster values specifically for each VM. For instructions, see [Customize an Individual Virtual Machine](#) in the VMware vSphere online documentation.

To add override options for other VMs on your cluster

1. On the **Summary** page in VMware vSphere, choose your cluster to open the cluster page, and then choose **Configure**.
2. Choose the **Configuration** tab, and then choose **VM Overrides**.
3. Add a new VM override option to change each value.

Set the following values for each option under **vSphere HA - VM Monitoring**:

- **VM Monitoring: Override Enabled - VM and Application Monitoring**
- **VM monitoring sensitivity: Override Enabled - VM and Application Monitoring**
- **VM Monitoring: Custom**
- **Failure interval: 30 seconds**
- **Minimum uptime: 120 seconds**
- **Maximum per-VM resets: 5**
- **Maximum resets time window: Within 1 hrs**

Activate Your Gateway

After the .ova for your gateway is deployed, activate your gateway. The instructions about how are different for each gateway type.

To activate your gateway

- Follow the procedures outlined in the following topics:
 - a. [Connect your Volume Gateway to AWS](#)
 - b. [Review settings and activate your Volume Gateway](#)
 - c. [Configure your Volume Gateway](#)

Test Your VMware High Availability Configuration

After you activate your gateway, test your configuration.

To test your VMware HA configuration

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. On the navigation pane, choose **Gateways**, and then choose the gateway that you want to test for VMware HA.
3. For **Actions**, choose **Verify VMware HA**.
4. In the **Verify VMware High Availability Configuration** box that appears, choose **OK**.

Note

Testing your VMware HA configuration reboots your gateway VM and interrupts connectivity to your gateway. The test might take a few minutes to complete.

If the test is successful, the status of **Verified** appears in the details tab of the gateway in the console.

5. Choose **Exit**.

You can find information about VMware HA events in the Amazon CloudWatch log groups. For more information, see [Getting Volume Gateway Health Logs with CloudWatch Log Groups](#).

Working with Volume Gateway storage resources

The topics in this section describe how you can manage the storage resources that are associated with your Volume Gateway appliance and its virtual host platform. This includes resources such as the physical disks attached to a gateway's hypervisor host platform, with specific procedures for removing disks from VMware vSphere ESXi, Microsoft Hyper-V, or Linux Kernel-based Virtual Machine (KVM) virtualization hosts. This also includes managing the Amazon EBS volumes attached to a gateway's Amazon EC2 instance for gateways hosted on Amazon EC2 in the AWS cloud.

Topics

- [Removing Disks from Your Gateway](#) - Learn about what to do if you need to remove a disk from the VMware vSphere ESXi, Microsoft Hyper-V, or Linux Kernel-based Virtual Machine (KVM) virtualization host platform for your gateway, for example if you have a physical disk failure.

- [Managing Amazon EBS volumes on Amazon EC2 gateways](#) - Learn about how you can increase or reduce the quantity of Amazon EBS volumes that are allocated for use as upload buffer or cache storage for a gateway that is hosted on an Amazon EC2 instance, for example, if your application storage needs increase or decrease over time.

Removing Disks from Your Gateway

Although we don't recommend removing the underlying disks from your gateway, you might want to remove a disk from your gateway, for example if you have a failed disk.

Removing a Disk from a Gateway Hosted on VMware ESXi

You can use the following procedure to remove a disk from your gateway hosted on VMware hypervisor.

To remove a disk allocated for the upload buffer (VMware ESXi)

1. In the vSphere client, open the context (right-click) menu, choose the name of your gateway VM, and then choose **Edit Settings**.
2. On the **Hardware** tab of the **Virtual Machine Properties** dialog box, select the disk allocated as upload buffer space, and then choose **Remove**.

Verify that the **Virtual Device Node** value in the **Virtual Machine Properties** dialog box has the same value that you noted previously. Doing this helps ensure that you remove the correct disk.

3. Choose an option in the **Removal Options** panel, and then choose **OK** to complete the process of removing the disk.

Removing a Disk from a Gateway Hosted on Microsoft Hyper-V

Using the following procedure, you can remove a disk from your gateway hosted on a Microsoft Hyper-V hypervisor.

To remove an underlying disk allocated for the upload buffer (Microsoft Hyper-V)

1. In the Microsoft Hyper-V Manager, open the context (right-click) menu, choose the name of your gateway VM, and then choose **Settings**.

2. In the **Hardware** list of the **Settings** dialog box, select the disk to remove, and then choose **Remove**.

The disks you add to a gateway appear under the **SCSI Controller** entry in the **Hardware** list. Verify that the **Controller** and **Location** value are the same value that you noted previously. Doing this helps ensure that you remove the correct disk.

The first SCSI controller displayed in the Microsoft Hyper-V Manager is controller 0.

3. Choose **OK** to apply the change.

Removing a Disk from a Gateway Hosted on Linux KVM

To detach a disk from your gateway hosted on Linux Kernel-based Virtual Machine (KVM) hypervisor, you can use a `virsh` command similar to the one following.

```
$ virsh detach-disk domain_name /device/path
```

For more details about managing KVM disks, see documentation of your Linux distribution.

Managing Amazon EBS volumes on Amazon EC2 gateways

When you initially configured your gateway to run as an Amazon EC2 instance, you allocated Amazon EBS volumes for use as an upload buffer and cache storage. Over time, as your applications needs change, you can allocate additional Amazon EBS volumes for this use. You can also reduce the storage you allocated by removing previously allocated Amazon EBS volumes. For more information about Amazon EBS, see [Amazon Elastic Block Store \(Amazon EBS\)](#) in the *Amazon EC2 User Guide*.

Before you add more storage to the gateway, you should review how to size your upload buffer and cache storage based on your application needs for a gateway. To do so, see [Determining the size of upload buffer to allocate](#) and [Determining the size of cache storage to allocate](#).

There are quotas on the maximum storage you can allocate as an upload buffer and cache storage. You can attach as many Amazon EBS volumes to your instance as you want, but you can only configure these volumes as upload buffer and cache storage space up to these storage quotas. For more information, see [AWS Storage Gateway quotas](#).

To add an Amazon EBS volume and configure it for your gateway

1. Create an Amazon EBS volume. For instructions, see [Creating or Restoring an Amazon EBS Volume](#) in the *Amazon EC2 User Guide*.
2. Attach the Amazon EBS volume to your Amazon EC2 instance. For instructions, see [Attaching an Amazon EBS Volume to an Instance](#) in the *Amazon EC2 User Guide*.
3. Configure the Amazon EBS volume you added as either an upload buffer or cache storage. For instructions, see [Managing local disks for your Storage Gateway](#).

There are times you might find you don't need the amount of storage you allocated for the upload buffer.

To remove an Amazon EBS volume

Warning

These steps apply only for Amazon EBS volumes allocated as upload buffer space, not for volumes allocated to cache.

1. Shut down the gateway by following the approach described in the [Shutting Down Your Gateway VM](#) section.
2. Detach the Amazon EBS volume from your Amazon EC2 instance. For instructions, see [Detaching an Amazon EBS Volume from an Instance](#) in the *Amazon EC2 User Guide*.
3. Delete the Amazon EBS volume. For instructions, see [Deleting an Amazon EBS Volume](#) in the *Amazon EC2 User Guide*.
4. Start the gateway by following the approach described in the [Shutting Down Your Gateway VM](#) section.

Getting an activation key for your gateway

To receive an activation key for your gateway, make a web request to the gateway virtual machine (VM). The VM returns a redirect that contains the activation key, which is passed as one of the parameters for the `ActivateGateway` API action to specify the configuration of your gateway. For more information, see [ActivateGateway](#) in the *Storage Gateway API Reference*.

Note

Gateway activation keys expire in 30 minutes if unused.

The request that you make to the gateway VM includes the AWS Region where the activation occurs. The URL that's returned by the redirect in the response contains a query string parameter called `activationkey`. This query string parameter is your activation key. The format of the query string looks like the following: `http://gateway_ip_address/?activationRegion=activation_region`. The output of this query returns both activation region and key.

The URL also includes `vpcEndpoint`, the VPC Endpoint ID for gateways that connect using the VPC endpoint type.

Note

The Storage Gateway Hardware Appliance, VM image templates, and Amazon EC2 Amazon Machine Images (AMI) come preconfigured with the HTTP services necessary to receive and respond to the web requests described on this page. It's not required or recommended to install any additional services on your gateway.

Topics

- [Linux \(curl\)](#)
- [Linux \(bash/zsh\)](#)
- [Microsoft Windows PowerShell](#)
- [Using your local console](#)

Linux (curl)

The following examples show you how to get an activation key using Linux (curl).

Note

Replace the highlighted variables with actual values for your gateway. Acceptable values are as follows:

- *gateway_ip_address* - The IPv4 address of your gateway, for example 172.31.29.201
- *gateway_type* - The type of gateway you want to activate, such as STORED, CACHED, VTL, FILE_S3, or FILE_FSX_SMB.
- *region_code* - The Region where you want to activate your gateway. See [Regional endpoints](#) in the *AWS General Reference Guide*. If this parameter is not specified, or if the value provided is misspelled or doesn't match a valid region, the command will default to the us-east-1 region.
- *vpc_endpoint* - The VPC endpoint name for your gateway, for example vpce-050f90485f28f2fd0-iep0e8vq.storagegateway.us-west-2.vpce.amazonaws.com.

To get the activation key for a public endpoint:

```
curl "http://gateway_ip_address?activationRegion=region_code&no_redirect"
```

To get the activation key for a VPC endpoint:

```
curl "http://gateway_ip_address?  
activationRegion=region_code&vpcEndpoint=vpc_endpoint&no_redirect"
```

Linux (bash/zsh)

The following example shows you how to use Linux (bash/zsh) to fetch the HTTP response, parse HTTP headers, and get the activation key.

```
function get-activation-key() {  
  local ip_address=$1  
  local activation_region=$2  
  if [[ -z "$ip_address" || -z "$activation_region" || -z "$gateway_type" ]]; then  
    echo "Usage: get-activation-key ip_address activation_region gateway_type"  
    return 1  
  fi  
}
```

```

if redirect_url=$(curl -f -s -S -w '%{redirect_url}' "http://$ip_address/?
activationRegion=$activation_region&gatewayType=$gateway_type"); then
    activation_key_param=$(echo "$redirect_url" | grep -oE 'activationKey=[A-Z0-9-]+')
    echo "$activation_key_param" | cut -f2 -d=
else
    return 1
fi
}

```

Microsoft Windows PowerShell

The following example shows you how to use Microsoft Windows PowerShell to fetch the HTTP response, parse HTTP headers, and get the activation key.

```

function Get-ActivationKey {
    [CmdletBinding()]
    Param(
        [parameter(Mandatory=$true)][string]$IpAddress,
        [parameter(Mandatory=$true)][string]$ActivationRegion,
        [parameter(Mandatory=$true)][string]$GatewayType
    )
    PROCESS {
        $request = Invoke-WebRequest -UseBasicParsing -Uri "http://$IpAddress/?
activationRegion=$ActivationRegion&gatewayType=$GatewayType" -MaximumRedirection 0 -
ErrorAction SilentlyContinue
        if ($request) {
            $activationKeyParam = $request.Headers.Location | Select-String -Pattern
"activationKey=([A-Z0-9-]+)"
            $activationKeyParam.Matches.Value.Split("=")[1]
        }
    }
}

```

Using your local console

The following example shows you how to use your local console to generate and display an activation key.

To get an activation key for your gateway from your local console

1. Log in to your local console. If you are connecting to your Amazon EC2 instance from a Windows computer, log in as *admin*.

2. After you log in and see the **AWS Appliance Activation - Configuration** main menu, select **0** to choose **Get activation key**.
3. Select **Storage Gateway** for gateway family option.
4. When prompted, enter the AWS Region where you want to activate your gateway.
5. Enter 1 for Public or 2 for VPC endpoint as the network type.
6. Enter 1 for Standard or 2 for Federal Information Processing Standard (FIPS) as the endpoint Type.

Connecting iSCSI Initiators

When managing your gateway, you work with volumes or virtual tape library (VTL) devices that are exposed as Internet Small Computer System Interface (iSCSI) targets. For Volume Gateways, the iSCSI targets are volumes. For Tape Gateways, the targets are VTL devices. As part of this work, you do such tasks as connecting to those targets, customizing iSCSI settings, connecting from a Red Hat Linux client, and configuring Challenge-Handshake Authentication Protocol (CHAP).

Topics

- [Connecting to your volumes from a Windows client](#)
- [Connecting your volumes to a Linux client](#)
- [Customizing iSCSI Settings](#)
- [Configuring CHAP Authentication for Your iSCSI Targets](#)

The iSCSI standard is an Internet Protocol (IP)–based storage networking standard for initiating and managing connections between IP-based storage devices and clients. The following list defines some of the terms that are used to describe the iSCSI connection and the components involved.

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. 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 volumes is exposed as an iSCSI target. Connect only one iSCSI initiator to each iSCSI target.

Microsoft iSCSI initiator

The software program on Microsoft Windows computers that allows you to connect a client computer (that is, the computer running the application whose data you want to write to the gateway) to an external iSCSI-based array (that is, the gateway). The connection is made using the host computer's Ethernet network adapter card. The Microsoft iSCSI initiator has been validated with Storage Gateway on Windows 8.1, Windows 10, Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019. The initiator is built into these operating systems.

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 Linux. The package includes a server daemon for the iSCSI protocol.

Each type of gateway can connect to iSCSI devices, and you can customize those connections, as described following.


Connecting to your volumes from a Windows client

A Volume Gateway exposes volumes you have created for the gateway as iSCSI targets. For more information, see [Connecting your volumes to your client](#).

Note

To connect to your volume target, your gateway must have an upload buffer configured. If an upload buffer is not configured for your gateway, then the status of your volumes is displayed as `UPLOAD BUFFER NOT CONFIGURED`. To configure an upload buffer for a gateway in a stored volumes setup, see [To configure additional upload buffer or cache storage for your gateway](#). To configure an upload buffer for a gateway in a cached volumes setup, see [To configure additional upload buffer or cache storage for your gateway](#).

The following diagram highlights the iSCSI target in the larger picture of the Storage Gateway architecture. For more information, see [How Volume Gateway works](#).



You can connect to your volume from either a Windows or Red Hat Linux client. You can optionally configure CHAP for either client type.

Your gateway exposes your volume as an iSCSI target with a name you specify, prepended by `iqn.1997-05.com.amazon:`. For example, if you specify a target name of `myvolume`, then the iSCSI target you use to connect to the volume is `iqn.1997-05.com.amazon:myvolume`. For more information about how to configure your applications to mount volumes over iSCSI, see [Connecting to your volumes from a Windows client](#).

To	See
Connect to your volume from Windows.	Connecting to a Microsoft Windows Client
Connect to your volume from Red Hat Linux.	Connecting to a Red Hat Enterprise Linux Client
Configure CHAP authentication for Windows and Red Hat Linux.	Configuring CHAP Authentication for Your iSCSI Targets

To connect your Windows client to a storage volume

1. On the **Start** menu of your Windows client computer, enter **iscsipl.exe** in the **Search Programs and files** box, locate the iSCSI initiator program, and then run it.

Note

You must have administrator rights on the client computer to run the iSCSI initiator.

2. If prompted, choose **Yes** to start the Microsoft iSCSI initiator service.
3. In the **iSCSI Initiator Properties** dialog box, choose the **Discovery** tab, and then choose **Discover Portal**.
4. In the **Discover Target Portal** dialog box, enter the IP address of your iSCSI target for **IP address or DNS name**, and then choose **OK**. To get the IP address of your gateway, check the **Gateway** tab on the Storage Gateway console. If you deployed your gateway on an Amazon EC2 instance, you can find the public IP or DNS address in the **Description** tab on the Amazon EC2 console.

The IP address now appears in the **Target portals** list on the **Discovery** tab.

Warning

For gateways that are deployed on an Amazon EC2 instance, accessing the gateway over a public internet connection is not supported. The Elastic IP address of the Amazon EC2 instance cannot be used as the target address.

5. Connect the new target portal to the storage volume target on the gateway:

- a. Choose the **Targets** tab.

The new target portal is shown with an inactive status. The target name shown should be the same as the name that you specified for your storage volume in step 1.

- b. Select the target, and then choose **Connect**.

If the target name is not populated already, enter the name of the target as shown in step 1. In the **Connect to Target** dialog box, select **Add this connection to the list of Favorite Targets**, and then choose **OK**.

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

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

 **Note**

Although it is not required for this exercise, we highly recommend that you customize your iSCSI settings for a real-world application as discussed in [Customizing Your Windows iSCSI Settings](#).

Connecting your volumes to a Linux client

When using Red Hat Enterprise Linux (RHEL), you use the `iscsi-initiator-utils` RPM package to connect to your gateway iSCSI targets (volumes or VTL devices).

To connect a Linux client to the iSCSI targets

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 one of the following commands.

For RHEL 5 or 6, use the following command.

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

For RHEL 7, use the following command.

```
sudo service iscsid status
```

- b. If the status command doesn't return a status of *running*, start the daemon using one of the following commands.

For RHEL 5 or 6, use the following command.

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

For RHEL 7, use the following command. For RHEL 7, you usually don't need to explicitly start the `iscsid` service.

```
sudo service iscsid start
```

3. To discover the volume or VTL device targets defined for a gateway, use the following discovery command.

```
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 on the Storage Gateway console.

The output of the discovery command will look like the following example output.

For Volume Gateways: `[GATEWAY_IP]:3260, 1 iqn.1997-05.com.amazon:myvolume`

For Tape Gateways: `iqn.1997-05.com.amazon:[GATEWAY_IP]-tapedrive-01`

Your iSCSI qualified name (IQN) will be different than what is shown preceding, because IQN values are unique to an organization. The name of the target is the name that you specified when you created the volume. You can also find this target name in the **iSCSI Target Info** properties pane when you select a volume on the Storage Gateway console.

4. To connect to a target, use the following command.

Note that you need to specify the correct `[GATEWAY_IP]` and IQN in the connect command.

Warning

For gateways that are deployed on an Amazon EC2 instance, accessing the gateway over a public internet connection is not supported. The Elastic IP address of the Amazon EC2 instance cannot be used as the target address.

```
sudo /sbin/iscsiadm --mode node --targetname  
iqn.1997-05.com.amazon:[ISCSI_TARGET_NAME] --portal [GATEWAY_IP]:3260,1 --login
```

5. To verify that the volume is attached to the client machine (the initiator), use the following command.

```
ls -l /dev/disk/by-path
```

The output of the command will look like the following example output.

```
lrwxrwxrwx. 1 root root 9 Apr 16 19:31 ip-[GATEWAY_IP]:3260-iscsi-  
iqn.1997-05.com.amazon:myvolume-lun-0 -> ../../sda
```

We highly recommend that after you set up your initiator, you customize your iSCSI settings as discussed in [Customizing Your Linux iSCSI Settings](#).

Customizing iSCSI Settings

After you set up your initiator, we highly recommend that you customize your iSCSI settings to prevent the initiator from disconnecting from targets.

By increasing the iSCSI timeout values as shown in the following steps, you make your application better at dealing with write operations 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 of the registry. For information on making a backup copy and other best practices to follow when working with the registry, see [Registry best practices](#) in the *Microsoft TechNet Library*.

Topics

- [Customizing Your Windows iSCSI Settings](#)
- [Customizing Your Linux iSCSI Settings](#)
- [Customizing Your Linux Disk Timeout Settings for Volume Gateways](#)

Customizing Your Windows iSCSI Settings

When using a Windows client, you use the Microsoft iSCSI initiator to connect to your gateway volume. For instructions on how to connect to your volumes, see [Connecting your volumes to your client](#).

1. Connect your Tape Gateway devices to your Windows client.
2. If you are using a backup application, configure the application to use the devices.

To customize your Windows iSCSI settings

1. Increase the maximum time for which requests are queued.
 - a. Start Registry Editor (`Regedit.exe`).
 - b. Navigate to the globally unique identifier (GUID) key for the device class that contains iSCSI controller settings, shown following.

Warning

Make sure that you are working in the **CurrentControlSet** subkey and not another control set, such as **ControlSet001** or **ControlSet002**.

```
HKEY_Local_Machine\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}
```

- c. Find the subkey for the Microsoft iSCSI initiator, shown following as [*<Instance Number>*].

The key is represented by a four-digit number, such as `0000`.

```
HKEY_Local_Machine\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\[<Instance Number]
```

Depending on what is installed on your computer, the Microsoft iSCSI initiator might not be the subkey `0000`. You can ensure that you have selected the correct subkey by verifying that the string `DriverDesc` has the value `Microsoft iSCSI Initiator`.

- d. To show the iSCSI settings, choose the **Parameters** subkey.
- e. Open the context (right-click) menu for the **MaxRequestHoldTime** DWORD (32-bit) value, choose **Modify**, and then change the value to **600**.

MaxRequestHoldTime specifies how many seconds Microsoft iSCSI initiator should hold and retry outstanding commands for, before notifying the upper layer of a `Device Removal` event. This value represents a hold time of 600 seconds.

2. You can increase the maximum amount of data that can be sent in iSCSI packets by modifying the following parameters:
 - **FirstBurstLength** controls the maximum amount of data that can be transmitted in an unsolicited write request. Set this value to **262144** or the Windows OS default, whichever is higher.
 - **MaxBurstLength** is similar to **FirstBurstLength**, but it sets the maximum amount of data that can be transmitted in solicited write sequences. Set this value to **1048576** or the Windows OS default, whichever is higher.
 - **MaxRecvDataSegmentLength** controls the maximum data segment size that is associated with a single protocol data unit (PDU). Set this value to **262144** or the Windows OS default, whichever is higher.

Note

Different backup software can be optimized to work best using different iSCSI settings. To verify which values for these parameters will provide the best performance, see the documentation for your backup software.

3. Increase the disk timeout value, as shown following:
 - a. Start Registry Editor (`Regedit.exe`), if you haven't already.

- b. Navigate to the **Disk** subkey in the **Services** subkey of the **CurrentControlSet**, shown following.

```
HKEY_Local_Machine\SYSTEM\CurrentControlSet\Services\Disk
```

- c. Open the context (right-click) menu for the **TimeOutValue** DWORD (32-bit) value, choose **Modify**, and then change the value to **600**.

TimeOutValue specifies how many seconds iSCSI initiator will wait for a response from the target before it attempts session recovery by dropping and re-establishing the connection. This value represents a timeout period of 600 seconds.

4. To ensure that the new configuration values take effect, restart your system.

Before restarting, you must make sure that the results of all write operations to volumes are flushed. To do this, take any mapped storage volume disks offline before restarting.

Customizing Your Linux iSCSI Settings

After setting up the initiator for your gateway, 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 following, you make your application better at dealing with write operations that take a long time and other transient issues such as network interruptions.

Note

Commands might be slightly different for other types of Linux. The following examples are based on Red Hat Linux.

To customize your Linux iSCSI settings

1. Increase the maximum time for which requests are queued.
 - a. Open the `/etc/iscsi/iscsid.conf` file and find the following lines.


```
node.session.timeo.replacement_timeout = [replacement_timeout_value]
node.conn[0].timeo.noop_out_interval = [noop_out_interval_value]
node.conn[0].timeo.noop_out_timeout = [noop_out_timeout_value]
```

- b. Set the `[replacement_timeout_value]` value to **600**.

Set the `[noop_out_interval_value]` value to **60**.

Set the `[noop_out_timeout_value]` value to **600**.

All three values are in seconds.

 **Note**

The `iscsid.conf` settings must be made before discovering the gateway. If you have already discovered your gateway or logged in to the target, or both, you can delete the entry from the discovery database using the following command. Then you can rediscover or log in again to pick up the new configuration.

```
iscsiadm -m discoverydb -t sendtargets -p [GATEWAY_IP]:3260 -o delete
```

2. Increase the maximum values for the amount of data that can be transmitted in each response.
 - a. Open the `/etc/iscsi/iscsid.conf` file and find the following lines.

```
node.session.iscsi.FirstBurstLength = [replacement_first_burst_length_value]
node.session.iscsi.MaxBurstLength = [replacement_max_burst_length_value]
node.conn[0].iscsi.MaxRecvDataSegmentLength
= [replacement_segment_length_value]
```

- b. We recommend the following values to achieve better performance. Your backup software might be optimized to use different values, so see your backup software documentation for best results.

Set the `[replacement_first_burst_length_value]` value to **262144** or the Linux OS default, whichever is higher.

Set the `[replacement_max_burst_length_value]` value to **1048576** or the Linux OS default, whichever is higher.

Set the `[replacement_segment_length_value]` value to **262144** or the Linux OS default, whichever is higher.

Note

Different backup software can be optimized to work best using different iSCSI settings. To verify which values for these parameters will provide the best performance, see the documentation for your backup software.

- Restart your system to ensure that the new configuration values take effect.

Before restarting, make sure that the results of all write operations to your tapes are flushed. To do this, unmount tapes before restarting.

Customizing Your Linux Disk Timeout Settings for Volume Gateways

If you are using a Volume Gateway, you can customize the following Linux disk timeout settings in addition to the iSCSI settings described in the preceding section.

To customize your Linux disk timeout settings

- Increase the disk timeout value in the rules file.
 - If you are using the RHEL 5 initiator, 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'"
```

This rules file does not exist in RHEL 6 or 7 initiators, so you must create it using the following rule.

```
ACTION=="add", SUBSYSTEMS=="scsi" , ATTRS{model}=="Storage Gateway", \  
RUN+="/bin/sh -c 'echo [timeout] > /sys$$DEVPATH/timeout'"
```

To modify the timeout value in RHEL 6, use the following command, and then add the lines of code shown preceding.

```
sudo vim /etc/udev/rules.d/50-udev.rules
```

To modify the timeout value in RHEL 7, use the following command, and then add the lines of code shown preceding.

```
sudo su -c "echo 600 > /sys/block/[device name]/device/timeout"
```

- b. Set the *[timeout]* value to **600**.

This value represents a timeout of 600 seconds.

2. Restart your system to ensure that the new configuration values take effect.

Before restarting, make sure that the results of all write operations to your volumes are flushed. To do this, unmount storage volumes before restarting.

3. You can test the configuration by using the following command.

```
udevadm test [PATH_TO_ISCSI_DEVICE]
```

This command shows the udev rules that are applied to the iSCSI device.

Configuring CHAP Authentication for Your iSCSI Targets

Storage Gateway supports authentication between your gateway and iSCSI initiators by using Challenge-Handshake Authentication Protocol (CHAP). CHAP provides protection against playback attacks by periodically verifying the identity of an iSCSI initiator as authenticated to access a volume and VTL device target.

Note

CHAP configuration is optional but highly recommended.

To set up CHAP, you must configure it both on the Storage Gateway console and in the iSCSI initiator software that you use to connect to the target. Storage Gateway uses mutual CHAP, which is when the initiator authenticates the target and the target authenticates the initiator.

To set up mutual CHAP for your targets

1. Configure CHAP on the Storage Gateway console, as discussed in [To configure CHAP for a volume target on the Storage Gateway console](#).

2. In your client initiator software, complete the CHAP configuration:
 - To configure mutual CHAP on a Windows client, see [To configure mutual CHAP on a Windows client](#).
 - To configure mutual CHAP on a Red Hat Linux client, see [To configure mutual CHAP on a Red Hat Linux client](#).

To configure CHAP for a volume target on the Storage Gateway console

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

1. On the Storage Gateway console, choose **Volumes** in the navigation pane.
2. For **Actions**, choose **Configure CHAP Authentication**.
3. Provide the requested information in the **Configure CHAP Authentication** dialog box.
 - a. For **Initiator Name**, enter the name of your iSCSI initiator. This name is an Amazon iSCSI qualified name (IQN) that is prepended by `iqn.1997-05.com.amazon:` followed by the target name. The following is an example.

```
iqn.1997-05.com.amazon:your-volume-name
```

You can find the initiator name by using your iSCSI initiator software. For example, for Windows clients, the name is the value on the **Configuration** tab of the iSCSI initiator. For more information, see [To configure mutual CHAP on a Windows client](#).

Note


To change an initiator name, you must first deactivate CHAP, change the initiator name in your iSCSI initiator software, and then activate CHAP with the new name.

- b. For **Secret used to Authenticate Initiator**, enter the secret requested.

This secret must be a minimum of 12 characters and a maximum of 16 characters long. This value is the secret key that the initiator (that is, the Windows client) must know to participate in CHAP with the target.

- c. For **Secret used to Authenticate Target (Mutual CHAP)**, enter the secret requested.

This secret must be a minimum of 12 characters and a maximum of 16 characters long. This value 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.

- d. Choose **Save**.
4. Choose the **Details** tab and confirm that **iSCSI CHAP authentication** is set to **true**.

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 volume on the console.

1. If the iSCSI initiator is not already started, on the **Start** menu of your Windows client computer, choose **Run**, enter **iscsicpl.exe**, and then choose **OK** to run the program.
2. Configure mutual CHAP configuration for the initiator (that is, the Windows client):
 - a. Choose the **Configuration** tab.

 **Note**

The **Initiator Name** value is unique to your initiator and company. The name shown preceding is the value that you used in the **Configure CHAP Authentication** dialog box of the Storage Gateway console.

The name shown in the example image is for demonstration purposes only.

- b. Choose **CHAP**.
- c. In the **iSCSI Initiator Mutual Chap Secret** dialog box, enter the mutual CHAP secret value.

In this dialog box, you enter the secret that the initiator (the Windows client) uses to authenticate the target (the storage volume). This secret allows the target to read and write to the initiator. This secret is the same as the secret entered into the **Secret used to**

Authenticate Target (Mutual CHAP) box in the **Configure CHAP Authentication** dialog box. For more information, see [Configuring CHAP Authentication for Your iSCSI Targets](#).

- d. If the key that you entered is fewer than 12 characters or more than 16 characters long, an **Initiator CHAP secret** error dialog box appears.

Choose **OK**, and then enter the key again.

3. Configure the target with the initiator's secret to complete the mutual CHAP configuration.
 - a. Choose the **Targets** tab.
 - b. If the target that you want to configure for CHAP is currently connected, disconnect the target by selecting it and choosing **Disconnect**.
 - c. Select the target that you want to configure for CHAP, and then choose **Connect**.
 - d. In the **Connect to Target** dialog box, choose **Advanced**.
 - e. In the **Advanced Settings** dialog box, configure CHAP.
 - i. Select **Activate CHAP log on**.
 - ii. Enter the secret that is required to authenticate the initiator. This secret is the same as the secret entered into the **Secret used to Authenticate Initiator** box in the **Configure CHAP Authentication** dialog box. For more information, see [Configuring CHAP Authentication for Your iSCSI Targets](#).
 - iii. Select **Perform mutual authentication**.
 - iv. To apply the changes, choose **OK**.
 - f. In the **Connect to Target** dialog box, choose **OK**.
4. If you provided the correct secret key, the target shows a status of **Connected**.

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 volume on the Storage Gateway console.

1. Ensure 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 to a Red Hat Enterprise Linux Client](#).
2. Disconnect and remove any existing configuration for the target for which you are about to configure CHAP.

- a. To find the target name and ensure it is a defined configuration, list the saved configurations using the following command.

```
sudo /sbin/iscsiadm --mode node
```

- b. Disconnect from the target.

The following command disconnects from the target named **myvolume** that is defined in the Amazon iSCSI qualified name (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
```

3. Edit the iSCSI configuration file to activate CHAP.

- a. Get the name of the initiator (that is, 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 looks 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 values for *username*, *password*, *username_in*, and *password_in*.

```
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
```

For guidance on what values to specify, see the following table.

Configuration Setting	Value
<i>username</i>	The initiator name that you found in a previous step in this procedure. The value starts with <i>iqn</i> . For example, iqn.1994-05.com.redhat:8e89b27b5b8 is a valid <i>username</i> value.
<i>password</i>	The secret key used to authenticate the initiator (the client you are using) when it communicates with the volume.
<i>username_in</i>	The IQN of the target volume. The value starts with <i>iqn</i> and ends with the target name. For example, iqn.1997-05.com.amazon:myvolume is a valid <i>username_in</i> value.
<i>password_in</i>	The secret key used to authenticate the target (the volume) when it communicates to the initiator.

- d. Save the changes in the configuration file, and then close the file.
4. Discover and log in to the target. To do so, follow the steps in [Connecting to a Red Hat Enterprise Linux Client](#).

Using AWS Direct Connect with Storage Gateway

AWS Direct Connect links your internal network to the Amazon Web Services Cloud. By using AWS Direct Connect with Storage Gateway, you can create a connection for high-throughput workload needs, providing a dedicated network connection between your on-premises gateway and AWS.

Storage Gateway uses public endpoints. With an AWS Direct Connect connection in place, you can create a public virtual interface to allow traffic to be routed to the Storage Gateway endpoints. The public virtual interface bypasses internet service providers in your network path. The Storage Gateway service public endpoint can be in the same AWS Region as the AWS Direct Connect location, or it can be in a different AWS Region.

The following illustration shows an example of how AWS Direct Connect works with Storage Gateway.

network architecture showing Storage Gateway connected to the cloud using AWS direct connect.

The following procedure assumes that you have created a functioning gateway.

To use AWS Direct Connect with Storage Gateway

1. Create and establish an AWS Direct Connect connection between your on-premises data center and your Storage Gateway endpoint. For more information about how to create a connection, see [Getting Started with AWS Direct Connect](#) in the *AWS Direct Connect User Guide*.
2. Connect your on-premises Storage Gateway appliance to the AWS Direct Connect router.
3. Create a public virtual interface, and configure your on-premises router accordingly. Even with Direct Connect, VPC endpoints must be created with the HAProxy. For more information, see [Creating a Virtual Interface](#) in the *AWS Direct Connect User Guide*.

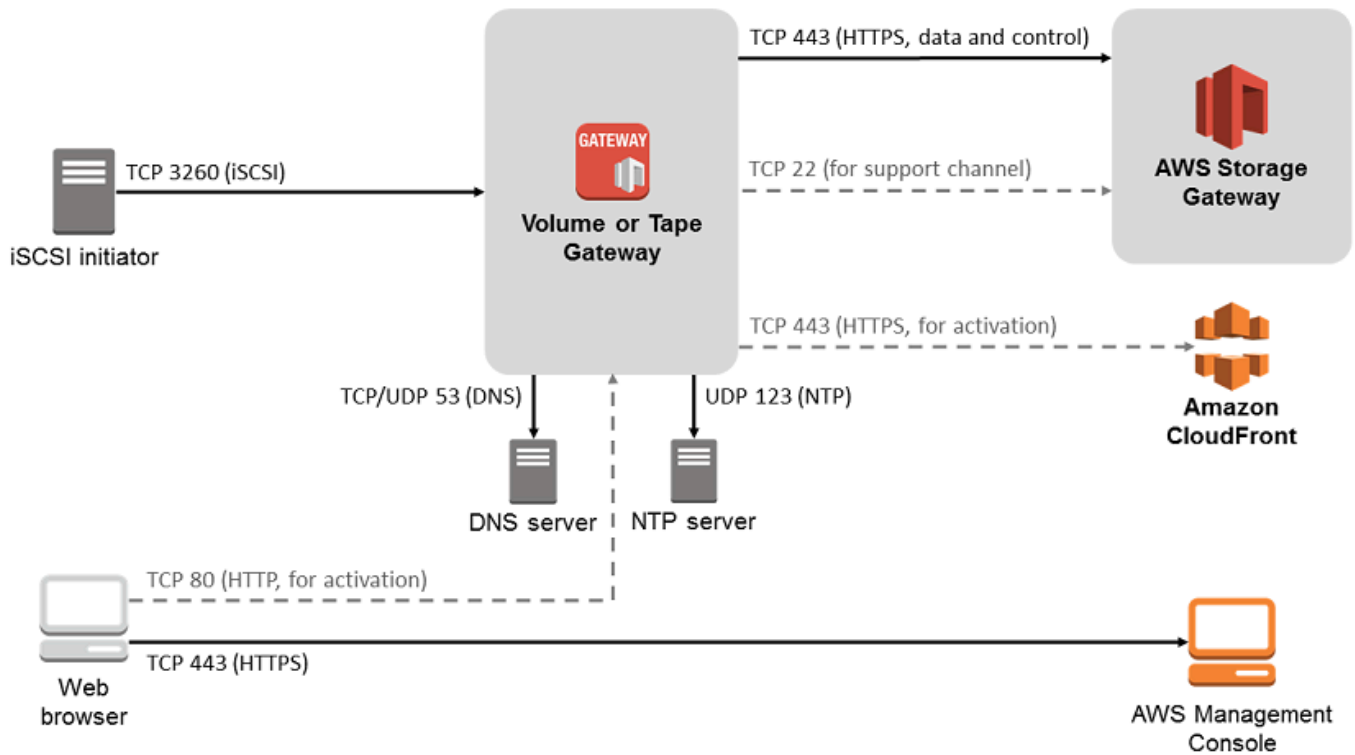
For details about AWS Direct Connect, see [What is AWS Direct Connect?](#) in the *AWS Direct Connect User Guide*.

Port requirements for Volume Gateway

Storage Gateway requires the following ports for its operation. Some ports are common to and required by all gateway types. Other ports are required by specific gateway types. In this section, you can find an illustration and a list of the required ports for Volume Gateway.

Volume Gateway

The following illustration shows all of the ports you need to open for Volume Gateway gateway operation.



The following ports are common to and required by all gateway types.

From	To	Protocol	Port	How Used
Storage Gateway VM	AWS	Transmission Control Protocol (TCP)	443 (HTTPS)	For communication from an Storage Gateway outbound VM to an AWS service endpoint. For information about service endpoints, see Allowing AWS Storage Gateway

From	To	Protocol	Port	How Used	
				access through firewalls and routers.	

From	To	Protocol	Port	How Used	
Your web browser	Storage Gateway VM	TCP	80 (HTTP)	<p>By local systems to obtain the Storage Gateway activation key. Port 80 is used only during activation of a Storage Gateway appliance.</p> <p>A Storage Gateway VM doesn't require port 80 to be publicly accessible. The required level of access to port 80 depends on your network configuration. If you activate your gateway from the Storage Gateway Management</p>	

From	To	Protocol	Port	How Used	
				Console, the host from which you connect to the console must have access to your gateway's port 80.	
Storage Gateway VM	Domain Name Service (DNS) server	User Datagram Protocol (UDP)/UDP	53 (DNS)	For communication between a Storage Gateway VM and the DNS server.	

From	To	Protocol	Port	How Used	
Storage Gateway VM	AWS	TCP	22 (Support channel)	Allows AWS Support to access your gateway to help you with troubleshooting gateway issues. You don't need this port open for the normal operation of your gateway, but it is required for troubleshooting.	

From	To	Protocol	Port	How Used
Storage Gateway VM	Network Time Protocol (NTP) server	UDP	123 (NTP)	Used by local systems to synchronize VM time to the host time. A Storage Gateway VM is configured to use the following NTP servers: <ul style="list-style-type: none"> 0.amazon.pool.ntp.org 1.amazon.pool.ntp.org 2.amazon.pool.ntp.org 3.amazon.pool.ntp.org
Storage Gateway Hardware Appliance	Hypertext Transfer Protocol (HTTP) proxy	TCP	8080 (HTTP)	Required briefly for activation.

In addition to the common ports, Volume Gateway also requires the following ports.

From	To	Protocol	Port	How Used
iSCSI initiators	Storage Gateway VM	TCP	3260 (iSCSI)	By local systems to connect to iSCSI targets exposed by a gateway.

Getting the IP address for your gateway appliance

After you choose a host and deploy your gateway VM, you connect and activate your gateway. To do this, you need the IP address of your gateway VM. You get the IP address from your gateway's local console. You log in to the local console and get the IP address from the top of the console page.

For gateways deployed on-premises, you can also get the IP address from your hypervisor. For Amazon EC2 gateways, you can also get the IP address of your Amazon EC2 instance from the Amazon EC2 Management Console. To find how to get your gateway's IP address, see one of the following:

- VMware host: [Accessing the Gateway Local Console with VMware ESXi](#)
- HyperV host: [Access the Gateway Local Console with Microsoft Hyper-V](#)
- Linux Kernel-based Virtual Machine (KVM) host: [Accessing the Gateway Local Console with Linux KVM](#)
- EC2 host: [Getting an IP Address from an Amazon EC2 Host](#)

When you locate the IP address, take note of it. Then return to the Storage Gateway console and type the IP address into the console.

Getting an IP Address from an Amazon EC2 Host

To get the IP address of the Amazon EC2 instance your gateway is deployed on, log in to the EC2 instance's local console. Then get the IP address from the top of the console page. For instructions, see [Logging In to Your Amazon EC2 Gateway Local Console](#).

You can also get the IP address from the Amazon EC2 Management Console. We recommend using the public IP address for activation. To get the public IP address, use procedure 1. If you choose to use the elastic IP address instead, see procedure 2.

Procedure 1: To connect to your gateway using the public IP address

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Instances**, and then select the EC2 instance that your gateway is deployed on.
3. Choose the **Description** tab at the bottom, and then note the public IP. You use this IP address to connect to the gateway. Return to the Storage Gateway console and type in the IP address.

If you want to use the elastic IP address for activation, use the procedure following.

Procedure 2: To connect to your gateway using the elastic IP address

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Instances**, and then select the EC2 instance that your gateway is deployed on.
3. Choose the **Description** tab at the bottom, and then note the **Elastic IP** value. You use this elastic IP address to connect to the gateway. Return to the Storage Gateway console and type in the elastic IP address.
4. After your gateway is activated, choose the gateway that you just activated, and then choose the **VTL devices** tab in the bottom panel.
5. Get the names of all your VTL devices.
6. For each target, run the following command to configure the target.

```
iscsiadm -m node -o new -T [$TARGET_NAME] -p [$Elastic_IP]:3260
```

7. For each target, run the following command to log in.

```
iscsiadm -m node -p [$ELASTIC_IP]:3260 --login
```

Your gateway is now connected using the elastic IP address of the EC2 instance.

Understanding Storage Gateway Resources and Resource IDs

In Storage Gateway, the primary resource is a *gateway* but other resource types include: *volume*, *virtual tape*, *iSCSI target*, and *vtl device*. These are referred to as *subresources* and they don't exist unless they are associated with a gateway.

These resources and subresources have unique Amazon Resource Names (ARNs) associated with them as shown in the following table.

Resource Type	ARN Format
Gateway ARN	arn:aws:storagegateway: <i>region:account-id</i> :gateway/ <i>gateway-id</i>
Volume ARN	arn:aws:storagegateway: <i>region:account-id</i> :gateway/ <i>gateway-id</i> /volume/ <i>volume-id</i>
Target ARN (iSCSI target)	arn:aws:storagegateway: <i>region:account-id</i> :gateway/ <i>gateway-id</i> /target/ <i>iSCSItarget</i>

Storage Gateway also supports the use of EC2 instances and EBS volumes and snapshots. These resources are Amazon EC2 resources that are used in Storage Gateway.

Working with Resource IDs

When you create a resource, Storage Gateway assigns the resource a unique resource ID. This resource ID is part of the resource ARN. A resource ID takes the form of a resource identifier, followed by a hyphen, and a unique combination of eight letters and numbers. For example, a gateway ID is of the form `sgw-12A3456B` where `sgw` is the resource identifier for gateways. A volume ID takes the form `vol-3344CCDD` where `vol` is the resource identifier for volumes.

For virtual tapes, you can prepend a up to a four character prefix to the barcode ID to help you organize your tapes.

Storage Gateway resource IDs are in uppercase. However, when you use these resource IDs with the Amazon EC2 API, Amazon EC2 expects resource IDs in lowercase. You must change your resource ID

to lowercase to use it with the EC2 API. For example, in Storage Gateway the ID for a volume might be `vol-1122AABB`. When you use this ID with the EC2 API, you must change it to `vol-1122aabb`. Otherwise, the EC2 API might not behave as expected.

Tagging Storage Gateway Resources

In Storage Gateway, you can use tags to manage your resources. Tags let you add metadata to your resources and categorize your resources to make them easier to manage. Each tag consists of a key-value pair, which you define. You can add tags to gateways, volumes, and virtual tapes. You can search and filter these resources based on the tags you add.

As an example, you can use tags to identify Storage Gateway resources used by each department in your organization. You might tag gateways and volumes used by your accounting department like this: (key=department and value=accounting). You can then filter with this tag to identify all gateways and volumes used by your accounting department and use the information to determine cost. For more information, see [Using Cost Allocation Tags](#) and [Working with Tag Editor](#).

If you archive a virtual tape that is tagged, the tape maintains its tags in the archive. Similarly, if you retrieve a tape from the archive to another gateway, the tags are maintained in the new gateway.

Tags don't have any semantic meaning but rather are interpreted as strings of characters.

The following restrictions apply to tags:

- Tag keys and values are case-sensitive.
- The maximum number of tags for each resource is 50.
- Tag keys cannot begin with `aws :`. This prefix is reserved for AWS use.
- Valid characters for the key property are UTF-8 letters and numbers, space, and special characters `+ - = . _ : /` and `@`.

Working with Tags

You can work with tags by using the Storage Gateway console, the Storage Gateway API, or the [Storage Gateway Command Line Interface \(CLI\)](#). The following procedures show you how to add, edit, and delete a tag on the console.

To add a tag

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose the resource you want to tag.

For example, to tag a gateway, choose **Gateways**, and then choose the gateway you want to tag from the list of gateways.
3. Choose **Tags**, and then choose **Add/edit tags**.
4. In the **Add/edit tags** dialog box, choose **Create tag**.
5. Type a key for **Key** and a value for **Value**. For example, you can type **Department** for the key and **Accounting** for the value.

Note

You can leave the **Value** box blank.

6. Choose **Create Tag** to add more tags. You can add multiple tags to a resource.
7. When you're done adding tags, choose **Save**.

To edit a tag

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose the resource whose tag you want to edit.
3. Choose **Tags** to open the **Add/edit tags** dialog box.
4. Choose the pencil icon next to the tag you want edit, and then edit the tag.
5. When you're done editing the tag, choose **Save**.

To delete a tag

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose the resource whose tag you want to delete.
3. Choose **Tags**, and then choose **Add/edit tags** to open the **Add/edit tags** dialog box.

4. Choose the **X** icon next to the tag you want to delete, and then choose **Save**.

Working with open-source components for Storage Gateway

This section describes third party tools and licenses that we depend on to deliver Storage Gateway functionality.

The source code for certain open-source software components that are included with the AWS Storage Gateway software is available for download at the following locations:

- For gateways deployed on VMware ESXi, download [sources.tar](#)
- For gateways deployed on Microsoft Hyper-V, download [sources_hyperv.tar](#)
- For gateways deployed on Linux Kernel-based Virtual Machine (KVM), download [sources_KVM.tar](#)

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>). For the relevant licenses for all dependent third party tools, see [Third Party Licenses](#).

AWS Storage Gateway quotas

In this topic, you can find information about volume and tape quotas, configuration, and performance limits for Storage Gateway.

Topics

- [Quotas for volumes](#)
- [Recommended local disk sizes for your gateway](#)

Quotas for volumes

The following table lists quotas for volumes.

Description	Cached volumes	Stored volumes
Maximum size of a volume	32 TiB	16 TiB

Description	Cached volumes	Stored volumes
<p>Note</p> <p>If you create a snapshot from a cached volume that is more than 16 TiB in size, you can restore it to a Storage Gateway volume but not to an Amazon Elastic Block Store (Amazon EBS) volume.</p>		
Maximum number of volumes per gateway	32	32
Total size of all volumes for a gateway	1,024 TiB	512 TiB

Recommended local disk sizes for your gateway

The following table recommends sizes for local disk storage for your deployed gateway.

Gateway Type	Cache (Minimum)	Cache (Maximum)	Upload Buffer (Minimum)	Upload Buffer (Maximum)	Other Required Local Disks
Cached volume gateway	150 GiB	64 TiB	150 GiB	2 TiB	—
Stored volume gateway	—	—	150 GiB	2 TiB	1 or more for stored volume or volumes

Note

You can configure one or more local drives for your cache and upload buffer, up to the maximum capacity.

When adding cache or upload buffer to an existing gateway, it's important to create new disks in your host (hypervisor or Amazon EC2 instance). Don't change the size of existing disks if the disks have been previously allocated as either a cache or upload buffer.

API Reference for Storage Gateway

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 Storage Gateway, see [AWS Storage Gateway Endpoints and Quotas](#) in the *AWS General Reference*.

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, see [Sample Code Libraries](#).

Topics

- [Storage Gateway Required Request Headers](#)
- [Signing Requests](#)
- [Error Responses](#)
- [Actions](#)

Storage Gateway Required Request Headers

This section describes the required headers that you must send with every POST request to 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](#) operation.

```
POST / HTTP/1.1
```

```
Host: storagegateway.us-east-2.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-2/
storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.ActivateGateway
```

The following are the headers that must include with your POST requests to Storage Gateway. Headers shown below that begin with "x-amz" are AWS-specific headers. All other headers listed are common header used in HTTP transactions.

Header	Description
Authorization	<p>The authorization header contains several of pieces of information about the request that allows 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/aw s4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-targ et, 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>
Content-Type	<p>Use <code>application/x-amz-json-1.1</code> as the content type for all requests to Storage Gateway.</p> <pre>Content-Type: application/x-amz-json-1.1</pre>

Header	Description
Host	<p>Use the host header to specify the Storage Gateway endpoint where you send your request. For example, <code>storagegateway.us-east-2.amazonaws.com</code> is the endpoint for the US East (Ohio) region. For more information about the endpoints available for Storage Gateway, see AWS Storage Gateway Endpoints and Quotas in the <i>AWS General Reference</i>.</p> <pre>Host: storagegateway. <i>region</i>.amazonaws.com</pre>
x-amz-date	<p>You must provide the time stamp in either the HTTP Date header or the AWS x-amz-date header. (Some HTTP client libraries don't let you set the Date header.) When an x-amz-date header is present, the Storage Gateway ignores any Date header during the request authentication. The x-amz-date format must be ISO8601 Basic in the YYYYMMDD'T'HHMMSS'Z' format. If both the Date and x-amz-date header are used, the format of the Date header does not have to be ISO8601.</p> <pre>x-amz-date: <i>YYYYMMDD'T'HHMMSS'Z'</i></pre>
x-amz-target	<p>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.</p> <pre>x-amz-target: StorageGateway_ <i>APIversion</i> .<i>operationName</i></pre> <p>The <i>operationName</i> value (e.g. "ActivateGateway") can be found from the API list, API Reference for Storage Gateway.</p>

Signing Requests

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, 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, Storage Gateway processes the request. Otherwise, the request is rejected.

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 Storage Gateway uses the same canonical form when it recalculates a signature to compare with the one you sent.

- [Task 2: Create a String to Sign](#)

Create a string that you will use as one of the input values to your cryptographic hash function. The string, called the *string to sign*, is a concatenation of the name of the hash algorithm, the request date, a *credential scope* string, and the canonicalized request from the previous task. The *credential scope* string itself is a concatenation of date, region, and service information.

- [Task 3: Create a Signature](#)

Create a signature for your request by using a cryptographic hash function that accepts two input strings: your *string to sign* and a *derived key*. The *derived key* is calculated by starting with your secret access key and using the *credential scope* string to create a series of Hash-based Message Authentication Codes (HMACs).

Example Signature Calculation

The following example walks you through the details of creating a signature for [ListGateways](#). The example could be used as a reference to check your signature calculation method. Other

reference calculations are included in the [Signature Version 4 Test Suite](#) of the Amazon Web Services Glossary.

The example assumes the following:

- The time stamp of the request is "Mon, 10 Sep 2012 00:00:00" GMT.
- The endpoint is the US East (Ohio) region.

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

```
POST / HTTP/1.1
Host: storagegateway.us-east-2.amazonaws.com
x-amz-Date: 20120910T000000Z
Authorization: SignatureToBeCalculated
Content-type: application/x-amz-json-1.1
x-amz-target: StorageGateway_20120630.ListGateways
{ }
```

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

```
POST
/

content-type:application/x-amz-json-1.1
host:storagegateway.us-east-2.amazonaws.com
x-amz-date:20120910T000000Z
x-amz-target:StorageGateway_20120630.ListGateways

content-type;host;x-amz-date;x-amz-target
44136fa355b3678a1146ad16f7e8649e94fb4fc21fe77e8310c060f61caaff8a
```

The last line of the canonical request is the hash of the request body. Also, note the empty third line in the canonical request. This is because there are no query parameters for this API (or any Storage Gateway APIs).

The *string to sign* for [Task 2: Create a String to Sign](#) is:

```
AWS4-HMAC-SHA256
```

```
20120910T000000Z
20120910/us-east-2/storagegateway/aws4_request
92c0effa6f9224ac752ca179a04cecbede3038b0959666a8160ab452c9e51b3e
```

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](#), the *derived key* can be represented as:

```
derived key = HMAC(HMAC(HMAC(HMAC("AWS4" + YourSecretAccessKey, "20120910"), "us-
east-2"), "storagegateway"), "aws4_request")
```

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

```
6d4c40b8f2257534dbdca9f326f147a0a7a419b63aff349d9d9c737c9a0f4c81
```

The final step is to construct the Authorization header. For the demonstration access key AKIAIOSFODNN7EXAMPLE, the header (with line breaks added for readability) is:

```
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120910/us-east-2/
storagegateway/aws4_request,
SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=6d4c40b8f2257534dbdca9f326f147a0a7a419b63aff349d9d9c737c9a0f4c81
```

Error Responses

Topics

- [Exceptions](#)
- [Operation Error Codes](#)
- [Error Responses](#)

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](#) API.

Depending on the type of error, 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](#).

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 [Operation Error Codes](#) 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 Operation Error Codes .	500 Internal Server Error
<code>InvalidAction</code>	The requested action or operation is not valid.	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 .	400 Bad Request
<code>InvalidSignatureException</code>	The request signature we calculate does not match the signature you	400 Bad Request

Exception	Message	HTTP Status Code
	provided. Check your AWS Access Key and signing method.	
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
TooManyRequests	Too many requests.	429 Too Many Requests
UnknownOperationException	An unknown operation was specified. Valid operations are listed in Operations in Storage Gateway .	400 Bad Request
UnrecognizedClientException	The security token included in the request is not valid.	400 Bad Request

Exception	Message	HTTP Status Code
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`—described in [Exceptions](#).

Operation Error Code	Message	Operations That Return this Error Code
ActivationKeyExpired	The specified activation key has expired.	ActivateGateway
ActivationKeyInvalid	The specified activation key is not valid.	ActivateGateway
ActivationKeyNotFound	The specified activation key was not found.	ActivateGateway
BandwidthThrottleScheduleNotFound	The specified bandwidth throttle was not found.	DeleteBandwidthRateLimit
CannotExportSnapshot	The specified snapshot cannot be exported.	CreateCachediSCSIVolume CreateStorediSCSIVolume
InitiatorNotFound	The specified initiator was not found.	DeleteChapCredentials
DiskAlreadyAllocated	The specified disk is already allocated.	AddCache AddUploadBuffer

Operation Error Code	Message	Operations That Return this Error Code
		AddWorkingStorage CreateStorediSCSIVolume
DiskDoesNotExist	The specified disk does not exist.	AddCache AddUploadBuffer AddWorkingStorage CreateStorediSCSIVolume
DiskSizeNotGigAligned	The specified disk is not gigabyte-aligned.	CreateStorediSCSIVolume
DiskSizeGreaterThanVolumeMaxSize	The specified disk size is greater than the maximum volume size.	CreateStorediSCSIVolume
DiskSizeLessThanVolumeSize	The specified disk size is less than the volume size.	CreateStorediSCSIVolume
DuplicateCertificateInfo	The specified certificate information is a duplicate.	ActivateGateway

Operation Error Code	Message	Operations That Return this Error Code
GatewayInternalError	A gateway internal error occurred.	AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateStorediSCSIVolume CreateSnapshotFromVolumeRecoveryPoint DeleteBandwidthRateLimit DeleteChapCredentials DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes DescribeWorkingStorage ListLocalDisks

Operation Error Code	Message	Operations That Return this Error Code
		ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewaySoftwareNow UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
GatewayNotConnected	The specified gateway is not connected.	AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateStorediSCSIVolume CreateSnapshotFromVolumeRecoveryPoint DeleteBandwidthRateLimit DeleteChapCredentials DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes DescribeWorkingStorage ListLocalDisks

Operation Error Code	Message	Operations That Return this Error Code
		ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewaySoftwareNow UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
GatewayNotFound	The specified gateway was not found.	AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteGateway DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes DescribeWorkingStorage

Operation Error Code	Message	Operations That Return this Error Code
		ListLocalDisks
		ListVolumes
		ListVolumeRecoveryPoints
		ShutdownGateway
		StartGateway
		UpdateBandwidthRateLimit
		UpdateChapCredentials
		UpdateMaintenanceStartTime
		UpdateGatewaySoftwareNow
		UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
GatewayProxyNetworkConnectionBusy	The specified gateway proxy network connection is busy.	AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes DescribeWorkingStorage ListLocalDisks

Operation Error Code	Message	Operations That Return this Error Code
		ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewaySoftwareNow UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
InternalError	An internal error occurred.	ActivateGateway AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteGateway DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes

Operation Error Code	Message	Operations That Return this Error Code
		DescribeWorkingStorage ListLocalDisks ListGateways ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewayInformation UpdateGatewaySoftwareNow UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
InvalidParameters	The specified request contains incorrect parameters.	ActivateGateway AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteGateway DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes

Operation Error Code	Message	Operations That Return this Error Code
		DescribeWorkingStorage ListLocalDisks ListGateways ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewayInformation UpdateGatewaySoftwareNow UpdateSnapshotSchedule
LocalStorageLimitExceeded	The local storage limit was exceeded.	AddCache AddUploadBuffer AddWorkingStorage
LunInvalid	The specified LUN is incorrect.	CreateStorediSCSIVolume

Operation Error Code	Message	Operations That Return this Error Code
MaximumVolumeCount Exceeded	The maximum volume count was exceeded.	CreateCachediSCSIVolume CreateStorediSCSIVolume DescribeCachediSCSIVolumes DescribeStorediSCSIVolumes
NetworkConfigurationChanged	The gateway network configuration has changed.	CreateCachediSCSIVolume CreateStorediSCSIVolume

Operation Error Code	Message	Operations That Return this Error Code
NotSupported	The specified operation is not supported.	ActivateGateway AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteGateway DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes

Operation Error Code	Message	Operations That Return this Error Code
		DescribeWorkingStorage ListLocalDisks ListGateways ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewayInformation UpdateGatewaySoftwareNow UpdateSnapshotSchedule
OutdatedGateway	The specified gateway is out of date.	ActivateGateway
SnapshotInProgressException	The specified snapshot is in progress.	DeleteVolume
SnapshotIdInvalid	The specified snapshot is not valid.	CreateCachediSCSIVolume CreateStorediSCSIVolume
StagingAreaFull	The staging area is full.	CreateCachediSCSIVolume CreateStorediSCSIVolume

Operation Error Code	Message	Operations That Return this Error Code
TargetAlreadyExists	The specified target already exists.	CreateCachediSCSIVolume CreateStorediSCSIVolume
TargetInvalid	The specified target is not valid.	CreateCachediSCSIVolume CreateStorediSCSIVolume DeleteChapCredentials DescribeChapCredentials UpdateChapCredentials
TargetNotFound	The specified target was not found.	CreateCachediSCSIVolume CreateStorediSCSIVolume DeleteChapCredentials DescribeChapCredentials DeleteVolume UpdateChapCredentials

Operation Error Code	Message	Operations That Return this Error Code
UnsupportedOperationForGatewayType	The specified operation is not valid for the type of the gateway.	AddCache AddWorkingStorage CreateCachediSCSIVolume CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteSnapshotSchedule DescribeCache DescribeCachediSCSIVolumes DescribeStorediSCSIVolumes DescribeUploadBuffer DescribeWorkingStorage ListVolumeRecoveryPoints
VolumeAlreadyExists	The specified volume already exists.	CreateCachediSCSIVolume CreateStorediSCSIVolume
VolumeIdInvalid	The specified volume is not valid.	DeleteVolume
VolumeInUse	The specified volume is already in use.	DeleteVolume

Operation Error Code	Message	Operations That Return this Error Code
VolumeNotFound	The specified volume was not found.	CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint DeleteVolume DescribeCachediSCSIVolumes DescribeSnapshotSchedule DescribeStorediSCSIVolumes UpdateSnapshotSchedule
VolumeNotReady	The specified volume is not ready.	CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint

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](#).

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 .

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.

Type: String

Error Response Examples

The following JSON body is returned if you use the DescribeStorediSCSIVolumes 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 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 Storage Gateway

For a list of Storage Gateway operations, see [Actions](#) in the *AWS Storage Gateway API Reference*.

Document history for the Volume Gateway User Guide

- **API version:** 2013-06-30
- **Latest documentation update:** November 24, 2020

The following table describes important changes in each release of the *AWS Storage Gateway User Guide* after April 2018. For notification about updates to this documentation, you can subscribe to an RSS feed.

Change	Description	Date
Notice of availability change for FSx File Gateway	AWS Storage Gateway's FSx File Gateway will no longer be available to new customers after 10/28/24. To use the service, you must sign up prior to that date. Existing customers of FSx File Gateway can continue to use the service normally. For capabilities similar to FSx File Gateway, visit this blog post .	September 26, 2024
Added option to turn maintenance updates on or off	Storage Gateway receives regular maintenance updates that can include operating system and software upgrades, fixes to address stability, performance, and security, and access to new features. You can now configure a setting to turn these updates on or off for each individual gateway in your deployment. For more information, see Managing	June 6, 2024

[gateway updates using the AWS Storage Gateway console.](#)

[Deprecated support for Tape Gateway on Snowball Edge](#)

It is no longer possible to host Tape Gateway on Snowball Edge devices.

March 14, 2024

[Updated instructions for testing your gateway setup using 3rd party applications](#)

The instructions for testing your gateway setup using 3rd party applications now describe the expected behavior if your gateway restarts during an ongoing backup job. For more information, see .

October 24, 2023

[Updated recommended CloudWatch alarms](#)

The CloudWatch HealthNotifications alarm now applies to and is recommended for all gateway types and host platforms. Recommended configuration settings have also been updated for HealthNotifications and AvailabilityNotifications . For more information see [Understanding CloudWatch alarms.](#)

October 2, 2023

[Separated Tape and Volume Gateway User Guides](#)

The Storage Gateway User Guide, which previously contained information about both the tape and Volume Gateway types, has been split into the Tape Gateway User Guide and the Volume Gateway User Guide, each containing information on only one type of gateway. For more information, see [Tape Gateway User Guide](#) and [Volume Gateway User Guide](#).

March 23, 2022

[Updated gateway creation procedures](#)

Procedures for creating all gateway types using the Storage Gateway console have been updated. For more information, see [Creating Your Gateway](#).

January 18, 2022

[New Tapes interface](#)

The **Tape overview** page in the AWS Storage Gateway console has been updated with new search and filtering features. All relevant procedures in this guide have been updated to describe the new functionality. For more information, see [Managing Your Tape Gateway](#).

September 23, 2021

[Support for Quest NetVault Backup 13 for Tape Gateway](#)

Tape Gateways now support Quest NetVault Backup 13 running on Microsoft Windows Server 2012 R2 or Microsoft Windows Server 2016. For more information, see, see [Testing Your Setup by Using Quest NetVault Backup](#).

August 22, 2021

[S3 File Gateway topics removed from Tape and Volume Gateway guides](#)

To help make the user guides for Tape Gateway and Volume Gateway easier to follow for customers setting up their respective gateway types, some unnecessary topics have been removed.

July 21, 2021

[Support for IBM Spectrum Protect 8.1.10 on Windows and Linux for Tape Gateway](#)

Tape Gateways now support IBM Spectrum Protect version 8.1.10 running on Microsoft Windows Server and Linux. For more information, see [Testing Your Setup by Using IBM Spectrum Protect](#).

November 24, 2020

[FedRAMP compliance](#)

Storage Gateway is now FedRAMP compliant. For more information, see [Compliance validation for Storage Gateway](#).

November 24, 2020

[Schedule-based bandwidth throttling](#)

Storage Gateway now supports schedule-based bandwidth throttling for tape and Volume Gateways. For more information, see [Scheduling bandwidth throttling using the Storage Gateway console](#).

November 9, 2020

[Cached volume and Tape Gateways local cache storage 4x increase](#)

Storage Gateway now supports a local cache of up to 64 TB for cached volume and Tape Gateways, improving performance for on-premises applications by providing low-latency access to larger working datasets. For more information, see [Recommended local disk sizes for your gateway](#).

November 9, 2020

[Gateway migration](#)

Storage Gateway now supports migrating cached Volume Gateways to new virtual machines. For more information, see [Moving Cached Volumes to a New Cached Volume Gateway Virtual Machine](#).

September 10, 2020

[Support for tape retention lock and write-once-read-many \(WORM\) tape protection](#)

Storage Gateway supports tape retention lock on virtual tapes and *write once read many* (WORM). Tape retention lock lets you specify the retention mode and period on archived virtual tapes, preventing them from being deleted for a fixed amount of time up to 100 years. It includes permission controls on who can delete tapes or modify retention settings. For more information, see [Using Tape Retention Lock](#). WORM-activated virtual tapes help ensure that data on active tapes in your virtual tape library cannot be overwritten or erased. For more information, see [Write Once, Read Many \(WORM\) Tape Protection](#).

August 19, 2020

[Order the hardware appliance through the console](#)

You can now order the hardware appliance through the AWS Storage Gateway console. For more information, see [Using the Storage Gateway Hardware Appliance](#).

August 12, 2020

[Support for Federal Information Processing Standard \(FIPS\) endpoints in new AWS Regions](#)

You can now activate a gateway with FIPS endpoints in the US East (Ohio), US East (N. Virginia), US West (N. California), US West (Oregon), and Canada (Central) Regions. For more information, see [AWS Storage Gateway endpoints and quotas](#) in the *AWS General Reference*.

July 31, 2020

[Gateway migration](#)

Storage Gateway now supports migrating tape and stored Volume Gateways to new virtual machines. For more information, see [Moving Your Data to a New Gateway](#).

July 31, 2020

[View Amazon CloudWatch alarms in the Storage Gateway console](#)

You can now view CloudWatch alarms in the Storage Gateway console. For more information, see [Understanding CloudWatch alarms](#).

May 29, 2020

[Support for Federal Information Processing Standard \(FIPS\) endpoints](#)

You can now activate a gateway with FIPS endpoints in the AWS GovCloud (US) Regions. To choose a FIPS endpoint for a Volume Gateway, see [Choosing a service endpoint](#). To choose a FIPS endpoint for a Tape Gateway, see [Connect your Tape Gateway to AWS](#).

May 22, 2020

[New AWS Regions](#)

Storage Gateway is now available in the Africa (Cape Town) and Europe (Milan) Regions. For more information, see [AWS Storage Gateway endpoints and quotas](#) in the *AWS General Reference*.

May 7, 2020

[Support for S3 Intelligent-Tiering storage class](#)

Storage Gateway now supports S3 Intelligent-Tiering storage class. The S3 Intelligent-Tiering storage class optimizes storage costs by automatically moving data to the most cost-effective storage access tier, without performance impact or operational overhead. For more information, see [Storage class for automatically optimizing frequently and infrequently accessed objects](#) in the *Amazon Simple Storage Service User Guide*.

April 30, 2020

[Tape Gateway write and read performance 2x increase](#)

Storage Gateway increases performance for reading from and writing to virtual tapes on Tape Gateway by 2x, allowing you to perform faster backup and recovery than before. For more information, see [Performance Guidance for Tape Gateways](#) in the *Storage Gateway User Guide*.

April 23, 2020

[Support for automatic tape creation](#)

Storage Gateway now provides the ability to automatically create new virtual tapes. Tape Gateway automatically creates new virtual tapes to maintain the minimum number of available tapes you configure and then makes these new tapes available for import by the backup application, allowing your backup jobs to run without interruption. For more information, see [Creating Tapes Automatically](#) in the *Storage Gateway User Guide*.

April 23, 2020

[New AWS Region](#)

Storage Gateway is now available in the AWS GovCloud (US-East) Region. For more information, see [AWS Storage Gateway Endpoints and Quotas](#) in the *AWS General Reference*.

March 12, 2020

[Support for Linux Kernel-based Virtual Machine \(KVM\) hypervisor](#)

Storage Gateway now provides the ability to deploy an on-premises gateway on the KVM virtualization platform. Gateways deployed on KVM have all the same functionality and features as the existing on-premises gateways. For more information, see [Supported Hypervisors and Host Requirements](#) in the *Storage Gateway User Guide*.

February 4, 2020

[Support for VMware vSphere High Availability](#)

Storage Gateway now provides support for high availability on VMware to help protect storage workloads against hardware, hypervisor, or network failures. For more information, see [Using VMware vSphere High Availability with Storage Gateway](#) in the *Storage Gateway User Guide*. This release also includes performance improvements. For more information, see [Performance](#) in the *Storage Gateway User Guide*.

November 20, 2019

[New AWS Region for Tape Gateway](#)

Tape Gateway is now available in the South America (Sao Paulo) Region. For more information, see [AWS Storage Gateway Endpoints and Quotas](#) in the *AWS General Reference*.

September 24, 2019

[Support for IBM Spectrum Protect version 7.1.9 on Linux, and for Tape Gateways an increased maximum tape size to 5 TiB](#)

Tape Gateways now support IBM Spectrum Protect (Tivoli Storage Manager) version 7.1.9 running on Linux, in addition to running on Microsoft Windows. For more information, see [Testing Your Setup by Using IBM Spectrum Protect](#) in the *Storage Gateway User Guide*. Also, for Tape Gateways, the maximum size of a virtual tape is now increased from 2.5 TiB to 5 TiB. For more information, see [Quotas for Tapes](#) in the *Storage Gateway User Guide*.

September 10, 2019

[Support for Amazon CloudWatch Logs](#)

You can now configure File Gateways with Amazon CloudWatch Log Groups to get notified about errors and the health of your gateway and its resources. For more information, see [Getting Notified About Gateway Health and Errors With Amazon CloudWatch Log Groups](#) in the *Storage Gateway User Guide*.

September 4, 2019

[New AWS Region](#)

Storage Gateway is now available in the Asia Pacific (Hong Kong) Region. For more information, see [AWS Storage Gateway Endpoints and Quotas](#) in the *AWS General Reference*.

August 14, 2019

[New AWS Region](#)

Storage Gateway is now available in the Middle East (Bahrain) Region. For more information, see [AWS Storage Gateway Endpoints and Quotas](#) in the *AWS General Reference*.

July 29, 2019

[Support for activating a gateway in a virtual private cloud \(VPC\)](#)

You can now activate a gateway in a VPC. You can create a private connection between your on-premises software appliance and cloud-based storage infrastructure. For more information, see [Activating a Gateway in a Virtual Private Cloud](#).

June 20, 2019

[Support for moving virtual tapes from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive](#)

You can now move your virtual tapes that are archived in the S3 Glacier Flexible Retrieval storage class to the S3 Glacier Deep Archive storage class for cost effective and long-term data retention. For more information, see [Moving a Tape from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive](#).

May 28, 2019

[SMB file share support for Microsoft Windows ACLs](#)

For File Gateways, you can now use Microsoft Windows access control lists (ACLs) to control access to Server Message Block (SMB) file shares. For more information, see [Using Microsoft Windows ACLs to Control Access to an SMB File Share](#).

May 8, 2019

[Integration with S3 Glacier Deep Archive](#)

Tape Gateway integrates with S3 Glacier Deep Archive. You can now archive virtual tapes in S3 Glacier Deep Archive for long-term data retention. For more information, see [Archiving Virtual Tapes](#).

March 27, 2019

[Availability of Storage Gateway Hardware Appliance in Europe](#)

The Storage Gateway Hardware Appliance is now available in Europe. For more information, see [AWS Storage Gateway Hardware Appliance Regions](#) in the *AWS General Reference*. In addition, you can now increase the useable storage on the Storage Gateway Hardware Appliance from 5 TB to 12 TB and replace the installed copper network card with a 10 Gigabit fiber optic network card. For more information, see [Setting Up Your Hardware Appliance](#).

February 25, 2019

[Integration with AWS Backup](#)

Storage Gateway integrates with AWS Backup. You can now use AWS Backup to back up on-premises business applications that use Storage Gateway volumes for cloud-backed storage. For more information, see [Backing Up Your Volumes](#).

January 16, 2019

[Support for Bacula Enterprise and IBM Spectrum Protect](#)

Tape Gateways now support Bacula Enterprise and IBM Spectrum Protect. Storage Gateway also now supports newer versions of Veritas NetBackup, Veritas Backup Exec and Quest NetVault backup. You can now use these backup applications to back up your data to Amazon S3 and archive directly to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). For more information, see [Using Your Backup Software to Test Your Gateway Setup](#).

November 13, 2018

[Support for Storage Gateway Hardware Appliance](#)

The Storage Gateway Hardware Appliance includes Storage Gateway software preinstalled on a third-party server. You can manage the appliance from the AWS Management Console. The appliance can host file, tape, and Volume Gateways. For more information, see [Using the Storage Gateway Hardware Appliance](#).

September 18, 2018

[Compatibility with Microsoft System Center 2016 Data Protection Manager \(DPM\)](#)

Tape Gateways are now compatible with Microsoft System Center 2016 Data Protection Manager (DPM). You can now use Microsoft DPM to back up your data to Amazon S3 and archive directly to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). For more information, see [Testing Your Setup by Using Microsoft System Center Data Protection Manager](#).

July 18, 2018

[Support for Server Message Block \(SMB\) protocol](#)

File Gateways added support for the Server Message Block (SMB) protocol to file shares. For more information, see [Creating a File Share](#).

June 20, 2018

[Support for file share, cached volumes, and virtual tape encryption](#)

You can now use AWS Key Management Service (AWS KMS) to encrypt data written to a file share, cached volume, or virtual tape. Currently, you can do this by using the AWS Storage Gateway API. For more information, see [Data encryption using AWS KMS](#).

June 12, 2018

[Support for NovaStor DataCenter/Network](#)

Tape Gateways now support NovaStor DataCenter/Network. You can now use NovaStor DataCenter/Network version 6.4 or 7.1 to back up your data to Amazon S3 and archive directly to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). For more information, see [Testing Your Setup by Using NovaStor DataCenter/Network](#).

May 24, 2018

Earlier updates

The following table describes important changes in each release of the *AWS Storage Gateway User Guide* before May 2018.

Change	Description	Date Changed
Support for S3 One Zone_IA storage class	For File Gateways, you can now choose S3 One Zone_IA as the default storage class for your file shares. Using this storage class, you can store your object data in a single Availability Zone in Amazon S3. For more information, see Create a file share .	April 4, 2018
New Region	Tape Gateway is now available in the Asia Pacific (Singapore) Region. For detailed information, see AWS Regions that support Storage Gateway .	April 3, 2018
Support for refresh cache notification, requester pays, and canned ACL	With File Gateways, you can now be notified when the gateway finishes refreshing the cache for your Amazon S3 bucket. For more information, see RefreshCache.html in the <i>Storage Gateway API Reference</i> .	March 1, 2018

Change	Description	Date Changed
<p>s for Amazon S3 buckets.</p>	<p>File Gateways now allow the requester or reader instead of the bucket owner to pay for access charges.</p> <p>File Gateways now allow you to give full control to the owner of the S3 bucket that maps to the NFS file share.</p> <p>For more information, see Create a file share.</p>	
<p>Support for Dell EMC NetWorker V9.x</p>	<p>Tape Gateways now support Dell EMC NetWorker V9.x. You can now use Dell EMC NetWorker V9.x to back up your data to Amazon S3 and archive directly to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). For more information, see Testing Your Setup by Using Dell EMC NetWorker.</p>	<p>February 27, 2018</p>
<p>New Region</p>	<p>Storage Gateway is now available in the Europe (Paris) Region. For detailed information, see AWS Regions that support Storage Gateway.</p>	<p>December 18, 2017</p>
<p>Support for file upload notification and guessing of the MIME type</p>	<p>File Gateways can now notify you when all files written to your NFS file share have been uploaded to Amazon S3. For more information, see NotifyWhenUploaded in the <i>Storage Gateway API Reference</i>.</p> <p>File Gateways now allow guessing of the MIME type for uploaded objects based on file extensions. For more information, see Create a file share.</p>	<p>November 21, 2017</p>
<p>Support for VMware ESXi Hypervisor version 6.5</p>	<p>AWS Storage Gateway now supports VMware ESXi Hypervisor version 6.5. This is in addition to version 4.1, 5.0, 5.1, 5.5, and 6.0. For more information, see Supported hypervisors and host requirements.</p>	<p>September 13, 2017</p>

Change	Description	Date Changed
Compatibility with Commvault 11	Tape Gateways are now compatible with Commvault 11. You can now use Commvault to back up your data to Amazon S3 and archive directly to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). For more information, see Testing Your Setup by Using Commvault .	September 12, 2017
File Gateway support for Microsoft Hyper-V hypervisor	You can now deploy a File Gateway on a Microsoft Hyper-V hypervisor. For information, see Supported hypervisors and host requirements .	June 22, 2017
Support for three to five hour tape retrieval from archive	For a Tape Gateway, you can now retrieve your tapes from archive in three to five hours. You can also determine the amount of data written to your tape from your backup application or your virtual tape library (VTL). For more information, see Viewing Tape Usage .	May 23, 2017
New Region	Storage Gateway is now available in the Asia Pacific (Mumbai) Region. For detailed information, see AWS Regions that support Storage Gateway .	May 02, 2017
Updates to file share settings Support for cache refresh for file shares	File Gateways now add mount options to the file share settings. You can now set squash and read-only options for your file share. For more information, see Create a file share . File Gateways now can find objects in the Amazon S3 bucket that were added or removed since the gateway last listed the bucket's contents and cached the results. For more information, see RefreshCache in the API Reference.	March 28, 2017

Change	Description	Date Changed
Support for cloning a volume	For cached Volume Gateways, AWS Storage Gateway now supports the ability to clone a volume from an existing volume. For more information, see Cloning a Volume .	March 16, 2017
Support for File Gateways on Amazon EC2	AWS Storage Gateway now provides the ability to deploy a File Gateway in Amazon EC2. You can launch a File Gateway in Amazon EC2 using the Storage Gateway Amazon Machine Image (AMI) now available as a community AMI. For information about how to create a File Gateway and deploy it on an EC2 instance, see Create and activate an Amazon S3 File Gateway or Create and activate an Amazon FSx File Gateway . For information about how to launch a File Gateway AMI, see Deploying an S3 File Gateway on an Amazon EC2 host or Deploying FSx File Gateway on an Amazon EC2 host .	February 08, 2017
Compatibility with Arcserve 17	Tape Gateway is now compatible with Arcserve 17. You can now use Arcserve to back up your data to Amazon S3 and archive directly to S3 Glacier Flexible Retrieval. For more information, see Testing Your Setup by Using Arcserve Backup r17.0 .	January 17, 2017
New Region	Storage Gateway is now available in the EU (London) Region. For detailed information, see AWS Regions that support Storage Gateway .	December 13, 2016
New Region	Storage Gateway is now available in the Canada (Central) Region. For detailed information, see AWS Regions that support Storage Gateway .	December 08, 2016

Change	Description	Date Changed
Support for File Gateway	In addition to Volume Gateways and Tape Gateway, Storage Gateway now provides File Gateway. File Gateway combines a service and virtual software appliance, allowing you to store and retrieve objects in Amazon S3 using industry-standard file protocols such as Network File System (NFS). The gateway provides access to objects in Amazon S3 as files on an NFS mount point.	November 29, 2016
Backup Exec 16	Tape Gateway is now compatible with Backup Exec 16. You can now use Backup Exec 16 to back up your data to Amazon S3 and archive directly to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). For more information, see Testing Your Setup by Using Veritas Backup Exec .	November 7, 2016
Compatibility with Micro Focus (HPE) Data Protector 9.x	Tape Gateway is now compatible with Micro Focus (HPE) Data Protector 9.x. You can now use HPE Data Protector to back up your data to Amazon S3 and archive directly to S3 Glacier Flexible Retrieval. For more information, see Testing Your Setup by Using Micro Focus (HPE) Data Protector .	November 2, 2016
New Region	Storage Gateway is now available in the US East (Ohio) Region. For detailed information, see AWS Regions that support Storage Gateway .	October 17, 2016
Storage Gateway console redesign	The Storage Gateway Management Console has been redesigned to make it easier to configure, manage, and monitor your gateways, volumes, and virtual tapes. The user interface now provides views that can be filtered and provides direct links to integrated AWS services such as CloudWatch and Amazon EBS. For more information, see Sign Up for AWS Storage Gateway .	August 30, 2016

Change	Description	Date Changed
Compatibility with Veeam Backup & Replication V9 Update 2 or later	Tape Gateway is now compatible with Veeam Backup & Replication V9 Update 2 or later (that is, version 9.0.0.1715 or later). You can now use Veeam Backup Replication V9 Update 2 or later to back up your data to Amazon S3 and archive directly to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). For more information, see Testing Your Setup by Using Veeam Backup & Replication .	August 15, 2016
Longer volume and snapshot IDs	Storage Gateway is introducing longer IDs for volumes and snapshots. You can activate the longer ID format for your volumes, snapshots, and other supported AWS resources. For more information, see Understanding Storage Gateway Resources and Resource IDs .	April 25, 2016
New Region Support for storage up to 512 TiB in size for stored volumes Other gateway updates and enhancements to the Storage Gateway local console	<p>Tape Gateway is now available in the Asia Pacific (Seoul) Region. For more information, see AWS Regions that support Storage Gateway.</p> <p>For stored volumes, you can now create up to 32 storage volumes up to 16 TiB in size each, for a maximum of 512 TiB of storage. For more information, see Stored volumes architecture and AWS Storage Gateway quotas.</p> <p>Total size of all tapes in a virtual tape library is increased to 1 PiB. For more information, see AWS Storage Gateway quotas.</p> <p>You can now set the password for your VM local console on the Storage Gateway Console. For information, see Setting the Local Console Password from the Storage Gateway Console.</p>	March 21, 2016

Change	Description	Date Changed
Compatibility with for Dell EMC NetWorker 8.x	Tape Gateway is now compatible with Dell EMC NetWorker 8.x. You can now use Dell EMC NetWorker to back up your data to Amazon S3 and archive directly to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). For more information, see Testing Your Setup by Using Dell EMC NetWorker .	February 29, 2016
Support for VMware ESXi Hypervisor version 6.0 and Red Hat Enterprise Linux 7 iSCSI initiator	AWS Storage Gateway now supports the VMware ESXi Hypervisor version 6.0 and the Red Hat Enterprise Linux 7 iSCSI initiator. For more information, see Supported hypervisors and host requirements and Supported iSCSI initiators .	October 20, 2015
Content restructuring	This release includes this improvement: The documentation now includes a Managing Your Activated Gateway section that combines management tasks that are common to all gateway solutions. Following, you can find instructions on how you can manage your gateway after you have deployed and activated it. For more information, see Managing Your Volume Gateway .	

Change	Description	Date Changed
<p>Support for storage up to 1,024 TiB in size for cached volumes</p> <p>Support for the VMXNET3 (10 GbE) network adapter type in VMware ESXi hypervisor</p> <p>Performance enhancements</p> <p>Miscellaneous enhancements and updates to the Storage Gateway local console</p>	<p>For cached volumes, you can now create up to 32 storage volumes at up to 32 TiB each for a maximum of 1,024 TiB of storage. For more information, see Cached volumes architecture and AWS Storage Gateway quotas.</p> <p>If your gateway is hosted on a VMware ESXi hypervisor, you can reconfigure the gateway to use the VMXNET3 adapter type. For more information, see Configuring network adapters for your gateway.</p> <p>The maximum upload rate for Storage Gateway has increased to 120 MB a second, and the maximum download rate has increased to 20 MB a second.</p> <p>The Storage Gateway local console has been updated and enhanced with additional features to help you perform maintenance tasks. For more information, see Configuring Your Gateway Network.</p>	<p>September 16, 2015</p>
<p>Support for tagging</p>	<p>Storage Gateway now supports resource tagging. You can now add tags to gateways, volumes, and virtual tapes to make them easier to manage. For more information, see Tagging Storage Gateway Resources.</p>	<p>September 2, 2015</p>

Change	Description	Date Changed
Compatibility with Quest (formerly Dell) NetVault Backup 10.0	Tape Gateway is now compatible with Quest NetVault Backup 10.0. You can now use Quest NetVault Backup 10.0 to back up your data to Amazon S3 and archive directly to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). For more information, see Testing Your Setup by Using Quest NetVault Backup .	June 22, 2015

Change	Description	Date Changed
Support for 16 TiB storage volumes for stored volumes gateway setups	Storage Gateway now supports 16 TiB storage volumes for stored volumes gateway setups. You can now create 12 16 TiB storage volumes for a maximum of 192 TiB of storage. For more information, see Stored volumes architecture .	June 3, 2015
Support for system resource checks on the Storage Gateway local console	You can now determine whether your system resources (virtual CPU cores, root volume size, and RAM) are sufficient for your gateway to function properly. For more information, see Viewing your gateway system resource status or Viewing your gateway system resource status .	
Support for the Red Hat Enterprise Linux 6 iSCSI initiator	Storage Gateway now supports the Red Hat Enterprise Linux 6 iSCSI initiator. For more information, see Requirements for setting up Volume Gateway .	
	<p>This release includes the following Storage Gateway improvements and updates:</p> <ul style="list-style-type: none"> • From the Storage Gateway console, you can now see the date and time the last successful software update was applied to your gateway. For more information, see Managing gateway updates. • Storage Gateway now provides an API you can use to list iSCSI initiators connected to your storage volumes. For more information, see ListVolumesInitiators in the API reference. 	

Change	Description	Date Changed
Support for Microsoft Hyper-V hypervisor versions 2012 and 2012 R2	Storage Gateway now supports Microsoft Hyper-V hypervisor versions 2012 and 2012 R2. This is in addition to support for Microsoft Hyper-V hypervisor version 2008 R2. For more information, see Supported hypervisors and host requirements .	April 30, 2015
Compatibility with Symantec Backup Exec 15	Tape Gateway is now compatible with Symantec Backup Exec 15. You can now use Symantec Backup Exec 15 to back up your data to Amazon S3 and archive directly to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). For more information, see Testing Your Setup by Using Veritas Backup Exec .	April 6, 2015
CHAP authentication support for storage volumes	Storage Gateway now supports configuring CHAP authentication for storage volumes. For more information, see Configure CHAP authentication for your volumes .	April 2, 2015
Support for VMware ESXi Hypervisor version 5.1 and 5.5	Storage Gateway now supports VMware ESXi Hypervisor versions 5.1 and 5.5. This is in addition to support for VMware ESXi Hypervisor versions 4.1 and 5.0. For more information, see Supported hypervisors and host requirements .	March 30, 2015
Support for Windows CHKDSK utility	Storage Gateway now supports the Windows CHKDSK utility. You can use this utility to verify the integrity of your volumes and fix errors on the volumes. For more information, see Troubleshooting volume issues .	March 04, 2015

Change	Description	Date Changed
Integration with AWS CloudTrail to capture API calls	<p>Storage Gateway is now integrated with AWS CloudTrail. AWS CloudTrail captures API calls made by or on behalf of Storage Gateway in your Amazon Web Services account and delivers the log files to an Amazon S3 bucket that you specify. For more information, see Logging and Monitoring in AWS Storage Gateway.</p> <p>This release includes the following Storage Gateway improvement and update:</p> <ul style="list-style-type: none">• Virtual tapes that have dirty data in cache storage (that is, that contain content that has not been uploaded to AWS) are now recovered when a gateway's cached drive changes. For more information, see Recovering a Virtual Tape From An Unrecoverable Gateway.	December 16, 2014

Change	Description	Date Changed
<p>Compatibility with additional backup software and medium changer</p>	<p>Tape Gateway is now compatible with the following backup software:</p> <ul style="list-style-type: none"> • Symantec Backup Exec 2014 • Microsoft System Center 2012 R2 Data Protection Manager • Veeam Backup & Replication V7 • Veeam Backup & Replication V8 <p>You can now use these four backup software products with the Storage Gateway virtual tape library (VTL) to back up to Amazon S3 and archive directly to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). For more information, see Using Your Backup Software to Test Your Gateway Setup.</p> <p>Storage Gateway now provides an additional medium changer that works with the new backup software.</p> <p>This release includes miscellaneous AWS Storage Gateway improvements and updates.</p>	<p>November 3, 2014</p>
<p>Europe (Frankfurt) Region</p>	<p>Storage Gateway is now available in the Europe (Frankfurt) Region. For detailed information, see AWS Regions that support Storage Gateway.</p>	<p>October 23, 2014</p>

Change	Description	Date Changed
Content restructure	Created a Getting Started section that is common to all gateway solutions. Following, you can find instructions for you to download, deploy, and activate a gateway. After you deploy and activate a gateway, you can proceed to further instructions specific to stored volumes, cached volumes, and Tape Gateway setups. For more information, see Creating a Tape Gateway .	May 19, 2014
Compatibility with Symantec Backup Exec 2012	Tape Gateway is now compatible with Symantec Backup Exec 2012. You can now use Symantec Backup Exec 2012 to back up your data to Amazon S3 and archive directly to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). For more information, see Testing Your Setup by Using Veritas Backup Exec .	April 28, 2014

Change	Description	Date Changed
<p>Support for Windows Server Failover Clustering</p> <p>Support for VMware ESX initiator</p> <p>Support for performing configuration tasks on Storage Gateway local console</p>	<ul style="list-style-type: none"> Storage Gateway now supports connecting multiple hosts to the same volume if the hosts coordinate access by using Windows Server Failover Clustering (WSFC). However, you can't connect multiple hosts to that same volume without using WSFC. Storage Gateway now allows you to manage storage connectivity directly through your ESX host. This provides an alternative to using initiator s resident in the guest OS of your VMs. Storage Gateway now provides support for performing configuration tasks in the Storage Gateway local console. For information about performing configuration tasks on gateways deployed on-premises, see Performing Tasks on the VM Local Console or Performing Tasks on the VM Local Console. For information about performing configuration tasks on gateways deployed on an EC2 instance, see Performing Tasks on the Amazon EC2 Local Console or Performing Tasks on the Amazon EC2 Local Console. 	<p>January 31, 2014</p>

Change	Description	Date Changed
Support for virtual tape library (VTL) and introduction of API version 2013-06-30	<p>Storage Gateway connects an on-premises software appliance with cloud-based storage to integrate your on-premises IT environment with the AWS storage infrastructure. In addition to Volume Gateways (cached volumes and stored volumes), Storage Gateway now supports gateway–virtual tape library (VTL). You can configure Tape Gateway with up to 10 virtual tape drives per gateway. Each virtual tape drive responds to the SCSI command set, so your existing on-premises backup applications will work without modification. For more information, see the following topics in the <i>AWS Storage Gateway User Guide</i>.</p> <ul style="list-style-type: none">• For an architectural overview, see How Tape Gateway works (architecture).• To get started with Tape Gateway, see Creating a Tape Gateway.	November 5, 2013
Support for Microsoft Hyper-V	<p>Storage Gateway now provides the ability to deploy an on-premises gateway on the Microsoft Hyper-V virtualization platform. Gateways deployed on Microsoft Hyper-V have all the same functionality and features as the existing on-premises Storage Gateway. To get started deploying a gateway with Microsoft Hyper-V, see Supported hypervisors and host requirements.</p>	April 10, 2013

Change	Description	Date Changed
Support for deploying a gateway on Amazon EC2	Storage Gateway now provides the ability to deploy a gateway in Amazon Elastic Compute Cloud (Amazon EC2). You can launch a gateway instance in Amazon EC2 using the Storage Gateway AMI available in AWS Marketplace . To get started deploying a gateway using the Storage Gateway AMI, see Deploy a customized Amazon EC2 instance for Volume Gateway .	January 15, 2013

Change	Description	Date Changed
Support for cached volumes and introduction of API Version 2012-06-30	<p>In this release, Storage Gateway introduces support for cached volumes. Cached volumes minimize the need to scale your on-premises storage infrastructure, while still providing your applications with low-latency access to their active data. You can create storage volumes up to 32 TiB in size and mount them as iSCSI devices from your on-premises application servers. Data written to your cached volumes is stored in Amazon Simple Storage Service (Amazon S3), with only a cache of recently written and recently read data stored locally on your on-premises storage hardware. Cached volumes allow you to utilize Amazon S3 for data where higher retrieval latencies are acceptable, such as for older, infrequently accessed data, while maintaining storage on-premises for data where low-latency access is required.</p> <p>In this release, Storage Gateway also introduces a new API version that, in addition to supporting the current operations, provides new operations to support cached volumes.</p> <p>For more information on the two Storage Gateway solutions, see How Volume Gateway works.</p> <p>You can also try a test setup. For instructions, see Creating a Tape Gateway.</p>	October 29, 2012

Change	Description	Date Changed
API and IAM support	<p>In this release, Storage Gateway introduces API support as well as support for AWS Identity and Access Management(IAM).</p> <ul style="list-style-type: none"> <p>API support—You can now programmatically configure and manage your Storage Gateway resources. For more information about the API, see API Reference for Storage Gateway in the <i>AWS Storage Gateway User Guide</i>.</p> <p>IAM support – AWS Identity and Access Management (IAM) lets you create users and manage user access to your Storage Gateway resources by means of IAM policies. For examples of IAM policies, see Identity and Access Management for AWS Storage Gateway. For more information about IAM, see AWS Identity and Access Management (IAM) detail page.</p> 	May 9, 2012
Static IP support	You can now specify a static IP for your local gateway. For more information, see Configuring Your Gateway Network .	March 5, 2012
New guide	This is the first release of <i>AWS Storage Gateway User Guide</i> .	January 24, 2012

Release notes for Volume Gateway appliance software

These release notes describe the new and updated features, improvements, and fixes that are included with each version of the Volume Gateway appliance. Each software version is identified by its release date and a unique version number.

You can determine a gateway's software version number by checking its **Details** page in the Storage Gateway console, or by calling the [DescribeGatewayInformation](#) API action using an AWS CLI command similar to the following:

```
aws storagegateway describe-gateway-information --gateway-arn
"arn:aws:storagegateway:us-west-2:123456789012:gateway/sgw-12A3456B"
```

The version number is returned in the `SoftwareVersion` field of the API response.

Note

A gateway won't report software version information under the following circumstances:

- The gateway is offline.
- The gateway is running older software that doesn't support version reporting.
- The gateway type is FSx File Gateway.

For more information about Volume Gateway updates, including how to modify the default automatic maintenance and update schedule for a gateway, see [Managing Gateway Updates Using the AWS Storage Gateway Console](#).

Release Date	Software Version	Release Notes
2024-08-30	2.11.0	<ul style="list-style-type: none">• Operating system updates for new and existing gateways
2024-07-29	2.10.0	<ul style="list-style-type: none">• Operating system updates for new and existing gateways

Release Date	Software Version	Release Notes
		<ul style="list-style-type: none">• Miscellaneous bug fixes and enhancements
2024-06-17	2.9.2	<ul style="list-style-type: none">• Operating system updates for new and existing gateways
2024-05-28	2.9.0	<ul style="list-style-type: none">• Reduced gateway restart time during software updates• Reduced the amount of data transferred for estimating network bandwidth
2024-05-08	2.8.3	<ul style="list-style-type: none">• Addressed cloud connectivity issue when using SOCKS5 proxy
2024-04-10	2.8.1	<ul style="list-style-type: none">• Addressed a memory usage issue introduced in 2.8.0• Security patch updates• Improved software update process• Addressed missing Network Time Protocol (NTP) component for new gateways
2024-03-06	2.8.0	<ul style="list-style-type: none">• Operating system updates for new gateways• Security patch updates
2023-12-19	2.7.0	<ul style="list-style-type: none">• Operating system updates for new gateways

Release Date	Software Version	Release Notes
2023-12-14	2.6.6	<ul style="list-style-type: none"><li data-bbox="1068 226 1398 260">• Maintenance release