



Tape 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: Tape 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

.....	x
What is Tape Gateway?	1
Tape Gateway	1
Are you a first-time Storage Gateway user?	2
How Tape Gateway works	2
Tape Gateways	2
Pricing	5
Plan your gateway deployment	5
Getting Started	7
Sign Up for AWS Storage Gateway	7
AWS Regions	8
Requirements	8
Hardware and storage requirements	8
Network and firewall requirements	10
Supported hypervisors and host requirements	22
Supported iSCSI initiators	23
Supported third-party backup applications	24
Accessing AWS Storage Gateway	25
Using the hardware appliance	27
Supported AWS regions	27
Setting up your hardware appliance	28
Rack-mounting and connecting the hardware appliance to power	29
Hardware appliance dimensions	29
Configuring network parameters	34
Activating your hardware appliance	37
Creating a gateway	38
Configuring an IP address for the gateway	39
Configuring your gateway	41
Removing a gateway	41
Deleting your hardware appliance	41
Creating Your Gateway	43
Overview - Gateway Activation	43
Set up gateway	43
Connect to AWS	43

Review and activate	43
Overview - Gateway Configuration	44
Overview - Storage Resources	44
Creating a Tape Gateway	44
Creating a Gateway	44
Creating Custom Tape Pools	50
Creating Tapes	53
Using Your Tape Gateway	59
Activating your gateway in a virtual private cloud	146
Creating a VPC endpoint for Storage Gateway	147
Managing Your Gateway	149
Managing Your Tape Gateway	149
Editing Gateway Information	150
Adding Tapes	150
Managing Automatic Tape Creation	150
Archiving Tapes	153
Moving a Tape from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive	153
Retrieving Archived Tapes	154
Viewing Tape Usage	156
Deleting Tapes	156
Deleting Custom Tape Pools	157
Deactivating Your Tape Gateway	158
Understanding Tape Status	159
Moving your data to a new gateway	161
Moving virtual tapes to a new Tape Gateway	162
Monitoring Storage Gateway	167
Understanding gateway metrics	167
Dimensions for Storage Gateway metrics	170
Monitoring the upload buffer	171
Monitoring cache storage	174
Understanding CloudWatch alarms	175
Creating recommended CloudWatch alarms	177
Creating a custom CloudWatch alarm	178
Monitoring Your Tape Gateway	180
Getting Tape Gateway Health Logs	180
Using Amazon CloudWatch Metrics	182

Understanding Virtual Tape Metrics	183
Measuring Performance Between Your Tape Gateway and AWS	185
Maintaining Your Gateway	189
Shutting Down Your Gateway VM	189
Starting and Stopping a Tape Gateway	190
Managing local disks	191
Deciding the amount of local disk storage	191
Optimize Performance	193
Sizing the upload buffer	193
Sizing cache storage	194
Add upload buffer or cache storage	195
Managing Bandwidth	196
Changing Bandwidth Throttling Using the Storage Gateway Console	197
Scheduling Bandwidth Throttling	197
Using the AWS SDK for Java	199
Using the AWS SDK for .NET	201
Using the AWS Tools for Windows PowerShell	203
Managing Gateway Updates	204
Turn maintenance updates on or off	205
Modify the gateway maintenance window schedule	206
Performing Maintenance Tasks on the Local Console	207
Performing Tasks on the VM Local Console	207
Performing Tasks on the EC2 Local Console	225
Accessing the Gateway Local Console	231
Configuring Network Adapters for Your Gateway	236
Deleting Your Gateway and Removing Resources	240
Deleting Your Gateway by Using the Storage Gateway Console	240
Removing Resources from a Gateway Deployed On-Premises	242
Removing Resources from a Gateway Deployed on an Amazon EC2 Instance	243
Performance	245
Performance guidance for Tape Gateways	245
Optimizing Gateway Performance	248
Recommended Configuration	248
Add Resources to Your Gateway	249
Optimize iSCSI Settings	252
Use a Larger Block Size for Tape Drives	252

Optimize the Performance of Virtual Tape Drives	252
Add Resources to Your Application Environment	253
Using VMware High Availability with Storage Gateway	254
Configure Your vSphere VMware HA Cluster	254
Download the .ova Image from the Storage Gateway console	256
Deploy the Gateway	256
(Optional) Add Override Options for Other VMs on Your Cluster	257
Activate Your Gateway	258
Test Your VMware High Availability Configuration	258
Security	260
Data protection	261
Data encryption	261
Identity and Access Management	263
Audience	264
Authenticating with identities	264
Managing access using policies	268
How AWS Storage Gateway works with IAM	270
Identity-based policy examples	277
Troubleshooting	280
Logging and Monitoring	281
Storage Gateway Information in CloudTrail	282
Understanding Storage Gateway Log File Entries	283
Compliance validation	285
Resilience	286
Infrastructure Security	286
AWS Security Best Practices	287
Troubleshooting gateway issues	288
Troubleshooting: gateway offline issues	288
Check the associated firewall or proxy	289
Check for an ongoing SSL or deep-packet inspection of your gateway's traffic	289
Check for a power outage or hardware failure on the hypervisor host	289
Check for issues with an associated cache disk	289
Troubleshooting: gateway activation issues	290
Resolve errors when activating your gateway using a public endpoint	291
Resolve errors when activating your gateway using an Amazon VPC endpoint	294

Resolve errors when activating your gateway using a public endpoint and there is a Storage Gateway VPC endpoint in the same VPC	298
Troubleshooting on-premises gateway issues	298
Activating AWS Support to help troubleshoot your gateway	303
Troubleshooting Microsoft Hyper-V setup issues	304
Troubleshooting Amazon EC2 gateway issues	308
Gateway activation hasn't occurred after a few moments	309
Can't find the EC2 gateway instance in the instance list	309
Can't attach a an Amazon EBS volume to the EC2 gateway instance	310
No disks available when you try to add storage volumes message	310
How to remove a disk allocated as upload buffer space to reduce upload buffer space	310
Throughput to or from the EC2 gateway drops to zero	310
Activating AWS Support to help troubleshoot the gateway	311
Connect to your Amazon EC2 gateway using the serial console	312
Troubleshooting hardware appliance issues	313
How to determine service IP address	313
How to perform a factory reset	313
How to perform a remote restart	313
How to obtain Dell iDRAC support	313
How to find the hardware appliance serial number	313
How to get hardware appliance support	314
Troubleshooting virtual tape issues	315
Recovering a Virtual Tape From An Unrecoverable Gateway	315
Troubleshooting Irrecoverable Tapes	319
High Availability Health Notifications	320
Troubleshooting high availability issues	320
Health notifications	320
Metrics	322
Recovering your data: best practices	322
Recovering from an unexpected VM shutdown	323
Recovering data from malfunctioning gateway or VM	323
Recovering data from an irrecoverable tape	324
Recovering data from a malfunctioning cache disk	324
Recovering data from an inaccessible data center	324
Additional Resources	326
Host Setup	326

Configuring VMware for Storage Gateway	326
Synchronizing Your Gateway VM Time	333
Deploy an Amazon EC2 host for Tape Gateway	335
Deploy Amazon EC2 with Default Settings	339
Modify Amazon EC2 instance metadata options	342
Tape Gateway	342
Removing Disks from Your Gateway	342
EBS Volumes for EC2 Gateways	346
Working with VTL Devices	347
Working with Tapes	352
Getting Activation Key	354
Linux (curl)	355
Linux (bash/zsh)	356
Microsoft Windows PowerShell	357
Using your local console	357
Connecting iSCSI Initiators	358
Connecting to VTL Devices	359
Connecting Your Volumes or VTL Devices to a Linux Client	365
Customizing iSCSI Settings	367
Configuring CHAP Authentication	375
Using AWS Direct Connect with Storage Gateway	384
Port Requirements	384
Connecting to Your Gateway	391
Getting an IP Address from an Amazon EC2 Host	391
Understanding Resources and Resource IDs	393
Working with Resource IDs	393
Tagging Your Resources	394
Working with Tags	394
Open-Source Components	396
Storage Gateway quotas	396
Quotas for tapes	396
Recommended local disk sizes for your gateway	397
API Reference	398
Required Request Headers	398
Signing Requests	401
Example Signature Calculation	401

Error Responses	403
Exceptions	404
Operation Error Codes	406
Error Responses	425
Operations	427
Document history	428
Earlier updates	444
Release Notes	463

Amazon S3 File Gateway documentation has been moved to [What is Amazon S3 File Gateway?](#)

Amazon FSx File Gateway documentation has been moved to [What is Amazon FSx File Gateway?](#)

Volume Gateway documentation has been moved to [What is Volume Gateway?](#)

What is Tape 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.

AWS Storage Gateway offers file-based File Gateways (Amazon S3 File and Amazon FSx File), volume-based (Cached and Stored), and tape-based storage solutions.

Topics

- [Tape Gateway](#)
- [Are you a first-time Storage Gateway user?](#)
- [How Tape Gateway works \(architecture\)](#)
- [Storage Gateway pricing](#)
- [Plan your Storage Gateway deployment](#)

Tape Gateway

Tape Gateway – A Tape Gateway provides cloud-backed virtual tape storage.

With a Tape Gateway, you can cost-effectively and durably archive backup data in S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive. A Tape Gateway provides a virtual tape infrastructure that scales seamlessly with your business needs and eliminates the operational burden of provisioning, scaling, and maintaining a physical tape infrastructure.

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 deploy your gateway on an EC2 instance to provision iSCSI storage volumes in AWS. You can use gateways hosted on EC2 instances for disaster recovery, data mirroring, and providing storage for applications hosted on Amazon EC2.

For an architectural overview, see [How Tape Gateway works \(architecture\)](#). To see the wide range of use cases that AWS Storage Gateway helps make possible, see [AWS Storage Gateway](#).

Documentation: For Tape Gateway documentation, see [Creating a Tape Gateway](#).

Are you a first-time Storage Gateway user?

In the following documentation, you can find a Getting Started section that covers setup information common to all gateways and also gateway-specific setup sections. The Getting Started section shows you how to deploy, activate, and configure storage for a gateway. The management section shows you how to manage your gateway and resources:

- [Creating a Tape Gateway](#) provides instructions on how to create and use a Tape Gateway. It shows you how to back up data to virtual tapes and archive the tapes.
- [Managing Your Gateway](#) describes how to perform management tasks for your gateway and its resources.

In this guide, you can primarily find how to work with gateway operations by using the AWS Management Console. If you want to perform these operations programmatically, see the [AWS Storage Gateway API Reference](#).


How Tape Gateway works (architecture)

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

Tape Gateways

Tape Gateway offers a durable, cost-effective solution to archive your data in the Amazon Web Services Cloud. With its virtual tape library (VTL) interface, you use your existing tape-based backup infrastructure to store data on virtual tape cartridges that you create on your Tape Gateway. Each Tape Gateway is preconfigured with a media changer and tape drives. These are available to your existing client backup applications as iSCSI devices. You add tape cartridges as you need to archive your data.

The following diagram provides an overview of Tape Gateway deployment.



The diagram identifies the following Tape Gateway components:

- **Virtual tape** – A virtual tape is like a physical tape cartridge. However, virtual tape data is stored in the Amazon Web Services Cloud. Like physical tapes, virtual tapes can be blank or can have data written on them. You can create virtual tapes either by using the Storage Gateway console or programmatically by using the Storage Gateway API. Each gateway can contain up to 1,500 tapes or up to 1 PiB of total tape data at a time. The size of each virtual tape, which you can configure when you create the tape, is between 100 GiB and 15 TiB.
- **Virtual tape library (VTL)** – A VTL is like a physical tape library available on-premises with robotic arms and tape drives. Your VTL includes the collection of stored virtual tapes. Each Tape Gateway comes with one VTL.

The virtual tapes that you create appear in your gateway's VTL. Tapes in the VTL are backed up by Amazon S3. As your backup software writes data to the gateway, the gateway stores data locally and then asynchronously uploads it to virtual tapes in your VTL—that is, Amazon S3.

- **Tape drive** – A VTL tape drive is analogous to a physical tape drive that can perform I/O and seek operations on a tape. Each VTL comes with a set of 10 tape drives, which are available to your backup application as iSCSI devices.
- **Media changer** – A VTL media changer is analogous to a robot that moves tapes around in a physical tape library's storage slots and tape drives. Each VTL comes with one media changer, which is available to your backup application as an iSCSI device.
- **Archive** – Archive is analogous to an offsite tape holding facility. You can archive tapes from your gateway's VTL to the archive. If needed, you can retrieve tapes from the archive back to your gateway's VTL.
 - **Archiving tapes** – When your backup software ejects a tape, your gateway moves the tape to the archive for long-term storage. The archive is located in the AWS Region in which you activated the gateway. Tapes in the archive are stored in the virtual tape shelf (VTS). The VTS is backed by [S3 Glacier Flexible Retrieval](#) or [S3 Glacier Deep Archive](#), low-cost storage service for data archiving, backup, and long-term data retention.
 - **Retrieving tapes** – You can't read archived tapes directly. To read an archived tape, you must first retrieve it to your Tape Gateway by using either the Storage Gateway console or the Storage Gateway API.

Important

If you archive a tape in S3 Glacier Flexible Retrieval, you can retrieve the tape typically within 3-5 hours. If you archive the tape in S3 Glacier Deep Archive, you can retrieve it typically within 12 hours.

After you deploy and activate a Tape Gateway, you mount the virtual tape drives and media changer on your on-premises application servers as iSCSI devices. You create virtual tapes as needed. Then you use your existing backup software application to write data to the virtual tapes. The media changer loads and unloads the virtual tapes into the virtual tape drives for read and write operations.

Allocating local disks for the gateway VM

Your gateway VM needs local disks, which you allocate for the following purposes:

- **Cache storage** – The cache storage acts as the durable store for data that is waiting to upload to Amazon S3 from the upload buffer.

If your application reads data from a virtual tape, the gateway saves the data to the cache storage. The gateway stores recently accessed data in the cache storage for low-latency access. If your application requests tape data, the gateway first checks the cache storage for the data before downloading the data from AWS.

- **Upload buffer** – The upload buffer provides a staging area for the gateway before it uploads the data to a virtual tape. The upload buffer is also critical for creating recovery points that you can use to recover tapes from unexpected failures. For more information, see [You Need to Recover a Virtual Tape from a Malfunctioning Tape Gateway](#).

As your backup application writes data to your gateway, the gateway copies data to both the cache storage and the upload buffer. It then acknowledges completion of the write operation to your backup application.

For guidelines on the amount of disk space to allocate for the cache storage and upload buffer, see [Deciding the amount of local disk storage](#).

Storage Gateway pricing

For current information about pricing, see [Pricing](#) on the AWS Storage Gateway details page.

Plan your Storage Gateway deployment

By using the Storage Gateway software appliance, you can connect your existing on-premises application infrastructure with scalable, cost-effective AWS cloud storage that provides data security features.

To deploy Storage Gateway, you first need to decide on the following two things:

1. **Your gateway type** – this guide covers the following gateway type:
 - **Tape Gateway** – If you are looking for a cost-effective, durable, long-term, offsite alternative for data archiving, deploy a Tape Gateway. With its virtual tape library (VTL) interface, you can use your existing tape-based backup software infrastructure to store data on virtual tape cartridges that you create. For more information, see [Supported third-party backup applications for a Tape Gateway](#). When you archive tapes, you don't worry about managing tapes on your premises and arranging shipments of tapes offsite. For an architectural overview, see [How Tape Gateway works \(architecture\)](#).

2. **Hosting option** – You can run Storage Gateway either on-premises as a VM appliance or hardware appliance, or in AWS as an Amazon EC2 instance. For more information, see [Requirements](#). If your data center goes offline and you don't have an available host, you can deploy a gateway on an EC2 instance. Storage Gateway provides an Amazon Machine Image (AMI) that contains the gateway VM image.

Additionally, as you configure a host to deploy a gateway software appliance, you need to allocate sufficient storage for the gateway VM.

Before you continue to the next step, make sure that you have done the following:

1. For a gateway deployed on-premises, choose the type of VM host and set it up. Your options are VMware ESXi Hypervisor, Microsoft Hyper-V, and Linux Kernel-based Virtual Machine (KVM). If you deploy the gateway behind a firewall, make sure that ports are accessible to the gateway VM. For more information, see [Requirements](#).
2. Install your client backup software. For more information, see [Supported third-party backup applications for a Tape Gateway](#).

Getting Started

In this section, you can find instructions about how to get started with Storage Gateway. To get started, you first sign up for AWS. If you are a first-time user, we recommend that you read the regions and requirements section.

Topics

- [Sign Up for AWS Storage Gateway](#)
- [AWS Regions](#)
- [Requirements](#)
- [Accessing AWS Storage Gateway](#)

Sign Up for AWS Storage Gateway

To use Storage Gateway, you need an Amazon Web Services account that gives you access to all AWS resources, forums, support, and usage reports. You aren't charged for any of the services unless you use them. If you already have an Amazon Web Services account, you can skip this step.

To sign up for Amazon Web Services 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](#).

For information about pricing, see [Pricing](#) on the Storage Gateway details page.

AWS Regions

Storage Gateway stores volume, snapshot, tape, and file data in the AWS Region in which your gateway is activated. File data is stored in the AWS Region where your Amazon S3 bucket is located. You select an AWS Region at the upper right of the Storage Gateway Management Console before you start deploying your gateway.

- 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

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](#)
- [Supported third-party backup applications for a Tape Gateway](#)

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 Tape 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**.

For Tape 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
Tape Gateway	150 GiB	64 TiB	150 GiB	2 TiB	—

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](#).

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.

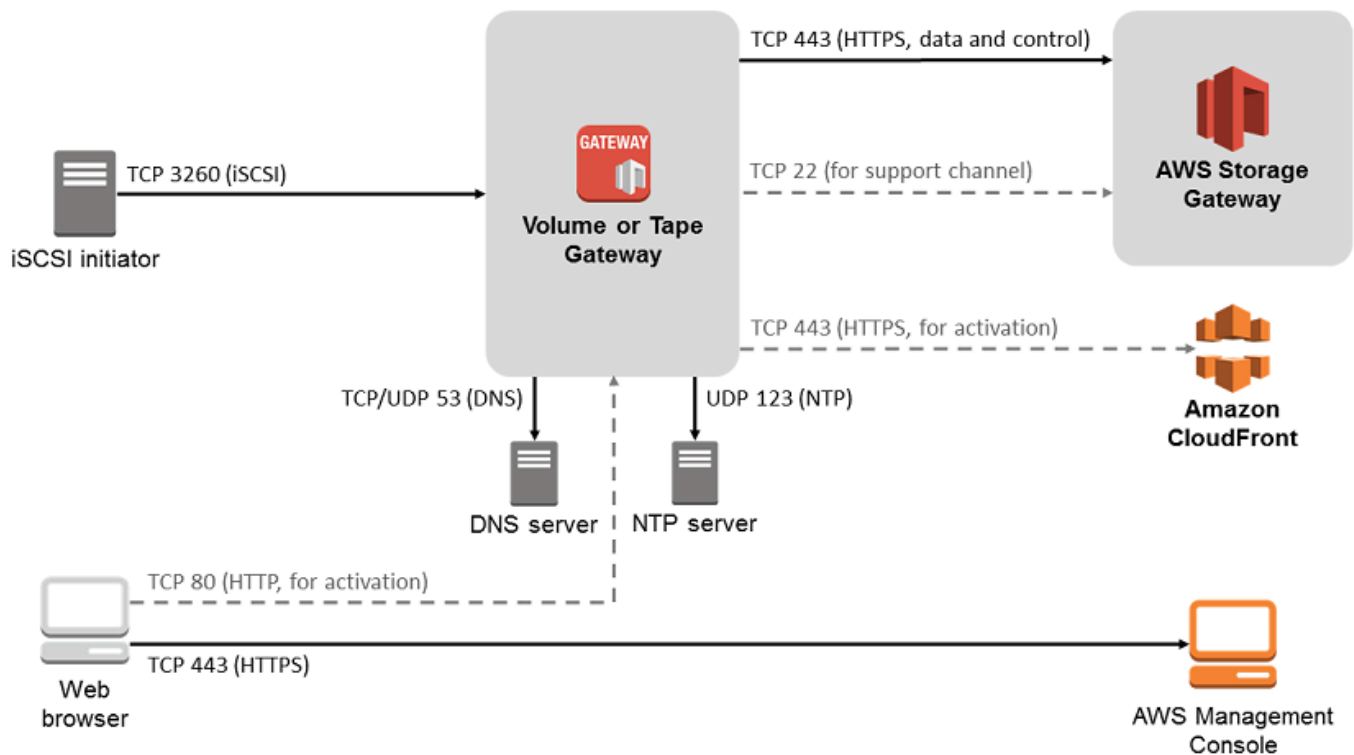
Protocol	Port	Direction	Source	Destination	How Used
TCP	80 (HTTP)	Inbound	The host from which you connect to the AWS Management Console.	Storage Gateway	<p>By local systems to obtain the Storage Gateway activation key. Port 80 is only used during activation of the Storage Gateway appliance.</p> <p>Storage 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 Tape Gateway.



In addition to the common ports, Tape 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](#) in the *Additional Storage Gateway resources* section.

Networking and firewall requirements for the Storage Gateway Hardware Appliance

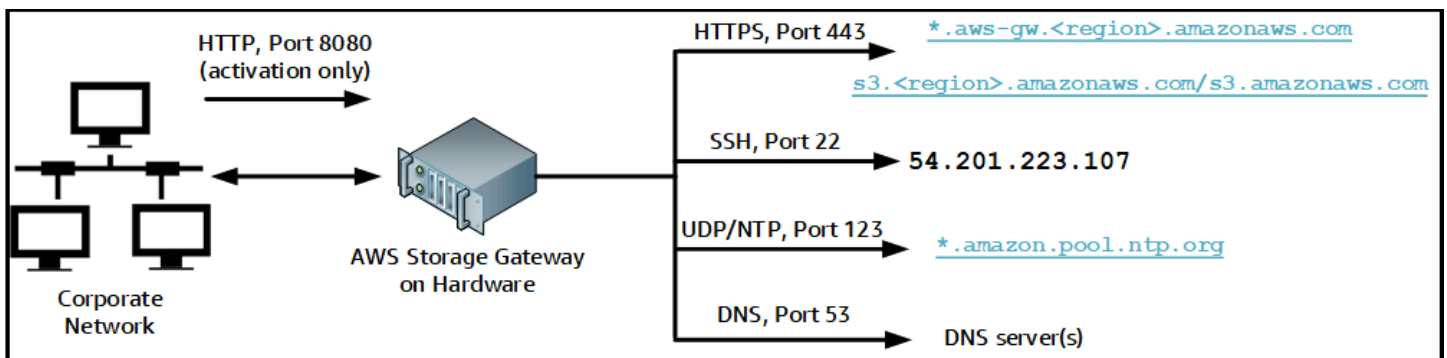
Each Storage Gateway Hardware Appliance requires the following network services:

- **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 network parameters](#).

Note

For an illustration showing the back of the server with its ports, see [Rack-mounting your hardware appliance and connecting it to power](#)

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](#).

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) – A free version of VMware is available on the [VMware website](#). 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 [Deploying an Amazon EC2 instance to host your Tape 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 Tape Gateway, the gateway is preconfigured with one media changer and 10 tape drives. These tape drives and the media changer are available to your existing client backup applications as iSCSI devices.

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.

Supported third-party backup applications for a Tape Gateway

You use a backup application to read, write, and manage tapes with a Tape Gateway. The following third-party backup applications are supported to work with Tape Gateways.

The type of medium changer you choose depends on the backup application you plan to use. The following table lists third-party backup applications that have been tested and found to be compatible with Tape Gateways. This table includes the medium changer type recommended for each backup application.

Backup Application	Medium Changer Type
Arcserve Backup	AWS-Gateway-VTL
Bacula Enterprise V10.x	AWS-Gateway-VTL or STK-L700
Commvault V11	STK-L700
Dell EMC NetWorker 19.5	AWS-Gateway-VTL
IBM Spectrum Protect v8.1.10	IBM-03584L32-0402
Micro Focus (HPE) Data Protector 9 or 11.x	AWS-Gateway-VTL
Microsoft System Center 2012 R2 or 2016 Data Protection Manager	STK-L700

Backup Application	Medium Changer Type
NovaStor DataCenter/Network 6.4 or 7.1	STK-L700
Quest NetVault Backup 12.4 or 13.x	STK-L700
Veeam Backup & Replication 11A	AWS-Gateway-VTL
Veritas Backup Exec 2014 or 15 or 16 or 20 or 22.x	AWS-Gateway-VTL
Veritas Backup Exec 2012	STK-L700
<div style="border: 1px solid #00a0e3; border-radius: 10px; padding: 10px; margin: 10px 0;"> <p>Note</p> <p>Veritas has ended support for Backup Exec 2012.</p> </div>	
Veritas NetBackup Version 7.x or 8.x	AWS-Gateway-VTL

⚠ Important

We highly recommend that you choose the medium changer that's listed for your backup application. Other medium changers might not function properly. You can choose a different medium changer after the gateway is activated. For more information, see [Selecting a Medium Changer After Gateway Activation](#).

Accessing AWS Storage Gateway

You can use the [Storage Gateway Management Console](#) to perform various gateway configuration and management tasks. The Getting Started section and various other sections of this guide use the console to illustrate gateway functionality.

To allow browser access to the Storage Gateway console, ensure that your browser has access to the Storage Gateway API endpoint. For more information, see [Storage Gateway endpoints and quotas](#) in the *AWS General Reference*.

Additionally, you can use the AWS Storage Gateway API to programmatically configure and manage your gateways. For more information about the API, see [API Reference for Storage Gateway](#).

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 [Sample Code Libraries](#).

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 your hardware appliances from the **Hardware appliance overview** page on 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 your hardware appliance with your Amazon Web Services account. After activation, your hardware appliance appears in the console as a gateway on the **Hardware appliance overview** page. You can configure your hardware appliance as a 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.

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

Topics

- [Supported AWS regions](#)
- [Setting up your hardware appliance](#)
- [Rack-mounting your hardware appliance and connecting it to power](#)
- [Configuring network parameters](#)
- [Activating your hardware appliance](#)
- [Creating a gateway](#)
- [Configuring an IP address for the gateway](#)
- [Configuring your gateway](#)
- [Removing a gateway from the hardware appliance](#)
- [Deleting your hardware appliance](#)

Supported AWS regions

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*.

Setting up your 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 Amazon Web Services 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.

Note

It is your responsibility to ensure the hardware appliance firmware is up-to-date.

To install and configure your hardware appliance

1. Rack-mount the appliance, and plug in power and network connections. For more information, see [Rack-mounting your hardware appliance and connecting it to power](#).
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 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 hardware appliance](#).
4. Install the Storage Gateway on your hardware appliance. For more information, see [Configuring your 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 Amazon Web Services 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

Rack-mounting your hardware appliance and connecting it to power


After you unbox your Storage Gateway Hardware Appliance, follow the instructions contained in the box to rack-mount the server. Your appliance has a 1U form factor and fits in a standard International Electrotechnical Commission (IEC) compliant 19-inch rack.

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.

Hardware appliance dimensions

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.


To connect the hardware appliance to power

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](#).

1. 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.

In the following image, you can see the hardware appliance with the different connections. hardware appliance rear with network and power connector labels.



hardware appliance rear with network and power connector labels.

2. 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.

3. Plug in the keyboard and monitor.
4. 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.


After the server boots up, the hardware console appears on the monitor. The hardware console presents a user interface specific to AWS that you can use to configure initial network parameters. You configure these parameters to connect the appliance to AWS and open up a support channel for troubleshooting by Amazon Web Services Support.

To work with the hardware 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.

To set a password for the first time

1. For **Set Password**, enter a password, and then press Down arrow.
2. For **Confirm**, re-enter your password, and then choose **Save Password**.

hardware appliance console set password dialog screen.



hardware appliance console set password dialog screen.

At this point, you are in the hardware console, shown following.

hardware appliance console main menu showing connections and menu options.



hardware appliance console main menu showing connections and menu options.

Next step

[Configuring network parameters](#)

Configuring network parameters

After the server boots up, you can enter your first password in the hardware console as described in [Rack-mounting your hardware appliance and connecting it to power](#).

Next, on the hardware console take the following steps to configure network parameters so your hardware appliance can connect to AWS.

To set a network address

1. Choose **Configure Network** and press the Enter key. The **Configure Network** screen shown following appears.
hardware appliance console configure network screen.



hardware appliance console configure network screen.

- For **IP Address**, enter a valid IPv4 address from one of the following sources:
 - Use the IPv4 address assigned by your Dynamic Host Configuration Protocol (DHCP) server to your physical network port.

If you do so, note this IPv4 address for later use in the activation step.

- Assign a static IPv4 address. To do so, choose **Static** in the em1 section and press Enter to view the Configure Static IP screen shown following.

The em1 section is at upper left section in the group of port settings.

After you have entered a valid IPv4 address, press the Down arrow or Tab.

Note

If you configure any other interface, it must provide the same always-on connection to the AWS endpoints listed in the requirements.

hardware appliance console configure NIC to static IP screen.



hardware appliance console configure NIC to static IP screen.

3. For **Subnet**, enter a valid subnet mask, and then press Down arrow.
4. For **Gateway**, enter your network gateway's IPv4 address, and then press Down arrow.
5. For **DNS1**, enter the IPv4 address for your Domain Name Service (DNS) server, and then press Down arrow.
6. (Optional) For **DNS2**, enter a second IPv4 address, and then press Down arrow. A second DNS server assignment would provide additional redundancy should the first DNS server become unavailable.
7. Choose **Save** and then press Enter to save your static IPv4 address setting for the appliance.

To log out of the hardware console

1. Choose **Back** to return to the Main screen.
2. Choose **Logout** to return to the Login screen.

Next step

[Activating your hardware appliance](#)

Activating your 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](#)

Creating a gateway

You can create an S3 File Gateway, FSx File Gateway, Tape Gateway, or Volume Gateway on the hardware appliance.

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. Choose **Hardware**.
3. Select the activated hardware appliance on which you want to create your gateway, then choose **Create Gateway**.
4. Follow the procedures described in [Creating Your Gateway](#) to set up, connect, and configure your chosen gateway type.

When you finish creating your gateway in the Storage Gateway console, the Storage Gateway software automatically starts installing on the hardware appliance. It can take 5–10 minutes for a gateway to display as **online** in the console.

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 an IP address for the gateway](#)

Configuring an IP address for the gateway

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, your iSCSI initiator, and so on) 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** to open a login screen for the gateway local console.
2. Enter the localhost **login** password, and then press Enter.

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

3. Change the default password. Choose **Actions** then **Set Local Password** and enter your new credentials in the **Set Local Password** dialog box.
4. (Optional) Configure your proxy settings. See [the section called "Setting the Local Console Password from the Storage Gateway Console"](#) for instructions.
5. Navigate to the Network Settings page of the gateway local console as shown following. gateway local console configuration page showing options including network configuration.

```

AWS Storage Gateway Configuration

#####
## Currently connected network adapters:
##
## eth0: 10.0.0.45
#####

1: SOCKS Proxy Configuration
2: Network Configuration
3: Test Network Connectivity
4: System Time Management
5: Gateway Console
6: View System Resource Check (0 Errors)

0: Stop AWS Storage Gateway

Press "x" to exit session

Enter command: _

```

gateway local console configuration page showing options including network configuration.

6. Type 2 to go to the **Network Configuration** page shown following.

gateway local console network configuration page with DHCP and static IP options.

```

AWS Storage Gateway Network Configuration

1: Describe Adapter
2: Configure DHCP
3: Configure Static IP
4: Reset all to DHCP
5: Set Default Adapter
6: View DNS Configuration
7: View Routes

Press "x" to exit

Enter command: _

```

gateway local console network configuration page with DHCP and static IP options.

7. Configure a static or DHCP IP address for the network port on your hardware appliance to present a file, volume, and Tape Gateway for applications. This IP address 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.

Next step

[Configuring your gateway](#)

Configuring your gateway

After your hardware appliance has been activated and configured, your appliance appears in the console. Now you can create the type of gateway that you want. Continue the installation on the **Configure gateway** page for your gateway type. For instructions, see [Configure your Tape Gateway](#).

Removing a gateway from the hardware appliance

To remove gateway software from your hardware appliance, use the following procedure. After you do so, the gateway software is uninstalled from your hardware appliance.

To remove a gateway from a hardware appliance

1. On the **Hardware** page of the Storage Gateway console, choose the hardware appliance you want to delete.
2. For **Actions**, choose **Remove Gateway**. The confirmation dialog box appears.
3. Verify that you want to remove the gateway software from specified hardware appliance, then type the word *remove* in the confirmation box and choose **Remove**.

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 by Using the AWS Storage Gateway Console 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 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 a gateway from the 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 topics 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 [Creating a Tape Gateway](#).

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 Tape Gateway

In this section, you can find instructions on how to create and use a Tape Gateway in AWS Storage Gateway.

Topics

- [Creating a Gateway](#)
- [Creating a Custom Tape Pool](#)
- [Creating Tapes](#)
- [Using Your Tape Gateway](#)

Creating a Gateway

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

Topics

- [Set up a Tape Gateway](#)

- [Connect your Tape Gateway to AWS](#)
- [Review settings and activate your Tape Gateway](#)
- [Configure your Tape Gateway](#)

Set up a Tape Gateway

To set up a new Tape 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 **Tape Gateway**.
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. In the **Backup application settings** section, for **Backup application**, choose the application you want to use to backup your tape data to the virtual tapes associated with your Tape Gateway.
7. 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 Tape Gateway to AWS](#).

Connect your Tape Gateway to AWS

To connect a new Tape Gateway to AWS

1. Complete the procedure described in [Set up a Tape 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 Tape Gateway](#).

Review settings and activate your Tape Gateway

To activate a new Tape Gateway

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

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 Tape Gateway](#).

Configure your Tape Gateway

To perform first-time configuration on a new Tape Gateway

1. Complete the procedures described in the following topics if you have not done so already:
 - [Set up a Tape Gateway](#)
 - [Connect your Tape Gateway to AWS](#)
 - [Review settings and activate your Tape 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 virtual tapes for it to use. For instructions, see [Creating Tapes](#).

Creating a Custom Tape Pool

This section describes how to create a new custom tape pool in AWS Storage Gateway.

Topics

- [Choosing a Tape Pool Type](#)
- [Using Tape Retention Lock](#)
- [Creating a Custom Tape Pool](#)

Choosing a Tape Pool Type

AWS Storage Gateway uses tape pools to determine the storage class that you want tapes to be archived in when they are ejected. Storage Gateway provides two standard tape pools:

- **Glacier Pool** – Archives the tape in the S3 Glacier Flexible Retrieval storage class. When your backup software ejects the tape, it is automatically archived in S3 Glacier Flexible Retrieval.

You use S3 Glacier Flexible Retrieval for more active archives, where you can retrieve the tapes typically within 3-5 hours. For more information, see [Storage classes for archiving objects](#) in the *Amazon Simple Storage Service User Guide*.

- **Deep Archive Pool** – Archives the tape in the S3 Glacier Deep Archive storage class. When your backup software ejects the tape, the tape is automatically archived in S3 Glacier Deep Archive. You use S3 Glacier Deep Archive for long-term data retention and digital preservation, where data is accessed once or twice a year. You can retrieve tapes archived in S3 Glacier Deep Archive typically within 12 hours. For detailed information, see [Storage classes for archiving objects](#) in the *Amazon Simple Storage Service User Guide*.

If you archive a tape in S3 Glacier Flexible Retrieval, you can move it to S3 Glacier Deep Archive later. For more information, see [Moving Your Tape from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive Storage Class](#).

Storage Gateway also supports creation of custom tape pools, which allow you to activate tape retention lock to prevent archived tapes from being deleted or moved to another pool for a fixed amount of time, up to 100 years. This includes locking permission controls on who can delete tapes or modify retention settings.

Using Tape Retention Lock

With tape retention lock, you can lock archived tapes. Tape retention lock is an option for tapes in a custom tape pool. Tapes that have tape retention lock activated can't be deleted or moved to another pool for a fixed amount of time, up to 100 years.

You can configure tape retention lock in one of two modes:

- **Governance mode** – When configured in governance mode, only AWS Identity and Access Management (IAM) users with the permissions to perform `storagegateway:BypassGovernanceRetention` can remove tapes from the pool. If you're using the AWS Storage Gateway API to remove the tape, you must also set `BypassGovernanceRetention` to `true`.
- **Compliance mode** – When configured in compliance mode, the protection cannot be removed by any user, including the root AWS account.

When a tape is locked in compliance mode, its retention lock type can't be changed, and its retention period can't be shortened. The compliance mode lock type helps ensure that a tape can't be overwritten or deleted for the duration of the retention period.

⚠ Important

A custom pool's configuration cannot be changed after it is created.

You can activate tape retention lock when you create a custom tape pool. Any new tapes that are attached to a custom pool inherit the retention lock type, period, and storage class for that pool.

You can also activate tape retention lock on tapes that were archived before the release of this feature by moving tapes between the default pool and a custom pool that you create. If the tape is archived, the tape retention lock is effective immediately.

ℹ Note

If you're moving archived tapes between the S3 Glacier Flexible Retrieval and S3 Glacier Deep Archive storage classes, you are charged a fee for moving a tape. There is no additional charge to move a tape from a default pool to a custom pool if the storage class remains the same.

Creating a Custom Tape Pool

Use the following steps to create a custom tape pool using the AWS Storage Gateway console.

To create a custom tape pool

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the left navigation pane, choose the **Tape Library** tab, and then choose the **Pools** tab.
3. Choose **Create pool** to open the **Create pool** pane.
4. For **Name**, enter a unique name to identify your custom tape pool. The pool name must be between 2 and 100 characters long.
5. For **Storage class**, choose **Glacier** or **Glacier Deep Archive**.
6. For **Retention lock type**, choose **None**, **Compliance**, or **Governance**.

Note

If you choose **Compliance**, tape retention lock cannot be removed by any user, including the root AWS account.

7. If you choose a tape retention lock type, enter the **Retention period** in days. The maximum retention period is 36,500 days (100 years).
8. (Optional) For **Tags**, choose **Add new tag** to add a tag to your custom tape pool. A tag is a case-sensitive key-value pair that helps you manage, filter, and search for your custom tape pools.

Enter a **Key**, and optionally, a **Value** for your tag. You can add up to 50 tags to the tape pool.

9. Choose **Create pool** to create your new custom tape pool.

Creating Tapes

This section describes how to create new virtual tapes using AWS Storage Gateway. You can create new virtual tapes manually using either the AWS Storage Gateway console or the Storage Gateway API. You can also configure your Tape Gateway to create them automatically, which helps decrease the need for manual tape management, makes your large deployments simpler, and helps scale on-premises and archive storage needs.

Tape Gateway supports *write once, read many* (WORM) and *tape retention lock* on virtual tapes. WORM-activated virtual tapes help ensure that the data on active tapes in your virtual tape library cannot be overwritten or erased. For more information about WORM protection for virtual tapes, see the section following, [the section called "WORM Tape Protection"](#).

With tape retention lock, you can 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 the retention settings. For more information about tape retention lock, see [the section called "Tape Retention Lock"](#).

Note

You are charged only for the amount of data that you write to the tape, not the tape capacity.

You can use AWS Key Management Service (AWS KMS) to encrypt data written to a virtual tape that is stored in Amazon Simple Storage Service (Amazon S3). Currently, you can do this by using the AWS Storage Gateway API or AWS Command Line Interface (AWS CLI). For more information, see [CreateTapes](#) or [create-tapes](#).

Topics

- [Write Once, Read Many \(WORM\) Tape Protection](#)
- [Creating Tapes Manually](#)
- [Allowing Automatic Tape Creation](#)

Write Once, Read Many (WORM) Tape Protection

You can prevent virtual tapes from being overwritten or erased by activating WORM protection for virtual tapes in AWS Storage Gateway. WORM protection for virtual tapes is activated when creating tapes.

Data that is written to WORM virtual tapes can't be overwritten. Only new data can be appended to WORM virtual tapes, and existing data can't be erased. Activating WORM protection for virtual tapes helps protect those tapes while they are in active use, before they are ejected and archived.

WORM configuration can only be set when tapes are created, and that configuration cannot be changed after the tapes are created.


Creating Tapes Manually

You can create new virtual tapes manually using either the AWS Storage Gateway console or the Storage Gateway API. The console offers a convenient interface for tape creation with the flexibility to specify a prefix for a randomly-generated tape barcode. If you need to fully customize your tape barcodes (for example, to match the serial number of a corresponding physical tape), you must use the API. For more information on creating tapes using the Storage Gateway API, see [CreateTapeWithBarcode](#) in the *Storage Gateway API Reference*.

To create virtual tapes manually using the Storage Gateway console

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


2. In the navigation pane, choose the **Gateways** tab.
3. Choose **Create tapes** to open the **Create tapes** pane.
4. For **Gateway**, choose a gateway. The tape is created for this gateway.
5. For **Tape type**, choose **Standard** to create standard virtual tapes. Choose **WORM** to create *write once read many* (WORM) virtual tapes.
6. For **Number of tapes**, choose the number of tapes that you want to create. For more information about tape quotas, see [AWS Storage Gateway quotas](#).
7. For **Capacity**, enter the size of the virtual tape that you want to create. Tapes must be larger than 100 GiB. For information about capacity quotas, see [AWS Storage Gateway quotas](#).
8. For **Barcode prefix**, enter the prefix that you want to prepend to the barcode of your virtual tapes.

 **Note**

Virtual tapes are uniquely identified by a barcode, and you can add a prefix to the barcode. You can use a prefix to help identify your virtual tapes. The prefix must be uppercase letters (A–Z) and must be one to four characters long.

9. For **Pool**, choose **Glacier Pool**, **Deep Archive Pool**, or a custom pool that you have created. The pool determines the storage class in which your tape is stored when it is ejected by your backup software.
 - Choose **Glacier Pool** if you want to archive the tape in the S3 Glacier Flexible Retrieval storage class. When your backup software ejects the tape, it is automatically archived in S3 Glacier Flexible Retrieval. You use S3 Glacier Flexible Retrieval for more active archives, where you can retrieve a tape typically within 3–5 hours. For more information, see [Storage classes for archiving objects](#) in the *Amazon Simple Storage Service User Guide*.
 - Choose **Deep Archive Pool** if you want to archive the tape in the S3 Glacier Deep Archive storage class. When your backup software ejects the tape, the tape is automatically archived in S3 Glacier Deep Archive. You use S3 Glacier Deep Archive for long-term data retention and digital preservation, where data is accessed once or twice a year. You can retrieve a tape archived in S3 Glacier Deep Archive typically within 12 hours. For more information, see [Storage classes for archiving objects](#) in the *Amazon Simple Storage Service User Guide*.
 - Choose a custom pool, if any are available. You configure custom tape pools to use either **Deep Archive Pool** or **Glacier Pool**. Tapes are archived to the configured storage class when they are ejected by your backup software.

If you archive a tape in S3 Glacier Flexible Retrieval, you can move it to S3 Glacier Deep Archive later. For more information, see [Moving Your Tape from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive Storage Class](#).

 **Note**

Tapes created before March 27, 2019, are archived directly in S3 Glacier Flexible Retrieval when your backup software ejects them.

10. (Optional) For **Tags**, choose **Add new tag** and enter a key and value to add tags to your tape. A tag is a case-sensitive key-value pair that helps you manage, filter, and search for your tapes.
11. Choose **Create tapes**.
12. In the navigation pane, choose **Tape Library > Tapes** to see your tapes. By default, this list displays up to 1,000 tapes at a time, but the searches that you perform apply to all of your tapes. You can use the search bar to find tapes that match a specific criteria, or to reduce the list to less than 1,000 tapes. When your list contains 1,000 tapes or fewer, you can then sort your tapes in ascending or descending order by various properties.

The status of the virtual tapes is initially set to **CREATING** when the virtual tapes are being created. After the tapes are created, their status changes to **AVAILABLE**. For more information, see [Managing Your Tape Gateway](#).

Allowing Automatic Tape Creation

The Tape Gateway can automatically create new virtual tapes to maintain the minimum number of available tapes that you configure. It then makes these new tapes available for import by the backup application so that your backup jobs can run without interruption. Allowing automatic tape creation removes the need for custom scripting in addition to the manual process of creating new virtual tapes.

The Tape Gateway spawns a new tape automatically when it has fewer tapes than the minimum number of available tapes specified for automatic tape creation. A new tape is spawned when:

- A tape is imported from an import/export slot.
- A tape is imported to the tape drive.

The gateway maintains a minimum number of tapes with the barcode prefix specified in the automatic tape creation policy. If there are fewer tapes than the minimum number of tapes with the barcode prefix, the gateway automatically creates enough new tapes to equal the minimum number of tapes specified in the automatic tape creation policy.

When you eject a tape and it goes into the import/export slot, that tape does not count toward the minimum number of tapes specified in your automatic tape creation policy. Only tapes in the import/export slot are counted as being "available." Exporting a tape does not initiate automatic tape creation. Only imports affect the number of available tapes.

Moving a tape from the import/export slot to a tape drive or storage slot reduces the number of tapes in the import/export slot with the same barcode prefix. The gateway creates new tapes to maintain the minimum number of available tapes for that barcode prefix.

To allow automatic tape creation

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose the **Gateways** tab.
3. Choose the gateway that you want to automatically create tapes for.
4. In the **Actions** menu, choose **Configure tape auto-create**.

The **Tape auto-create** page appears. You can add, change, or remove tape auto-create options here.

5. To allow automatic tape creation, choose **Add new item** then configure the settings for automatic tape creation.
6. For **Tape type**, choose **Standard** to create standard virtual tapes. Choose **WORM** to create *write-once-read-many* (WORM) virtual tapes.
7. For **Minimum number of tapes**, enter the minimum number of virtual tapes that should be available on the Tape Gateway at all times. The valid range for this value is a minimum of 1 and a maximum of 10.
8. For **Capacity**, enter the size, in bytes, of the virtual tape capacity. The valid range is a minimum of 100 GiB and a maximum of 15 TiB.
9. For **Barcode prefix**, enter the prefix that you want to prepend to the barcode of your virtual tapes.

Note

Virtual tapes are uniquely identified by a barcode, and you can add a prefix to the barcode. The prefix is optional, but you can use it to help identify your virtual tapes. The prefix must be uppercase letters (A–Z) and must be one to four characters long.

10. For **Pool**, choose **Glacier Pool**, **Deep Archive Pool**, or a custom pool that you have created. The pool determines the storage class in which your tape is stored when it is ejected by your backup software.
 - Choose **Glacier Pool** if you want to archive the tape in the S3 Glacier Flexible Retrieval storage class. When your backup software ejects the tape, it is automatically archived in S3 Glacier Flexible Retrieval. You use S3 Glacier Flexible Retrieval for more active archives, where you can retrieve a tape typically within 3–5 hours. For more information, see [Storage classes for archiving objects](#) in the *Amazon Simple Storage Service User Guide*.
 - Choose **Deep Archive Pool** if you want to archive the tape in the S3 Glacier Deep Archive storage class. When your backup software ejects the tape, the tape is automatically archived in S3 Glacier Deep Archive. You use S3 Glacier Deep Archive for long-term data retention and digital preservation, where data is accessed once or twice a year. You can retrieve a tape archived in S3 Glacier Deep Archive typically within 12 hours. For more information, see [Storage classes for archiving objects](#) in the *Amazon Simple Storage Service User Guide*.
 - Choose a custom pool, if any are available. You configure custom tape pools to use either **Deep Archive Pool** or **Glacier Pool**. Tapes are archived to the configured storage class when they are ejected by your backup software.

If you archive a tape in S3 Glacier Flexible Retrieval, you can move it to S3 Glacier Deep Archive later. For more information, see [Moving Your Tape from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive Storage Class](#).

Note

Tapes created before March 27, 2019, are archived directly in S3 Glacier Flexible Retrieval when your backup software ejects them.

11. When finished configuring settings, choose **Save changes**.

12. In the navigation pane, choose **Tape Library > Tapes** to see your tapes. By default, this list displays up to 1,000 tapes at a time, but the searches that you perform apply to all of your tapes. You can use the search bar to find tapes that match a specific criteria, or to reduce the list to less than 1,000 tapes. When your list contains 1,000 tapes or fewer, you can then sort your tapes in ascending or descending order by various properties.

The status of available virtual tapes is initially set to **CREATING** when the tapes are being created. After the tapes are created, their status changes to **AVAILABLE**. For more information, see [Managing Your Tape Gateway](#).

For more information about changing automatic tape creation policies, or deleting automatic tape creation from a Tape Gateway, see [Managing Automatic Tape Creation](#).

Next Step

[Using Your Tape Gateway](#)

Using Your Tape Gateway

Following, you can find instructions about how to use your Tape Gateway.

Topics

- [Connecting Your VTL Devices](#)
- [Using Your Backup Software to Test Your Gateway Setup](#)
- [Where Do I Go from Here?](#)

Connecting Your VTL Devices

Following, you can find instructions about how to connect your virtual tape library (VTL) devices to your Microsoft Windows or Red Hat Enterprise Linux (RHEL) client.

Topics

- [Connecting to a Microsoft Windows Client](#)
- [Connecting to a 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.

To connect your VTL devices to a Windows client

1. Start `iscsicpl.exe`.

Note

You must have administrator rights on the client computer to run the iSCSI initiator.

2. Start the Microsoft iSCSI initiator service.
3. In the **iSCSI Initiator Properties** dialog box, choose the **Discovery** tab, and then choose **Discover Portal**.
4. Provide the IP address of your Tape Gateway for **IP address or DNS name**.
5. Choose the **Targets** tab, and then choose **Refresh**. All 10 tape drives and the medium changer appear in the **Discovered targets** box. The status for the targets is **Inactive**.
6. Choose the first device and connect it. You connect the devices one at a time.
7. Connect all of the targets.

On a Windows client, the driver provider for the tape drive must be Microsoft. Use the following procedure to verify the driver provider, and update the driver and provider if necessary:

To verify and update the driver and provider

1. On your Windows client, start Device Manager.
2. Expand **Tape drives**, open the context (right-click) menu for a tape drive, and choose **Properties**.
3. In the **Driver** tab of the **Device Properties** dialog box, verify **Driver Provider** is Microsoft.
4. If **Driver Provider** is not Microsoft, set the value as follows:
 - a. Choose **Update Driver**.
 - b. In the **Update Driver Software** dialog box, choose **Browse my computer for driver software**.

- c. In the **Update Driver Software** dialog box, choose **Let me pick from a list of device drivers on my computer**.
 - d. Choose **LTO Tape drive** and choose **Next**.
5. Choose **Close** to close the **Update Driver Software** window, and verify that the **Driver Provider** value is now set to Microsoft.
 6. Repeat the steps to update driver and provider for all the tape drives.

Connecting to a Linux Client

The following procedure shows a summary of the steps that you follow to connect to an RHEL client.

To connect a Linux client to VTL devices

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 looks 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.

Be 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
```

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

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

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

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

```
total 0
lrwxrwxrwx 1 root root 10 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-mediachanger-lun-0-changer -> ../../sg20
lrwxrwxrwx 1 root root 9 Sep 8 11:19 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-01-lun-0 -> ../../st6
lrwxrwxrwx 1 root root 10 Sep 8 11:19 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-01-lun-0-nst -> ../../nst6
lrwxrwxrwx 1 root root 9 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-02-lun-0 -> ../../st7
```



```
lrwxrwxrwx 1 root root 10 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-02-lun-0-nst -> ../../nst7
lrwxrwxrwx 1 root root 9 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-03-lun-0 -> ../../st8
lrwxrwxrwx 1 root root 10 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-03-lun-0-nst -> ../../nst8
lrwxrwxrwx 1 root root 9 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-04-lun-0 -> ../../st9
lrwxrwxrwx 1 root root 10 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-04-lun-0-nst -> ../../nst9
lrwxrwxrwx 1 root root 10 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-05-lun-0 -> ../../st10
lrwxrwxrwx 1 root root 11 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-05-lun-0-nst -> ../../nst10
lrwxrwxrwx 1 root root 10 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-06-lun-0 -> ../../st11
lrwxrwxrwx 1 root root 11 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-06-lun-0-nst -> ../../nst11
lrwxrwxrwx 1 root root 10 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-07-lun-0 -> ../../st12
lrwxrwxrwx 1 root root 11 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-07-lun-0-nst -> ../../nst12
lrwxrwxrwx 1 root root 10 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-08-lun-0 -> ../../st13
lrwxrwxrwx 1 root root 11 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-08-lun-0-nst -> ../../nst13
lrwxrwxrwx 1 root root 10 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-09-lun-0 -> ../../st14
lrwxrwxrwx 1 root root 11 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-09-lun-0-nst -> ../../nst14
lrwxrwxrwx 1 root root 10 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-10-lun-0 -> ../../st15
lrwxrwxrwx 1 root root 11 Sep 8 11:20 ip-10.6.56.90:3260-iscsi-
iqn.1997-05.com.amazon:sgw-9999999c-tapedrive-10-lun-0-nst -> ../../nst15
lrwxrwxrwx 1 root root 9 Aug 19 10:15 pci-0000:12:00.0-fc-0x0000000000000012-lun-0-
changer -> ../../sg6
lrwxrwxrwx 1 root root 9 Aug 19 10:15 pci-0000:12:00.0-fc-0x000000000000001c-lun-0
-> ../../st0
lrwxrwxrwx 1 root root 10 Aug 19 10:15 pci-0000:12:00.0-fc-0x000000000000001c-
lun-0-nst -> ../../nst0
lrwxrwxrwx 1 root root 9 Aug 19 10:15 pci-0000:12:00.0-fc-0x000000000000001f-lun-0
-> ../../st1
lrwxrwxrwx 1 root root 10 Aug 19 10:15 pci-0000:12:00.0-fc-0x000000000000001f-
lun-0-nst -> ../../nst1
```

```
lrwxrwxrwx 1 root root 9 Aug 19 10:15 pci-0000:12:00.0-fc-0x0000000000000022-lun-0
-> ../../st2
lrwxrwxrwx 1 root root 10 Aug 19 10:15 pci-0000:12:00.0-fc-0x0000000000000022-
lun-0-nst -> ../../nst2
lrwxrwxrwx 1 root root 9 Aug 19 10:15 pci-0000:12:00.1-fc-0x0000000000000025-lun-0
-> ../../st5
lrwxrwxrwx 1 root root 10 Aug 19 10:15 pci-0000:12:00.1-fc-0x0000000000000025-
lun-0-nst -> ../../nst5
lrwxrwxrwx 1 root root 9 Aug 19 10:15 pci-0000:12:00.1-fc-0x0000000000000028-lun-0
-> ../../st3
lrwxrwxrwx 1 root root 10 Aug 19 10:15 pci-0000:12:00.1-fc-0x0000000000000028-
lun-0-nst -> ../../nst3
lrwxrwxrwx 1 root root 9 Aug 19 10:15 pci-0000:12:00.1-fc-0x000000000000002b-lun-0
-> ../../st4
lrwxrwxrwx 1 root root 10 Aug 19 10:15 pci-0000:12:00.1-fc-0x000000000000002b-
lun-0-nst -> ../../nst4
```

Next Step

[Using Your Backup Software to Test Your Gateway Setup](#)

Using Your Backup Software to Test Your Gateway Setup

You test your Tape Gateway setup by performing the following tasks using your backup application:

1. Configure the backup application to detect your storage devices.

Note

To improve I/O performance, we recommend setting the block size of the tape drives in your backup application to 1 MB. For more information, see [Use a Larger Block Size for Tape Drives](#).

2. Back up data to a tape.
3. Archive the tape.
4. Retrieve the tape from the archive.
5. Restore data from the tape.

To test your setup, use a compatible backup application, as described following.

Note

Unless otherwise stated, all backup applications were qualified on Microsoft Windows.

Topics


- [Testing Your Setup by Using Arcserve Backup r17.0](#)
- [Testing Your Setup by Using Bacula Enterprise](#)
- [Testing Your Setup by Using Commvault](#)
- [Testing Your Setup by Using Dell EMC NetWorker](#)
- [Testing Your Setup by Using IBM Spectrum Protect](#)
- [Testing Your Setup by Using Micro Focus \(HPE\) Data Protector](#)
- [Testing Your Setup by Using Microsoft System Center Data Protection Manager](#)
- [Testing Your Setup by Using NovaStor DataCenter/Network](#)
- [Testing Your Setup by Using Quest NetVault Backup](#)
- [Testing Your Setup by Using Veeam Backup & Replication](#)
- [Testing Your Setup by Using Veritas Backup Exec](#)
- [Testing Your Setup by Using Veritas NetBackup](#)

For more information about compatible backup applications, see [Supported third-party backup applications for a Tape Gateway](#).

Testing Your Setup by Using Arcserve Backup r17.0

You can back up your data to virtual tapes, archive the tapes, and manage your virtual tape library (VTL) devices by using Arcserve Backup r17.0. In this topic, you can find basic documentation to configure Arcserve Backup with a Tape Gateway and perform a backup and restore operation. For detailed information about to use Arcserve Backup r17.0, see [Arcserve Backup r17 documentation](#) in the *Arcserve Administration Guide*.

The following screen shot shows the Arcserve menus.



Topics

- [Configuring Arcserve to Work with VTL Devices](#)
- [Loading Tapes into a Media Pool](#)
- [Backing Up Data to a Tape](#)
- [Archiving a Tape](#)
- [Restoring Data from a Tape](#)

Configuring Arcserve to Work with VTL Devices

After you have connected your virtual tape library (VTL) devices to your client, you scan for your devices.

To scan for VTL devices

1. In the Arcserve Backup Manager, choose the **Utilities** menu.
2. Choose **Media Assure and Scan**.

Loading Tapes into a Media Pool

When the Arcserve software connects to your gateway and your tapes become available, Arcserve automatically loads your tapes. If your gateway is not found in the Arcserve software, try restarting the tape engine in Arcserve.

To restart the tape engine

1. Choose **Quick Start**, choose **Administration**, and then choose **Device**.
2. On the navigation menu, open the context (right-click) menu for your gateway and choose an import/export slot.
3. Choose **Quick Import** and assign your tape to an empty slot.
4. Open the context (right-click) menu for your gateway and choose **Inventory/Offline Slots**.
5. Choose **Quick Inventory** to retrieve media information from the database.

If you add a new tape, you need to scan your gateway for the new tape to have it appear in Arcserve. If the new tapes don't appear, you must import the tapes.

To import tapes

1. Choose the **Quick Start** menu, choose **Back up**, and then choose **Destination tap**.
2. Choose your gateway, open the context (right-click) menu for one tape, and then choose **Import/Export Slot**.
3. Open the context (right-click) menu for each new tape and choose **Inventory**.
4. Open the context (right-click) menu for each new tape and choose **Format**.

Each tape's barcode now appears in your Storage Gateway console, and each tape is ready to use.

Backing Up Data to a Tape

When your tapes have been loaded into Arcserve, you can back up data. The backup process is the same as backing up physical tapes.

To back up data to a tape

1. From the **Quick Start** menu, open the restore a backup session.
2. Choose the **Source** tab, and then choose the file system or database system that you want to back up.

3. Choose the **Schedule** tab and choose the repeat method you want to use.
4. Choose the **Destination** tab and then choose the tape you want to use. If the data you are backing up is larger than the tape can hold, Arcserve prompts you to mount a new tape.
5. Choose **Submit** to back up your data.

Note

If your Tape Gateway restarts for any reason during an ongoing backup job, the backup job might fail. To complete the failed backup job, you must resubmit it.

Archiving a Tape

When you archive a tape, your Tape Gateway moves the tape from the tape library to the offline storage. Before you eject and archive a tape, you might want to check the content on it.

To archive a tape

1. From the **Quick Start** menu, open the restore a backup session.
2. Choose the **Source** tab, and then choose the file system or database system you want to back up.
3. Choose the **Schedule** tab and choose the repeat method you want to use.
4. Choose your gateway, open the context (right-click) menu for one tape, and then choose **Import/Export Slot**.
5. Assign a mail slot to load the tape. The status in the Storage Gateway console changes to **Archive**. The archive process might take some time.

The archiving process can take some time to complete. The initial status of the tape appears as **IN TRANSIT TO VTS**. When archiving starts, the status changes to **ARCHIVING**. When archiving is completed, the tape is no longer listed in the VTL but is archived in S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive.

Restoring Data from a Tape

Restoring your archived data is a two-step process.

To restore data from an archived tape

1. Retrieve the archived tape to a Tape Gateway. For instructions, see [Retrieving Archived Tapes](#).
2. Use Arcserve to restore the data. This process is the same as restoring data from physical tapes. For instructions, see the [Arcserve Backup r17 documentation](#).

To restore data from a tape, use the following procedure.

To restore data from a tape

1. From the **Quick Start** menu, open the restore a restore session.
2. Choose the **Source** tab, and then choose the file system or database system you want to restore.
3. Choose the **Destination** tab and accept the default settings.
4. Choose the **Schedule** tab, choose the repeat method that you want to use, and then choose **Submit**.

Next Step

[Cleaning Up Resources You Don't Need](#)

Testing Your Setup by Using Bacula Enterprise

You can back up your data to virtual tapes, archive the tapes, and manage your virtual tape library (VTL) devices by using Bacula Enterprise version 10. In this topic, you can find basic documentation on how to configure the Bacula version 10 backup application for a Tape Gateway and perform backup and restore operations. For detailed information about how to use Bacula version 10, see [Bacula Systems Manuals and Documentation](#) or contact Bacula Systems.

Note

Bacula is only supported on Linux.

Setting Up Bacula Enterprise

After you have connected your virtual tape library (VTL) devices to your Linux client, you configure the Bacula software to recognize your devices. For information about how to connect VTL devices to your client, see [Connecting Your VTL Devices](#).

To set up Bacula

1. Get a licensed copy of the Bacula Enterprise backup software from Bacula Systems.
2. Install the Bacula Enterprise software on your on-premises or in-cloud computer.

For information about how to get the installation software, see [Enterprise Backup for Amazon S3 and Storage Gateway](#). For additional installation guidance, see the Bacula whitepaper [Using Cloud Services and Object Storage with Bacula Enterprise Edition](#).

Configuring Bacula to Work with VTL Devices

Next, configure Bacula to work with your VTL devices. Following, you can find basic configuration steps.

To configure Bacula

1. Install the Bacula Director and the Bacula Storage daemon. For instructions, see chapter 7 of the [Using Cloud Services and Object Storage with Bacula Enterprise Edition](#) Bacula white paper.
2. Connect to the system that is running Bacula Director and configure the iSCSI initiator. To do so, use the script provided in step 7.4 in the [Using Cloud Services and Object Storage with Bacula Enterprise Edition](#) Bacula whitepaper.
3. Configure the storage devices. Use the script provided in the Bacula whitepaper discussed preceding.
4. Configure the local Bacula Director, add storage targets, and define media pools for your tapes. Use the script provided in the Bacula whitepaper discussed preceding.

Backing Up Data to Tape

1. Create tapes in the Storage Gateway console. For information on how to create tapes, see [Creating Tapes](#).
2. Transfer tapes from the I/E slot to the storage slot by using the following command.


```
/opt/bacula/scripts/mtx-changer
```

For example, the following command transfers tapes from I/E slot 1601 to storage slot 1.

```
/opt/bacula/scripts/mtx-changer transfer 1601 1
```

3. Launch the Bacula console by using the following command.

```
/opt/bacula/bin/bconsole
```

Note

When you create and transfer a tape to Bacula, use the Bacula console (bconsole) command `update slots storage=VTL` so that Bacula knows about the new tapes that you created.

4. Label the tape with the barcode as the volume name or label by using the following bconsole command.

```
label storage=VTL pool=pool.VTL barcodes === label the tapes with the  
barcode as the volume name / label
```

5. Mount the tape by using the following command.

```
mount storage=VTL slot=1 drive=0
```

6. Create a backup job that uses the media pools you created, and then write data to the virtual tape by using the same procedures that you do with physical tapes.
7. Unmount the tape from the Bacula console by using the following command.

```
umount storage=VTL slot=1 drive=0
```

Note

If your Tape Gateway restarts for any reason during an ongoing backup job, the backup job will fail, and the tape status in Bacula Enterprise will change to **FULL**. If you know the tape has not been fully utilized, you can manually change the tape status back to **APPEND** and continue the backup job using the same tape. You can also continue the job on a different tape if other tapes in **APPEND** status are available.

Archiving a Tape

When all backup jobs for a particular tape are done and you can archive the tape, use the `mtx-changer` script to move the tape from the storage slot to the I/E slot. This action is similar to the `eject` action in other backup applications.

To archive a tape

1. Transfer the tape from the storage slot to the I/E slot by using the `/opt/bacula/scripts/mtx-changer` command.

For example, the following command transfers a tape from the storage slot 1 to I/E slot 1601.

```
/opt/bacula/scripts/mtx-changer transfer 1 1601
```

2. Verify that the tape is archived in the offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive) and that the tape has the status **Archived**.

Restoring Data from an Archived and Retrieved Tape

Restoring your archived data is a two-step process.

To restore data from an archived tape

1. Retrieve the archived tape from archive to a Tape Gateway. For instructions, see [Retrieving Archived Tapes](#).
2. Restore your data by using the Bacula software:
 - a. Import the tapes into the storage slot by using the `/opt/bacula/scripts/mtx-changer` command to transfer tapes from the I/E slot.

For example, the following command transfers tapes from I/E slot 1601 to storage slot 1.

```
/opt/bacula/scripts/mtx-changer transfer 1601 1
```

- b. Use the Bacula console to update the slots, and then mount the tape.
- c. Run the restore command to restore your data. For instructions, see the Bacula documentation.

Testing Your Setup by Using Commvault

You can back up your data to virtual tapes, archive the tapes, and manage your virtual tape library (VTL) devices by using Commvault version 11. In this topic, you can find basic documentation on how to configure the Commvault backup application for a Tape Gateway, perform a backup archive, and retrieve your data from archived tapes. For detailed information about how to use Commvault, see the [Commvault Quick Start Guide](#) on the Commvault website.

Topics

- [Configuring Commvault to Work with VTL Devices](#)
- [Creating a Storage Policy and a Subclient](#)
- [Backing Up Data to a Tape in Commvault](#)
- [Archiving a Tape in Commvault](#)
- [Restoring Data from a Tape](#)


Configuring Commvault to Work with VTL Devices

After you connect the VTL devices to the Windows client, you configure Commvault to recognize them. For information about how to connect VTL devices to the Windows client, see [Connecting Your VTL Devices to a Windows client](#).


The Commvault backup application doesn't automatically recognize VTL devices. You must manually add devices to expose them to the Commvault backup application and then discover the devices.

To configure Commvault

1. In the CommCell console main menu, choose **Storage**, and then choose **Expert Storage Configuration** to open the **Select MediaAgents** dialog box.



2. Choose the available media agent you want to use, choose **Add**, and then choose **OK**.
3. In the **Expert Storage Configuration** dialog box, choose **Start**, and then choose **Detect/Configure Devices**.
4. Leave the **Device Type** options selected, choose **Exhaustive Detection**, and then choose **OK**.
5. In the **Confirm Exhaustive Detection** confirmation box, choose **Yes**.
6. In the **Device Selection** dialog box, choose your library and all its drives, and then choose **OK**. Wait for your devices to be detected, and then choose **Close** to close the log report.
7. Right-click your library, choose **Configure**, and then choose **Yes**. Close the configuration dialog box.
8. In the **Does this library have a barcode reader?** dialog box, choose **Yes**, and then for device type, choose **IBM ULTRIUM V5**.
9. In the CommCell browser, choose **Storage Resources**, and then choose **Libraries** to see your tape library.



10. To see your tapes in your library, open the context (right-click) menu for your library, and then choose **Discover Media, Media location, Media Library**.
11. To mount your tapes, open the context (right-click) menu for your media, and then choose **Load**.

Creating a Storage Policy and a Subclient

Every backup and restore job is associated with a storage policy and a subclient policy.

A storage policy maps the original location of the data to your media.

To create a storage policy

1. In the CommCell browser, choose **Policies**.
2. Open the context (right-click) menu for **Storage Policies**, and then choose **New Storage Policy**.
3. In the Create Storage Policy wizard, choose **Data Protection and Archiving**, and then choose **Next**.

4. Type a name for **Storage Policy Name**, and then choose **Incremental Storage Policy**. To associate this storage policy with incremental loads, choose one of the options. Otherwise, leave the options unchecked, and then choose **Next**.
5. In the **Do you want to Use Global Deduplication Policy?** dialog box, choose your **Deduplication** preference, and then choose **Next**.
6. From **Library for Primary Copy**, choose your VTL library, and then choose **Next**.
7. Verify that your media agent settings are correct, and then choose **Next**.
8. Verify that your scratch pool settings are correct, and then choose **Next**.
9. Configure your retention policies in **iData Agent Backup data**, and then choose **Next**.
10. Review the encryption settings, and then choose **Next**.
11. To see your storage policy, choose **Storage Policies**.

You create a subclient policy and associate it with your storage policy. A subclient policy allows you to configure similar file system clients from a central template, so that you don't have to set up many similar file systems manually.

To create a subclient policy

1. In the CommCell browser, choose **Client Computers**, and then choose your client computer. Choose **File System**, and then choose **defaultBackupSet**.
2. Right-click **defaultBackupSet**, choose **All Tasks**, and then choose **New Subclient**.
3. In the **Subclient** properties box, type a name in **SubClient Name**, and then choose **OK**.
4. Choose **Browse**, navigate to the files that you want to back up, choose **Add**, and then close the dialog box.
5. In the **Subclient** property box, choose the **Storage Device** tab, choose a storage policy from **Storage policy**, and then choose **OK**.
6. In the **Backup Schedule** window that appears, associate the new subclient with a backup schedule.
7. Choose **Do Not Schedule** for one time or on-demand backups, and then choose **OK**.

You should now see your subclient in the **defaultBackupSet** tab.

Backing Up Data to a Tape in Commvault

You create a backup job and write data to a virtual tape by using the same procedures you use with physical tapes. For more information, see the [Commvault documentation](#).

Note

If your Tape Gateway restarts for any reason during an ongoing backup job, the backup job might fail. In some cases, you can select an option to resume the failed job. Otherwise, you must submit a new job. If Commvault marks the tape as unusable after a job fails, you must reload the tape into the drive to continue writing to it. If multiple tapes are available, Commvault might continue the failed backup job on a different tape.

Archiving a Tape in Commvault

You start the archiving process by ejecting the tape. When you archive a tape, Tape Gateway moves the tape from the tape library to offline storage. Before you eject and archive a tape, you might want to first check the content on the tape.

To archive a tape

1. In the CommCell browser, choose **Storage Resources, Libraries**, and then choose **Your library**. Choose **Media By Location**, and then choose **Media In Library**.
2. Open the context (right-click) menu for the tape you want to archive, choose **All Tasks**, choose **Export**, and then choose **OK**.

The archiving process can take some time to complete. The initial status of the tape appears as **IN TRANSIT TO VTS**. When archiving starts, the status changes to **ARCHIVING**. When archiving is completed, the tape is no longer listed in the VTL.

In the Commvault software, verify that the tape is no longer in the storage slot.

In the navigation pane of the Storage Gateway console, choose **Tapes**. Verify that your archived tape's status is **ARCHIVED**.

Restoring Data from a Tape

You can restore data from a tape that has never been archived and retrieved, or from a tape that has been archived and retrieved. For tapes that have never been archived and retrieved (nonretrieved tapes), you have two options to restore the data:

- Restore by subclient
- Restore by job ID

To restore data from a nonretrieved tape by subclient

1. In the CommCell browser, choose **Client Computers**, and then choose your client computer. Choose **File System**, and then choose **defaultBackupSet**.
2. Open the context (right-click) menu for your subclient, choose **Browse and Restore**, and then choose **View Content**.
3. Choose the files you want to restore, and then choose **Recover All Selected**.
4. Choose **Home**, and then choose **Job Controller** to monitor the status of your restore job.

To restore data from a nonretrieved tape by job ID

1. In the CommCell browser, choose **Client Computers**, and then choose your client computer. Right-click **File System**, choose **View**, and then choose **Backup History**.
2. In the **Backup Type** category, choose the type of backup jobs you want, and then choose **OK**. A tab with the history of backup jobs appears.
3. Find the **Job ID** you want to restore, right-click it, and then choose **Browse and Restore**.
4. In the **Browse and Restore Options** dialog box, choose **View Content**.
5. Choose the files that you want to restore, and then choose **Recover All Selected**.
6. Choose **Home**, and then choose **Job Controller** to monitor the status of your restore job.

To restore data from an archived and retrieved tape

1. In the CommCell browser, choose **Storage Resources**, choose **Libraries**, and then choose **Your library**. Choose **Media By Location**, and then choose **Media In Library**.
2. Right-click the retrieved tape, choose **All Tasks**, and then choose **Catalog**.
3. In the **Catalog Media** dialog box, choose **Catalog only**, and then choose **OK**.

4. Choose **CommCell Home**, and then choose **Job Controller** to monitor the status of your restore job.
5. After the job succeeds, open the context (right-click) menu for your tape, choose **View**, and then choose **View Catalog Contents**. Take note of the **Job ID** value for use later.
6. Choose **Recatalog/Merge**. Make sure that **Merge only** is chosen in the **Catalog Media** dialog box.
7. Choose **Home**, and then choose **Job Controller** to monitor the status of your restore job.
8. After the job succeeds, choose **CommCell Home**, choose **Control Panel**, and then choose **Browse/Search/Recovery**.
9. Choose **Show aged data during browse and recovery**, choose **OK**, and then close the **Control Panel**.
10. In the CommCell browser, right-click **Client Computers**, and then choose your client computer. Choose **View**, and then choose **Job History**.
11. In the **Job History Filter** dialog box, choose **Advanced**.
12. Choose **Include Aged Data**, and then choose **OK**.
13. In the **Job History** dialog box, choose **OK** to open the **history of jobs** tab.
14. Find the job that you want to restore, open the context (right-click) menu for it, and then choose **Browse and Restore**.
15. In the **Browse and Restore** dialog box, choose **View Content**.
16. Choose the files that you want to restore, and then choose **Recover All Selected**.
17. Choose **Home**, and then choose **Job Controller** to monitor the status of your restore job.

Testing Your Setup by Using Dell EMC NetWorker

You can back up your data to virtual tapes, archive the tapes and manage your virtual tape library (VTL) devices by using Dell EMC NetWorker 19.5. In this topic, you can find basic documentation on how to configure the Dell EMC NetWorker software to work with a Tape Gateway and perform a backup, including how to configure storage devices, write data to a tape, archive a tape and restore data from a tape.

For detailed information about how to install and use the Dell EMC NetWorker software, see the [Administration Guide](#).

For more information about compatible backup applications, see [Supported third-party backup applications for a Tape Gateway](#).

Topics


- [Configuring to Work with VTL Devices](#)
- [Allowing Import of WORM Tapes into Dell EMC NetWorker](#)
- [Backing Up Data to a Tape in Dell EMC NetWorker](#)
- [Archiving a Tape in Dell EMC NetWorker](#)
- [Restoring Data from an Archived Tape in Dell EMC NetWorker](#)


Configuring to Work with VTL Devices

After you have connected your virtual tape library (VTL) devices to your Microsoft Windows client, you configure to recognize your devices. For information about how to connect VTL devices to the Windows client, see [Connecting Your VTL Devices](#).

doesn't automatically recognize Tape Gateway devices. To expose your VTL devices to the NetWorker software and get the software to discover them, you manually configure the software. Following, we assume that you have correctly installed the software and that you are familiar with the Management Console. For more information about the Management Console, see the NetWorker Management Console interface section of the [Dell EMC NetWorker Administration Guide](#).

The following screen shot shows Dell EMC NetWorker 19.5.

 NetWorker Management Console V19.5.0.3 - localhost



To configure the Dell EMC NetWorker software for VTL devices

1. Start the Dell EMC NetWorker Management Console application, choose **Enterprise** from the menu, and then choose **localhost** from the left pane.

2. Open the context (right-click) menu for **localhost**, and then choose **Launch Application**.
3. Choose the **Devices** tab, open the context (right-click) menu for **Libraries**, and then choose **Scan for Devices**.
4. In the Scan for Devices wizard, choose **Start Scan**, and then choose **OK** from the dialog box that appears.
5. Expand the **Libraries** folder tree to see all your libraries and hit F5 to refresh. This process might take a few seconds to load the devices into the library.
6. Open a command window (CMD.exe) with admin privileges and run "jbconfig" utility that is installed with Dell EMC NetWorker 19.5.

```
Microsoft Windows [Version 10.0.17763.2366]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>jbconfig

Jbconfig is running on host dellnetworker (Windows Server 2019 Datacenter 10.0),
and is using dellnetworker as the NetWorker server.

    1) Configure an Autodetected SCSI Jukebox.
    2) Configure an Autodetected NDMP SCSI Jukebox.
    3) Configure an SJI Jukebox.
    4) Configure an STL Silo.
    5) Exit.

Which activity do you want to perform? [1]
14484:jbconfig: Scanning SCSI buses; this may take a while ...
Installing 'Standard SCSI Jukebox' jukebox - scsidev@1.0.0.

What name do you want to assign to this jukebox device? AWSVTL
15814:jbconfig: Attempting to detect serial numbers on the jukebox and drives ...

15815:jbconfig: Will try to use SCSI information returned by jukebox to configure drives.

Turn NetWorker auto-cleaning on (yes / no) [yes]? no

The following drive(s) can be auto-configured in this jukebox:
 1> LTO Ultrium-5 @ 1.1.0 ==> \\.\Tape0
 2> LTO Ultrium-5 @ 1.2.0 ==> \\.\Tape1
 3> LTO Ultrium-5 @ 1.3.0 ==> \\.\Tape2
 4> LTO Ultrium-5 @ 1.4.0 ==> \\.\Tape3
 5> LTO Ultrium-5 @ 1.5.0 ==> \\.\Tape4
 6> LTO Ultrium-5 @ 1.6.0 ==> \\.\Tape5
 7> LTO Ultrium-5 @ 1.7.0 ==> \\.\Tape6
 8> LTO Ultrium-5 @ 1.8.0 ==> \\.\Tape7
 9> LTO Ultrium-5 @ 1.9.0 ==> \\.\Tape8
10> LTO Ultrium-5 @ 1.10.0 ==> \\.\Tape9
These are all the drives that this jukebox has reported.
```

```

To change the drive model(s) or configure them as shared or NDMP drives,
you need to bypass auto-configure. Bypass auto-configure? (yes / no) [no]

Jukebox has been added successfully

The following configuration options have been set:

> Jukebox description to the control port and model.
> Autochanger control port to the port at which we found it.
> Autocleaning off.
> Barcode reading to on.
> Volume labels that match the barcodes.

You can review and change the characteristics of the autochanger and its
associated devices using the NetWorker Management Console.

Would you like to configure another jukebox? (yes/no) [no]

C:\Users\Administrator>_

```

7. When "jbconfig" completes, return to the NetWorker GUI and hit F5 to refresh.
8. Choose your library to see your tapes in the left pane and the corresponding empty volume slots list in the right pane. In this screen shot, the "AWSVTL" library is selected.

Device	Volume	Write...	Message	Slot	Volume	Barc
\\.\Tape0				1		
\\.\Tape1				2		
\\.\Tape2				3		
\\.\Tape3				4		
\\.\Tape4				5		
\\.\Tape5				6		
\\.\Tape6				7		
\\.\Tape7				8		
\\.\Tape8				9		
\\.\Tape9				10		
				11		

9. In the volume list, select the volumes you want to activate (selected volumes are highlighted), open the context (right-click) menu for the selected volumes, and then choose **Deposit**. This action moves the tape from the I/E slot into the volume slot.

10. In the dialog box that appears, choose **Yes**, and then in the **Load the Cartridges into** dialog box, choose **Yes**.
11. If you don't have any more tapes to deposit, choose **No** or **Ignore**. Otherwise, choose **Yes** to deposit additional tapes.

Allowing Import of WORM Tapes into Dell EMC NetWorker

You are now ready to import tapes from your Tape Gateway into the Dell EMC NetWorker library.

The virtual tapes are write once read many (WORM) tapes, but Dell EMC NetWorker expects non-WORM tapes. For Dell EMC NetWorker to work with your virtual tapes, you must activate import of tapes into non-WORM media pools.

To allow import of WORM tapes into non-WORM media pools

1. On NetWorker Console, choose **Media**, open the context (right-click) menu for **localhost**, and then choose **Properties**.
2. In the **NetWorker Sever Properties** window, choose the **Configuration** tab.
3. In the **Worm tape handling** section, clear the **WORM tapes only in WORM pools** box, and then choose **OK**.

Backing Up Data to a Tape in Dell EMC NetWorker

Backing up data to a tape is a two-step process.

1. Label the tapes you want to back up your data to, create the target media pool, and add the tapes to the pool.

You create a media pool and write data to a virtual tape by using the same procedures you do with physical tapes. For detailed information, see the Backing Up Data section of the [Dell EMC NetWorker Administration Guide](#).

2. Write data to the tape. You back up data by using the Dell EMC NetWorker User application instead of the Dell EMC NetWorker Management Console. The Dell EMC NetWorker User application installs as part of the NetWorker installation.

Note

You use the Dell EMC NetWorker User application to perform backups, but you view the status of your backup and restore jobs in the EMC Management Console. To view status, choose the **Devices** menu and view the status in the **Log** window.

Note

If your Tape Gateway restarts for any reason during an ongoing backup job, the backup job will be suspended, and the tape status in Dell EMC NetWorker will change to **Write Protected**. You can archive the tape or continue to read data from it. You can resume the suspended backup job on a different tape.

Archiving a Tape in Dell EMC NetWorker

When you archive a tape, Tape Gateway moves the tape from the Dell EMC NetWorker tape library to the offline storage. You begin tape archival by ejecting a tape from the tape drive to the storage slot. You then withdraw the tape from the slot to the archive by using your backup application—that is, the Dell EMC NetWorker software.

To archive a tape by using Dell EMC NetWorker

1. On the **Devices** tab in the NetWorker Administration window, choose **localhost** or your EMC server, and then choose **Libraries**.
2. Choose the library you imported from your virtual tape library.
3. From the list of tapes that you have written data to, open the context (right-click) menu for the tape you want to archive, and then choose **Eject/Withdraw**.
4. In the confirmation box that appears, choose **OK**.

The archiving process can take some time to complete. The initial status of the tape appears as **IN TRANSIT TO VTS**. When archiving starts, the status changes to **ARCHIVING**. When archiving is completed, the tape is no longer listed in the VTL.

In the Dell EMC NetWorker software, verify that the tape is no longer in the storage slot.

In the navigation pane of the Storage Gateway console, choose **Tapes**. Verify that your archived tape's status is **ARCHIVED**.

Restoring Data from an Archived Tape in Dell EMC NetWorker

Restoring your archived data is a two-step process:

1. Retrieve the archived tape a Tape Gateway. For instructions, see [Retrieving Archived Tapes](#).
2. Use the Dell EMC NetWorker software to restore the data. You do this by creating a restoring a folder file, as you do when restoring data from physical tapes. For instructions, see the Using the NetWorker User program section of the [Dell EMC NetWorker Administration Guide](#).

Next Step

[Cleaning Up Resources You Don't Need](#)

Testing Your Setup by Using IBM Spectrum Protect

You can back up your data to virtual tapes, archive the tapes, and manage your virtual tape library (VTL) devices by using IBM Spectrum Protect with AWS Storage Gateway. (IBM Spectrum Protect was formerly known as Tivoli Storage Manager.)

This topic contains basic information about how to configure the IBM Spectrum Protect version 8.1.10 backup software for a Tape Gateway. It also includes basic information about performing backup and restore operations with IBM Spectrum Protect. For more information about how to administer IBM Spectrum Protect backup software, see IBM's [Overview of administration tasks](#) for IBM Spectrum Protect.

The IBM Spectrum Protect backup software supports AWS Storage Gateway on the following operating systems.

- **Microsoft Windows Server**
- **Red Hat Linux**

For information about IBM Spectrum Protect supported devices for Windows, see [IBM Spectrum Protect \(formerly Tivoli Storage Manager\) Supported Devices for AIX, HP-UX, Solaris, and Windows](#).

For information about IBM Spectrum Protect supported devices for Linux, see [IBM Spectrum Protect \(formerly Tivoli Storage Manager\) Supported Devices for Linux](#).

Topics

- [Setting Up IBM Spectrum Protect](#)
- [Configuring IBM Spectrum Protect to Work with VTL Devices](#)
- [Writing Data to a Tape in IBM Spectrum Protect](#)
- [Restoring Data from a Tape Archived in IBM Spectrum Protect](#)

Setting Up IBM Spectrum Protect

After you connect your VTL devices to your client, you configure the IBM Spectrum Protect version 8.1.10 software to recognize them. For more information about connecting VTL devices to your client, see [Connecting Your VTL Devices](#).

To set up IBM Spectrum Protect

1. Get a licensed copy of the IBM Spectrum Protect version 8.1.10 software from IBM.
2. Install the IBM Spectrum Protect software on your on-premises environment or in-cloud Amazon EC2 instance. For more information, see IBM's [Installing and upgrading](#) documentation for IBM Spectrum Protect.

For more information about configuring IBM Spectrum Protect software, see [Configuring AWS Tape Gateway virtual tape libraries for an IBM Spectrum Protect server](#).

Configuring IBM Spectrum Protect to Work with VTL Devices

Next, configure IBM Spectrum Protect to work with your VTL devices. You can configure IBM Spectrum Protect to work with VTL devices on Microsoft Windows Server or Red Hat Linux.

Configuring IBM Spectrum Protect for Windows

For complete instructions on how to configure IBM Spectrum Protect on Windows, see [Tape Device Driver-W12 6266 for Windows 2012](#) on the Lenovo website. Following is basic documentation on the process.

To configure IBM Spectrum Protect for Microsoft Windows

1. Get the correct driver package for your media changer. For the tape-device driver, IBM Spectrum Protect requires version W12 6266 for Windows 2012. For instructions on how to get the drivers, see [Tape Device Driver-W12 6266 for Windows 2012](#) on the Lenovo website.

Note

Make sure that you install the "non-exclusive" set of drivers.

2. On your computer, open **Computer Management**, expand **Media Changer devices**, and verify that the media changer type is listed as **IBM 3584 Tape Library**.
3. Ensure that the barcode for any tape in the virtual tape library is eight characters or less. If you try to assign your tape a barcode that is longer than eight characters, you get this error message: "Tape barcode is too long for media changer".
4. Ensure that all your tape drives and media changer appear in IBM Spectrum Protect. To do so, use the following command: `\Tivoli\TSM\server>tsmdlst.exe`

Configure IBM Spectrum Protect for Linux

Following is basic documentation on configuring IBM Spectrum to work with VTL devices on Linux.

To configure IBM Spectrum Protect for Linux

1. Go to [IBM Fix Central](#) on the IBM Support website, and choose **Select product**.
2. For **Product Group**, choose **System Storage**.
3. For **Select from System Storage**, choose **Tape systems**.
4. For **Tape systems**, choose **Tape drivers and software**.
5. For **Select from Tape drivers and software**, choose **Tape device drivers**.
6. For **Platform**, choose your operating system and choose **Continue**.
7. Choose the device driver version that you want to download. Then follow the instructions on the **Fix Central** download page to download and configure IBM Spectrum Protect.
8. Ensure that the barcode for any tape in the virtual tape library is eight characters or less. If you try to assign your tape a barcode that is longer than eight characters, you get this error message: "Tape barcode is too long for media changer".

Writing Data to a Tape in IBM Spectrum Protect

You write data to a Tape Gateway virtual tape by using the same procedure and backup policies that you do with physical tapes. Create the necessary configuration for backup and restore jobs. For

more information about configuring IBM Spectrum Protect, see [Overview of administration tasks](#) for IBM Spectrum Protect.

Note

If your Tape Gateway restarts for any reason during an ongoing backup job, the backup job might fail. If the backup job fails, the tape status in IBM Spectrum Protect changes to **ReadOnly**. If you know the tape has not been fully utilized, you can manually change the tape status back to **ReadWrite**, and either resume or resubmit the backup job using the same tape. IBM Spectrum Protect might continue the failed backup job on a different tape if other tapes in **ReadWrite** status are available.

Restoring Data from a Tape Archived in IBM Spectrum Protect

Restoring your archived data is a two-step process.

To restore data from an archived tape

1. Retrieve the archived tape from archive to a Tape Gateway. For instructions, see [Retrieving Archived Tapes](#).
2. Restore the data by using the IBM Spectrum Protect backup software. You do this by creating a recovery point, as you do when restoring data from physical tapes. For more information about configuring IBM Spectrum Protect, see [Overview of administration tasks](#) for IBM Spectrum Protect.

Next Step

[Cleaning Up Resources You Don't Need](#)

Testing Your Setup by Using Micro Focus (HPE) Data Protector

You can back up your data to virtual tapes, archive the tapes, and manage your virtual tape library (VTL) devices by using Micro Focus (HPE) Data Protector v9.x. In this topic, you can find basic documentation on how to configure the Micro Focus (HPE) Data Protector software for a Tape Gateway and perform a backup and restore operation. For detailed information about how to use the Micro Focus (HPE) Data Protector software, see the Hewlett Packard documentation. For more information about compatible backup applications, see [Supported third-party backup applications for a Tape Gateway](#).

Topics

- [Configuring Micro Focus \(HPE\) Data Protector to Work with VTL Devices](#)
- [Preparing Virtual Tapes for Use with HPE Data Protector](#)
- [Loading Tapes into a Media Pool](#)
- [Backing Up Data to a Tape](#)
- [Archiving a Tape](#)
- [Restoring Data from a Tape](#)

Configuring Micro Focus (HPE) Data Protector to Work with VTL Devices


After you have connected the virtual tape library (VTL) devices to the client, you configure Micro Focus (HPE) Data Protector to recognize your devices. For information about how to connect VTL devices to the client, see [Connecting Your VTL Devices](#).

The Micro Focus (HPE) Data Protector software doesn't automatically recognize Tape Gateway devices. To have the software recognize these devices, manually add the devices and then discover the VTL devices, as described following.

To add the VTL devices

1. In the Micro Focus (HPE) Data Protector main window, choose the **Devices & Media** shelf in the list at top left.

Open the context (right-click) menu for **Devices**, and choose **Add Device**.



2. On the **Add Device** tab, type a value for **Device Name**. For **Device Type**, choose **SCSI Library**, and then choose **Next**.
3. On the next screen, do the following:
 - a. For **SCSI address of the library robotic**, select your specific address.
 - b. For **Select what action Data Protector should take if the drive is busy**, choose "Abort" or your preferred action.
 - c. Choose to activate these options:
 - **Barcode reader support**
 - **Automatically discover changed SCSI address**
 - **SCSI Reserve/Release (robotic control)**
 - d. Leave **Use barcode as medium label on initialization** clear (unchecked), unless your system requires it.
 - e. Choose **Next** to continue.

4. On the next screen, specify the slots that you want to use with HP Data Protector. Use a hyphen ("-") between numbers to indicate a range of slots, for example 1–6. When you've specified slots to use, choose **Next**.
5. For the standard type of media used by the physical device, choose **LTO_Ultrium**, and then choose **Finish** to complete the setup.

Your tape library is now ready to use. To load tapes into it, see the next section.

Preparing Virtual Tapes for Use with HPE Data Protector

Before you can back up data to a virtual tape, you need to prepare the tape for use. Doing this involves the following actions:

- Load a virtual tape into a tape library
- Load the virtual tape into a slot
- Create a media pool
- Load the virtual tape into media pool

In the following sections, you can find steps to guide you through this process.

Loading Virtual Tapes into a Tape Library

Your tape library should now be listed under **Devices**. If you don't see it, press F5 to refresh the screen. When your library is listed, you can load virtual tapes into the library.

To load virtual tapes into your tape library

1. Choose the plus sign next to your tape library to display the nodes for robotics paths, drives, and slots.
2. Open the context (right-click) menu for **Drives**, choose **Add Drive**, type a name for your tape, and then choose **Next** to continue.
3. Choose the tape drive you want to add for **SCSI address of data drive**, choose **Automatically discover changed SCSI address**, and then choose **Next**.
4. On the following screen, choose **Advanced**. The **Advanced Options** pop-up screen appears.
 - a. On the **Settings** tab, you should consider the following options:
 - **CRC Check** (to detect accidental data changes)

- **Detect dirty drive** (to ensure the drive is clean before backup)
- **SCSI Reserve/Release(drive)** (to avoid tape contention)

For testing purposes, you can leave these options deactivated (unchecked).

- b. On the **Sizes** tab, set the **Block size (kB)** to **Default (256)**.
 - c. Choose **OK** to close the advanced options screen, and then choose **Next** to continue.
5. On the next screen, choose these options under **Device Policies**:
 - **Device may be used for restore**
 - **Device may be used as source device for object copy**
 6. Choose **Finish** to finish adding your tape drive to your tape library.

Loading Virtual Tapes into Slots

Now that you have a tape drive in your tape library, you can load virtual tapes into slots.

To load a tape into a slot

1. In the tape library tree node, open the node labeled **Slots**. Each slot has a status represented by an icon:
 - A green tape means that a tape is already loaded into the slot.
 - A gray slot means that the slot is empty.
 - A cyan question mark means that the tape in that slot is not formatted.
2. For an empty slot, open the context (right-click) menu, and then choose **Enter**. If you have existing tapes, choose one to load into that slot.

Creating a Media Pool

A *media pool* is a logical group used to organize your tapes. To set up tape backup, you create a media pool.

To create a media pool

1. In the **Devices & Media** shelf, open the tree node for **Media**, open the context (right-click) menu for the **Pools** node, and then choose **Add Media Pool**.

2. For **Pool name**, type a name.
3. For **Media Type**, choose **LTO_Ultrium**, and then choose **Next**.
4. On the following screen, accept the default values, and then choose **Next**.
5. Choose **Finish** to finish creating a media pool.

Loading Tapes into a Media Pool

Before you can back up data onto your tapes, you must load the tapes into the media pool that you created.

To load a virtual tape into a media pool

1. On your tape library tree node, choose the **Slots** node.
2. Choose a loaded tape, one that has a green icon showing a loaded tape. Open the context (right-click) menu and choose **Format**, and then choose **Next**.
3. Choose the media pool you created, and then choose **Next**.
4. For **Medium Description**, choose **Use barcode**, and then choose **Next**.
5. For **Options**, choose **Force Operation**, and then choose **Finish**.

You should now see your chosen slot change from a status of unassigned (gray) to a status of tape inserted (green). A series of messages appear to confirm that your media is initialized.

At this point, you should have everything configured to begin using your virtual tape library with HPE Data Protector. To double-check that this is the case, use the following procedure.

To verify that your tape library is configured for use

- Choose **Drives**, then open the context (right-click) menu for your drive, and choose **Scan**.

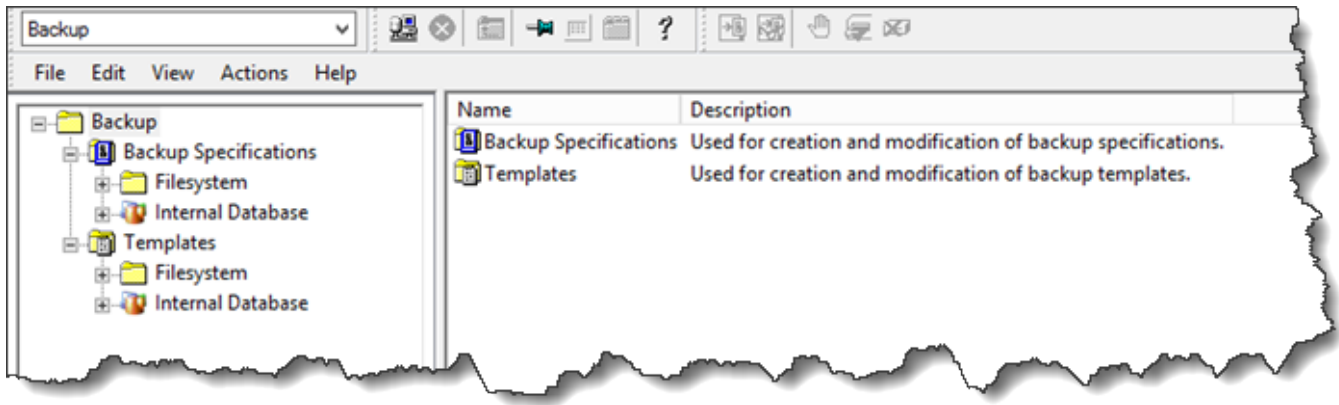
If your configuration is correct, a message confirms that your media was successfully scanned.

Backing Up Data to a Tape

When your tapes have been loaded into a media pool, you can back up data to them.

To back up data to a tape

1. Choose the **Backup** shelf at top left of the screen.



- Open the context (right-click) menu for **Filesystem**, and choose **Add Backup**.
- On the **Create New Backup** screen, under **Filesystem**, choose **Blank File System Backup**, and then choose **OK**.
- On the tree node that shows your host system, select the file system or file systems that you want to back up, and choose **Next** to continue.
- Open the tree node for the tape library you want to use, open the context (right-click) menu for the tape drive you want to use, and then choose **Properties**.
- Choose your media pool, choose **OK**, and then choose **Next**.
- For the next three screens, accept the default settings and choose **Next**.
- On the **Perform finishing steps in your backup/template design** screen, choose **Save as** to save this session. In the pop-up window, give the backup a name and assign it to the group where you want to save your new backup specification.
- Choose **Start Interactive Backup**.

If the host system contains a database system, you can choose it as your target backup system. The screens and selections are similar to the file-system backup just described.

Note

If your Tape Gateway restarts for any reason during an ongoing backup job, the backup job will fail, and the tape drive in Data Protector is marked as **Dirty**. Data Protector also marks the tape quality as **Poor**, and prevents writing to the tape. To continue reading data from the tape, you must clean the drive and re-mount the tape. To complete the failed backup job, you must resubmit it on a new tape.

Archiving a Tape

When you archive a tape, Tape Gateway moves the tape from the tape library to the offline storage. Before you eject and archive a tape, you might want to check the content on it.

To check a tape's content before archiving it

1. Choose **Slots** and then choose the tape you want to check.
2. Choose **Objects** and check what content is on the tape.

When you have chosen a tape to archive, use the following procedure.

To eject and archive a tape

1. Open the context (right-click) menu for that tape, and choose **Eject**.
2. On the Storage Gateway console, choose your gateway, and then choose **VTL Tape Cartridges** and verify the status of the virtual tape you are archiving.

After the tape is ejected, it will be automatically archived in the offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive). The archiving process can take some time to complete. The initial status of the tape is shown as **IN TRANSIT TO VTS**. When archiving starts, the status changes to **ARCHIVING**. When archiving is completed, the tape is no longer listed in the VTL but is archived in S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive.

Restoring Data from a Tape

Restoring your archived data is a two-step process.


To restore data from an archived tape

1. Retrieve the archived tape to a Tape Gateway. For instructions, see [Retrieving Archived Tapes](#).
2. Use HPE Data Protector to restore the data. This process is the same as restoring data from physical tapes.

To restore data from a tape, use the following procedure.

To restore data from a tape

1. Choose the **Restore** shelf at the top left of the screen.



2. Choose the file system or database system you want to restore. For the backup that you want to restore, make sure that the box is selected. Choose **Restore**.
3. In the **Start Restore Session** window, choose **Needed Media**. Choose **All media**, and you should see the tape originally used for the backup. Choose that tape, and then choose **Close**.
4. In the **Start Restore Session** window, accept the default settings, choose **Next**, and then choose **Finish**.

Next Step

[Cleaning Up Resources You Don't Need](#)

Testing Your Setup by Using Microsoft System Center Data Protection Manager

You can back up your data to virtual tapes, archive the tapes, and manage your virtual tape library (VTL) devices by using Microsoft System Center 2012 R2 or 2016 Data Protection Manager (DPM). In this topic, you can find basic documentation on how to configure the DPM backup application for a Tape Gateway and perform a backup and restore operation.

For detailed information about how to use DPM, see the [DPM documentation](#) on the Microsoft System Center website. For more information about compatible backup applications, see [Supported third-party backup applications for a Tape Gateway](#).

Topics

- [Configuring DPM to Recognize VTL Devices](#)
- [Importing a Tape into DPM](#)
- [Writing Data to a Tape in DPM](#)
- [Archiving a Tape by Using DPM](#)
- [Restoring Data from a Tape Archived in DPM](#)

Configuring DPM to Recognize VTL Devices

After you have connected the virtual tape library (VTL) devices to the Windows client, you configure DPM to recognize your devices. For information about how to connect VTL devices to the Windows client, see [Connecting Your VTL Devices](#).

By default, the DPM server does not recognize Tape Gateway devices. To configure the server to work with the Tape Gateway devices, you perform the following tasks:

1. Update the device drivers for the VTL devices to expose them to the DPM server.
2. Manually map the VTL devices to the DPM tape library.

To update the VTL device drivers

- In Device Manager, update the driver for the medium changer. For instructions, see [Updating the Device Driver for Your Medium Changer](#).


You use the DPMDriveMappingTool to map your tape drives to the DPM tape library.

To map tape drives to the DPM server tape library

1. Create at least one tape for your gateway. For information on how to do this on the console, see [Creating Tapes](#).
2. Import the tape into the DPM library. For information on how to do this, see [Importing a Tape into DPM](#).
3. If the DPMLA service is running, stop it by opening a command terminal and typing the following on the command line.

```
net stop DPMLA
```

4. Locate the following file on the DPM server: %ProgramFiles%\System Center 2016 R2\DPM\DPM\Config\DPMLA.xml.

 **Note**

If this file exists, the DPMDriveMappingTool overwrites it. If you want to preserve your original file, create a backup copy.

5. Open a command terminal, change the directory to %ProgramFiles%\System Center 2016 R2\DPM\DPM\Bin, and run the following command.

```
C:\Microsoft System Center 2016 R2\DPM\DPM\bin>DPMDriveMappingTool.exe
```

The output for the command looks like the following.


```
Performing Device Inventory ...
Mapping Drives to Library ...
Adding Standalone Drives ...
Writing the Map File ...
Drive Mapping Completed Successfully.
```

Importing a Tape into DPM

You are now ready to import tapes from your Tape Gateway into the DPM backup application library.

To import tapes into the DPM backup application library

1. On the DPM server, open the Management Console, choose **Rescan**, and then choose **Refresh**. Doing this displays your medium changer and tape drives.



- Open the context (right-click) menu for the media changer in the **Library** section, and then choose **Add tape (I/E port)** to add a tape to the **Slots** list.

Note

The process of adding tapes can take several minutes to complete.

The tape label appears as **Unknown**, and the tape is not usable. For the tape to be usable, you must identify it.

- Open the context (right-click) menu for the tape you want to identify, and then choose **Identify unknown tape**.

Note

The process of identifying tapes can take a few seconds or a few minutes.

If the tapes don't display barcodes correctly, you need to change the media changer driver to Sun/StorageTek Library. For more information, see [Displaying Barcodes for Tapes in Microsoft System Center DPM](#).

When identification is complete, the tape label changes to **Free**. That is, the tape is free for data to be written to it.

In the following screenshot, the tape in slot 2 has been identified and is free to use but the tape in slot 3 is not.

Name /	Status	Tape Label	Barcode
Library: Spectra Logic Medium Changer			
Drives (In use: 0 Idle: 2)			
LTO Tape drive (AMZN_SGW-...	Idle	-	None
LTO Tape drive (AMZN_SGW-...	Idle	-	None
Slots (Tape available: 7 Empty: 1593)			
Slot 1	Empty	-	None
Slot 2	Tape available	Free	AMZN9FA53A
Slot 3	Tape available	Unknown	PH27A582
Slot 4	Tape available	Free	AMZN9FA537

Writing Data to a Tape in DPM

You write data to a Tape Gateway virtual tape by using the same protection procedures and policies you do with physical tapes. You create a protection group and add the data you want to back up, and then back up the data by creating a recovery point. For detailed information about how to use DPM, see the [DPM documentation](#) on the Microsoft System Center website.

By default, the capacity of a tape is 30GB. When you backup data that is larger than a tape's capacity, a device I/O error occurs. If the position where the error occurred is larger than the size of the tape, Microsoft DPM treats the error as an indication of end of tape. If the position where the error occurred is less than the size of the tape, the backup job fails. To resolve the issue, change the TapeSize value in the registry entry to match the size of your tape. For information about how to do this, see [Error ID: 30101](#) at the Microsoft System Center.

Note

If your Tape Gateway restarts for any reason during an ongoing backup job, the backup job will fail. To complete the failed backup job, you must resubmit it.

Archiving a Tape by Using DPM

When you archive a tape, Tape Gateway moves the tape from the DPM tape library to offline storage. You begin tape archival by removing the tape from the slot using your backup application—that is, DPM.

To archive a tape in DPM

1. Open the context (right-click) menu for the tape you want to archive, and then choose **Remove tape (I/E port)**.



2. In the dialog box that appears, choose **Yes**. Doing this ejects the tape from the medium changer's storage slot and moves the tape into one of the gateway's I/E slots. When a tape is moved into the gateway's I/E slot, it is immediately sent for archiving.
3. On the Storage Gateway console, choose your gateway, and then choose **VTL Tape Cartridges** and verify the status of the virtual tape you are archiving.

The archiving process can take some time to complete. The initial status of the tape is shown as **IN TRANSIT TO VTS**. When archiving starts, the status changes to **ARCHIVING**. When archiving is completed, the tape is no longer listed in the VTL.

Restoring Data from a Tape Archived in DPM

Restoring your archived data is a two-step process.

To restore data from an archived tape

1. Retrieve the archived tape from archive to a Tape Gateway. For instructions, see [Retrieving Archived Tapes](#).
2. Use the DPM backup application to restore the data. You do this by creating a recovery point, as you do when restoring data from physical tapes. For instructions, see [Recovering Client Computer Data](#) on the DPM website.

Next Step

[Cleaning Up Resources You Don't Need](#)

Testing Your Setup by Using NovaStor DataCenter/Network

You can back up your data to virtual tapes, archive the tapes, and manage your virtual tape library (VTL) devices by using NovaStor DataCenter/Network version 6.4 or 7.1. In this topic, you can find basic documentation on how to configure the NovaStor DataCenter/Network version 7.1 backup application for a Tape Gateway and perform backup and restore operations. For detailed information about how to use NovaStor DataCenter/Network version 7.1, see [Documentation NovaStor DataCenter/Network](#).

Setting Up NovaStor DataCenter/Network

After you have connected your virtual tape library (VTL) devices to your Microsoft Windows client, you configure the NovaStor software to recognize your devices. For information about how to connect VTL devices to your Windows client, see [Connecting Your VTL Devices](#).

NovaStor DataCenter/Network requires drivers from the driver manufacturers. You can use the Windows drivers, but you must first deactivate other backup applications.


Configuring NovaStor DataCenter/Network to Work with VTL Devices

When configuring your VTL devices to work with NovaStor DataCenter/Network version 6.4 or 7.1, you might see an error message that reads `External Program did not exit correctly`. This issue requires a workaround, which you need to perform before you continue.

You can prevent the issue by creating the workaround before you start configuring your VTL devices. For information about how to create the workaround, see [Resolving an "External Program Did Not Exit Correctly" Error](#).

To configure NovaStor DataCenter/Network to work with VTL devices

1. In the NovaStor DataCenter/Network Admin console, choose **Media Management**, and then choose **Storage Management**.



2. In the **Storage Targets** menu, open the context menu (right-click) for **Media Management Servers**, choose **New**, and choose **OK** to create and prepopulate a **storage** node.

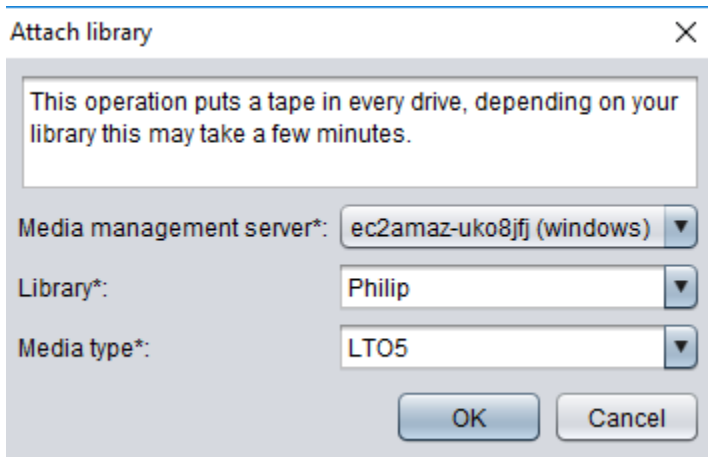
If you see an error message that says External Program did not exit correctly, resolve the issue before you continue. This issue requires a workaround. For information about how to resolve this issue, see [Resolving an "External Program Did Not Exit Correctly" Error](#).

⚠ Important

This error occurs because the element assignment range from AWS Storage Gateway for storage drives and tape drives exceeds the number that NovaStor DataCenter/Network allows.

3. Open the context (right-click) menu for the **storage** node that was created, and choose **New Library**.
4. Choose the library server from the list. The library list is automatically populated.

5. Name the library and choose **OK**.
6. Choose the library to display all the properties of the Storage Gateway virtual tape library.
7. In the **Storage Targets** menu, expand **Backup Servers**, open the context (right-click) menu for the server, and choose **Attach Library**.
8. In the **Attach Library** dialog box that appears, choose the **LTO5** media type, and then choose **OK**.



9. Expand **Backup Servers** to see the Storage Gateway virtual tape library and the library partition that shows all the mounted tape drives.

Creating a Tape Pool

A tape pool is dynamically created in the NovaStor DataCenter/Network software and so doesn't contain a fixed number of media. A tape pool that needs a tape gets it from its scratch pool. A *scratch pool* is a reservoir of tapes that are freely available for one or more tape pools to use. A tape pool returns to the scratch pool any media that have exceeded their retention times and that are no longer needed.

Creating a tape pool is a three-step task:

1. You create a scratch pool.
2. You assign tapes to the scratch pool.
3. You create a tape pool.

To create a scratch pool

1. In the left navigation menu, choose the **Scratch Pools** tab.

2. Open the context (right-click) menu for **Scratch Pools**, and choose **Create Scratch Pool**.
3. In the **Scratch Pools** dialog box, name your scratch pool, and then choose your media type.
4. Choose **Label Volume**, and create a low water mark for the scratch pool. When the scratch pool is emptied down to the low water mark, a warning appears.
5. In the warning dialog box that appears, choose **OK** to create the scratch pool.

To assign tapes to a scratch pool

1. In the left navigation menu, choose **Tape Library Management**.
2. Choose the **Library** tab to see your library's inventory.
3. Choose the tapes that you want to assign to the scratch pool. Make sure that the tapes are set to the correct media type.
4. Open the context (right-click) menu for the library and choose **Add to Scratch Pool**.

You now have a filled scratch pool that you can use for tape pools.

To create a tape pool

1. From the left navigation menu, choose **Tape Library Management**.
2. Open the context (right-click) menu for the **Media Pools** tab and choose **Create Media Pool**.
3. Name the media pool and choose **Backup Server**.
4. Choose a library partition for the media pool.
5. Choose the scratch pool that you want the pool to get the tapes from.
6. For **Schedule**, choose **Not Scheduled**.

Configuring Media Import and Export to Archive Tapes

NovaStor DataCenter/Network can use import/export slots if they are part of the media changer.

For an export, NovaStor DataCenter/Network must know which tapes are going to be physically taken out of the library.

For an import, NovaStor DataCenter/Network recognizes tape media that are exported in the tape library and offers to import them all, either from a data slot or an export slot. Your Tape Gateway archives tapes in the offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive).

To configure media import and export

1. Navigate to **Tape Library Management**, choose a server for **Media Management Server**, and then choose **Library**.
2. Choose the **Off-site Locations** tab.
3. Open the context (right-click) menu for the white area, and choose **Add** to open a new panel.
4. In the panel, type **S3 Glacier Flexible Retrieval** or **S3 Glacier Deep Archive** and add an optional description in the text box.

Backing Up Data to Tape

You create a backup job and write data to a virtual tape by using the same procedures that you do with physical tapes. For detailed information about how to back up data using the NovaStor software, see [Documentation NovaStor DataCenter/Network](#).

Note

If your Tape Gateway restarts for any reason during an ongoing backup job, the backup job will fail, and the tape will become unwriteable. You can archive the tape or continue to read data from it. To complete the failed backup job, you must resubmit it on a new tape.

Archiving a Tape

When you archive a tape, a Tape Gateway ejects the tape from the tape drive to the storage slot. It then exports the tape from the slot to the archive by using your backup application—that is, NovaStor DataCenter/Network.

To archive a tape

1. In the left navigation menu, choose **Tape Library Management**.
2. Choose the **Library** tab to see the library's inventory.
3. Highlight the tapes you want to archive, open the context (right-click) menu for the tapes, and choose your off-site archive location.

The archiving process can take some time to complete. The initial status of the tape appears as **IN TRANSIT TO VTS**. When archiving starts, the status changes to **ARCHIVING**. When archiving is completed, the tape is no longer listed in the VTL.

In NovaStor DataCenter/Network, verify that the tape is no longer in the storage slot.

In the navigation pane of the Storage Gateway console, choose **Tapes**. Verify that your archived tape's status is **ARCHIVED**.

Restoring Data from an Archived and Retrieved Tape

Restoring your archived data is a two-step process.

To restore data from an archived tape

1. Retrieve the archived tape from archive to a Tape Gateway. For instructions, see [Retrieving Archived Tapes](#).
2. Use the NovaStor DataCenter/Network software to restore the data. You do this by refreshing the mail slot and moving each tape you want to retrieve into an empty slot, as you do when restoring data from physical tapes. For information about restoring data, see [Documentation NovaStor DataCenter/Network](#).

Writing Several Backup Jobs to a Tape Drive at the Same Time

In the NovaStor software, you can write several jobs to a tape drive at the same time using the multiplexing feature. This feature is available when a multiplexer is available for a media pool. For information about how to use multiplexing, see [Documentation NovaStor DataCenter/Network](#).

Resolving an "External Program Did Not Exit Correctly" Error

When configuring your VTL devices to work with NovaStor DataCenter/Network version 6.4 or 7.1, you might see an error message that reads `External Program did not exit correctly`. This error occurs because the element assignment range from Storage Gateway for storage drives and tape drives exceeds the number that NovaStor DataCenter/Network allows.

Storage Gateway returns 3200 storage and import/export slots, which is more than the 2400 limit that NovaStor DataCenter/Network allows. To resolve this issue, you add a configuration file that activates the NovaStor software to limit the number of storage and import/export slots and preconfigures the element assignment range.


```
1 Configuration
2 Status Handler
3 Status Import/Export
4 Status Drive
5 Status Slot
6 Mount Medium
7 Unmount Medium
8 Find Address by Tag

9 Reset Stacker
11 Move Element
88 Inventory
99 Exit

What ([#,#[,#]])? 11
Source Address? 30000
Destination Address? 20000

1 Configuration
2 Status Handler
3 Status Import/Export
4 Status Drive
5 Status Slot
6 Mount Medium
7 Unmount Medium
8 Find Address by Tag
9 Reset Stacker
```

6. Attach the library to the backup server.
7. In the NovaStor software, import all the tapes from import/export slots into the library.

Testing Your Setup by Using Quest NetVault Backup

You can back up your data to virtual tapes, archive the tapes, and manage your virtual tape library (VTL) devices by using the following Quest (formerly Dell) NetVault Backup versions:

- Quest NetVault Backup 12.4
- Quest NetVault Backup 13.x

In this topic, you can find basic documentation on how to configure the Quest NetVault Backup application for a Tape Gateway and perform a backup and restore operation.

For detailed information about how to use the Quest NetVault Backup application, see the [Quest NetVault Backup – Administration Guide](#). For more information about compatible backup applications, see [Supported third-party backup applications for a Tape Gateway](#).

Topics

- [Configuring Quest NetVault Backup to Work with VTL Devices](#)
- [Backing Up Data to a Tape in the Quest NetVault Backup](#)
- [Archiving a Tape by Using the Quest NetVault Backup](#)
- [Restoring Data from a Tape Archived in Quest NetVault Backup](#)

Configuring Quest NetVault Backup to Work with VTL Devices


After you have connected the virtual tape library (VTL) devices to the Windows client, you configure Quest NetVault Backup to recognize your devices. For information about how to connect VTL devices to the Windows client, see [Connecting Your VTL Devices](#).

The Quest NetVault Backup application doesn't automatically recognize Tape Gateway devices. You must manually add the devices to expose them to the Quest NetVault Backup application and then discover the VTL devices.

Adding VTL Devices

To add the VTL devices

1. In Quest NetVault Backup, choose **Manage Devices** in the **Configuration** tab.
2. On the Manage Devices page, choose **Add Devices**.
3. In the Add Storage Wizard, choose **Tape library / media changer**, and then choose **Next**.



4. On the next page, choose the client machine that is physically attached to the library and choose **Next** to scan for devices.
5. If devices are found, they are displayed. In this case, your medium changer is displayed in the device box.
6. Choose your medium changer and choose **Next**. Detailed information about the device is displayed in the wizard.
7. On the Add Tapes to Bays page, choose **Scan For Devices**, choose your client machine, and then choose **Next**.

All your drives are displayed on the page. Quest NetVault Backup displays the 10 bays to which you can add your drives. The bays are displayed one at a time.

Device	Serial Number
3-0.5.0 (IBM ULT3580-TD5)	AMZN_SGW- 54A94C3D_TD_00005
3-0.29.0 (IBM ULT3580-TD5)	AMZN_SGW- 54A94C3D_TD_00007
3-0.30.0 (IBM ULT3580-TD5)	AMZN_SGW- 54A94C3D_TD_00008
3-0.31.0 (IBM ULT3580-TD5)	AMZN_SGW- 54A94C3D_TD_00009
3-0.32.0 (IBM ULT3580-TD5)	AMZN_SGW- 54A94C3D_TD_00010

1 - 5 of 5 items

8. Choose the drive you want to add to the bay that is displayed, and then choose **Next**.

⚠ Important

When you add a drive to a bay, the drive and bay numbers must match. For example, if bay 1 is displayed, you must add drive 1. If a drive is not connected, leave its matching bay empty.

9. When your client machine appears, choose it, and then choose **Next**. The client machine can appear multiple times.
10. When the drives are displayed, repeat steps 7 through 9 to add all the drives to the bays.
11. In the **Configuration** tab, choose **Manage devices** and on the **Manage Devices** page, expand your medium changer to see the devices that you added.

Backing Up Data to a Tape in the Quest NetVault Backup

You create a backup job and write data to a virtual tape by using the same procedures you do with physical tapes. For detailed information about how to back up data, see the [Quest NetVault Backup - Administration Guide](#).

Note


If your Tape Gateway restarts for any reason during an ongoing backup job, the backup job will fail. To complete the failed backup job, you must resubmit it.

Archiving a Tape by Using the Quest NetVault Backup

When you archive a tape, a Tape Gateway ejects the tape from the tape drive to the storage slot. It then exports the tape from the slot to the archive by using your backup application—that is, the Quest NetVault Backup.

To archive a tape in Quest NetVault Backup

1. In the Quest NetVault Backup Configuration tab, choose and expand your medium changer to see your tapes.
2. On the **Slots** row, choose the settings icon to open the **Slots Browser** for the medium changer.



3. In the slots, locate the tape you want to archive, choose it, and then choose **Export**.

Slot ▲	Status	Barcode	Media
1	Reserved		
2	Has Blank Media	AMZND1A774	
3	Has Blank Media	AMZND6A773	
4	Empty		
5	Empty		
6	Empty		
7	Empty		
8	Empty		
9	Empty		
10	Empty		

Navigation icons: Home, Previous, Next, End

Buttons: < Back, > Ports, Set Slot, Export, 🔍 Scan

The archiving process can take some time to complete. The initial status of the tape appears as **IN TRANSIT TO VTS**. When archiving starts, the status changes to **ARCHIVING**. When archiving is completed, the tape is no longer listed in the VTL.

In the Quest NetVault Backup software, verify that the tape is no longer in the storage slot.

In the navigation pane of the Storage Gateway console, choose **Tapes**. Verify that your archived tape's status is **ARCHIVED**.

Restoring Data from a Tape Archived in Quest NetVault Backup

Restoring your archived data is a two-step process.

To restore data from an archived tape

1. Retrieve the archived tape from archive to a Tape Gateway. For instructions, see [Retrieving Archived Tapes](#).

2. Use the Quest NetVault Backup application to restore the data. You do this by creating a restoring a folder file, as you do when restoring data from physical tapes. For instructions on creating a restore job, see [Quest NetVault Backup - Administration Guide](#).

Next Step

[Cleaning Up Resources You Don't Need](#)

Testing Your Setup by Using Veeam Backup & Replication

You can back up your data to virtual tapes, archive the tapes, and manage your virtual tape library (VTL) devices by using Veeam Backup & Replication 11A. In this topic, you can find basic documentation on how to configure the Veeam Backup & Replication software for a Tape Gateway and perform a backup and restore operation. For detailed information about how to use the Veeam software, see the [About Veeam Backup & Replication](#) in the Veeam Help Center. For more information about compatible backup applications, see [Supported third-party backup applications for a Tape Gateway](#).

Topics

- [Configuring Veeam to Work with VTL Devices](#)
- [Importing a Tape into Veeam](#)
- [Backing Up Data to a Tape in Veeam](#)
- [Archiving a Tape by Using Veeam](#)
- [Restoring Data from a Tape Archived in Veeam](#)

Configuring Veeam to Work with VTL Devices

After you have connected your virtual tape library (VTL) devices to the Windows client, you configure Veeam Backup & Replication to recognize your devices. For information about how to connect VTL devices to the Windows client, see [Connecting Your VTL Devices](#).

Updating VTL Device Drivers


To configure the software to work with Tape Gateway devices, you update the device drivers for the VTL devices to expose them to the Veeam software and then discover the VTL devices. In Device Manager, update the driver for the medium changer. For instructions, see [Updating the Device Driver for Your Medium Changer](#).

Discovering VTL Devices

You must use native SCSI commands instead of a Windows driver to discover your tape library if your media changer is unknown. For detailed instructions, see [Tape Libraries](#).

To discover VTL devices

1. In the Veeam software, choose **Tape Infrastructure**. When the Tape Gateway is connected, virtual tapes are listed in the **Tape Infrastructure** tab.



2. Expand the **Tape** tree to see your tape drives and medium changer.
3. Expand the medium changer tree. If your tape drives are mapped to the medium changer, the drives appear under **Drives**. Otherwise, your tape library and tape drives appear as separate devices.

If the drives are not mapped automatically, follow the [instructions on the Veeam website](#) to map the drives.

Importing a Tape into Veeam

You are now ready to import tapes from your Tape Gateway into the Veeam backup application library.

To import a tape into the Veeam library

1. Open the context (right-click) menu for the medium changer, and choose **Import** to import the tapes to the I/E slots.
2. Open the context (right-click) menu for the medium charger, and choose **Inventory Library** to identify unrecognized tapes. When you load a new virtual tape into a tape drive for the first time, the tape is not recognized by the Veeam backup application. To identify the unrecognized tape, you inventory the tapes in the tape library.

Backing Up Data to a Tape in Veeam

Backing data to a tape is a two-step process:

1. You create a media pool and add the tape to the media pool.
2. You write data to the tape.

You create a media pool and write data to a virtual tape by using the same procedures you do with physical tapes. For detailed information about how to back up data, see the [Getting Started with Tapes](#) in the Veeam Help Center.

Note

If your Tape Gateway restarts for any reason during an ongoing backup job, the backup job will fail. To complete the failed backup job, you must resubmit it.

Archiving a Tape by Using Veeam

When you archive a tape, Tape Gateway moves the tape from the Veeam tape library to the offline storage. You begin tape archival by ejecting from the tape drive to the storage slot and then exporting the tape from the slot to the archive by using your backup application—that is, the Veeam software.

To archive a tape in the Veeam library

1. Choose **Tape Infrastructure**, and choose the media pool that contains the tape you want to archive.

The screenshot shows the AWS Tape Gateway console interface. At the top, there is a navigation bar with 'Home' and 'Tape Media' tabs. Below this is a ribbon of action buttons: 'Inventory', 'Catalog', 'Restore Content', 'Verify', 'Copy', 'Move to', 'Export', 'Eject', 'Erase', 'Mark as Free', 'Remove', and 'Protect'. The main area is divided into two panes. The left pane, titled 'Tape Infrastructure', shows a tree view with 'Tape Infrastructure' expanded to 'Media Pools' > 'Media Pool 1 (4)' > 'Tape Servers' > 'Libraries' > 'AWS Gateway-VTL 0100' > 'Drives' > 'Media' > 'Online'. The right pane shows a table of tapes:

Name ↑	Location	Expires
Tape 2	Slot 2	Not def
Tape 3	Slot 3	Not def
Tape 4	Slot 4	Not def

A context menu is open over 'Tape 4', listing actions: 'Inventory', 'Catalog', 'Restore content...', 'Copy...', 'Verify', 'Export' (highlighted), 'Eject', 'Move to', 'Erase', 'Mark as free', 'Protect', and 'Properties'.

2. Open the context (right-click) menu for the tape that you want to archive, and then choose **Eject Tape**.
3. For **Ejecting tape**, choose **Close**. The location of the tape changes from a tape drive to a slot.

4. Open the context (right-click) menu for the tape again, and then choose **Export**. The status of the tape changes from **Tape drive** to **Offline**.
5. For **Exporting tape**, choose **Close**. The location of the tape changes from **Slot** to **Offline**.
6. On the Storage Gateway console, choose your gateway, and then choose **VTL Tape Cartridges** and verify the status of the virtual tape you are archiving.

The archiving process can take some time to complete. The initial status of the tape appears as **IN TRANSIT TO VTS**. When archiving starts, the status changes to **ARCHIVING**. When archiving is completed, the tape is no longer listed in the VTL but is archived in S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive.

Restoring Data from a Tape Archived in Veeam

Restoring your archived data is a two-step process.

To restore data from an archived tape

1. Retrieve the archived tape from archive to a Tape Gateway. For instructions, see [Retrieving Archived Tapes](#).
2. Use the Veeam software to restore the data. You do this by creating a restoring a folder file, as you do when restoring data from physical tapes. For instructions, see [Restoring Files from Tape](#) in the Veeam Help Center.

Next Step

[Cleaning Up Resources You Don't Need](#)

Testing Your Setup by Using Veritas Backup Exec

You can back up your data to virtual tapes, archive the tapes, and manage your virtual tape library (VTL) devices by using Veritas Backup Exec. In this topic, you can find basic documentation needed to perform backup and restore operations using the following versions of Backup Exec:

- Veritas Backup Exec 2014
- Veritas Backup Exec 15
- Veritas Backup Exec 16
- Veritas Backup Exec 20.x

- Veritas Backup Exec 22.x

The procedure for using these versions of Backup Exec with a Tape Gateway is the same. See the [Veritas support website](#) for detailed information about how to use Backup Exec, including how to create secure backups with Backup Exec, software and hardware compatibility lists, and administrator guides for Backup Exec.

For more information about supported backup applications, see [Supported third-party backup applications for a Tape Gateway](#).

Topics

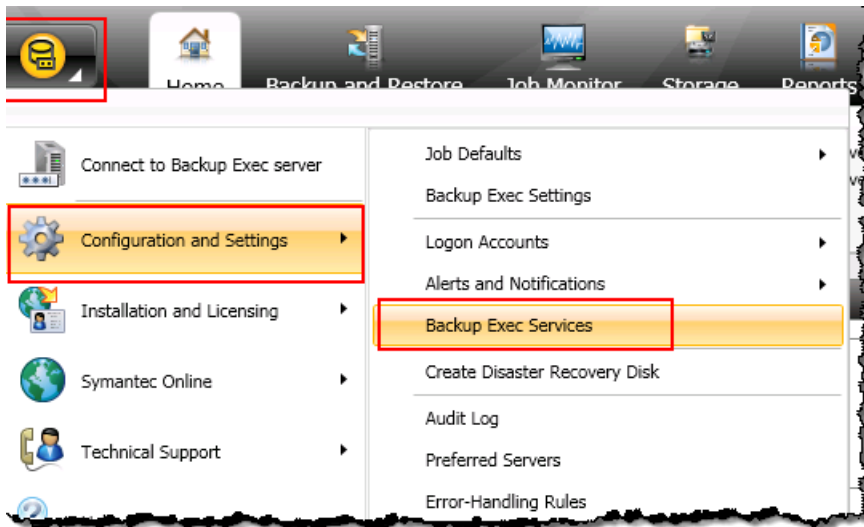
- [Configuring Storage in Backup Exec](#)
- [Importing a Tape in Backup Exec](#)
- [Writing Data to a Tape in Backup Exec](#)
- [Archiving a Tape Using Backup Exec](#)
- [Restoring Data from a Tape Archived in Backup Exec](#)
- [Deactivating a Tape Drive in Backup Exec](#)

Configuring Storage in Backup Exec

After you have connected the virtual tape library (VTL) devices to the Windows client, you configure Backup Exec storage to recognize your devices. For information about how to connect VTL devices to the Windows client, see [Connecting Your VTL Devices](#).

To configure storage


1. Start the Backup Exec software, and then choose the yellow icon in top-left corner on the toolbar.
2. Choose **Configuration and Settings**, and then choose **Backup Exec Services** to open the Backup Exec Service Manager.



3. Choose **Restart All Services**. Backup Exec then recognizes the VTL devices (that is, the medium changer and tape drives). The restart process might take a few minutes.

Note

Tape Gateway provides 10 tape drives. However, your Backup Exec license agreement might require your backup application to work with fewer than 10 tape drives. In that case, you must deactivate tape drives in the Backup Exec robotic library to leave only the number of tape drives allowed by your license agreement activated. For instructions, see [Deactivating a Tape Drive in Backup Exec](#).



4. After the restart is completed, close the Backup Exec Service Manager.

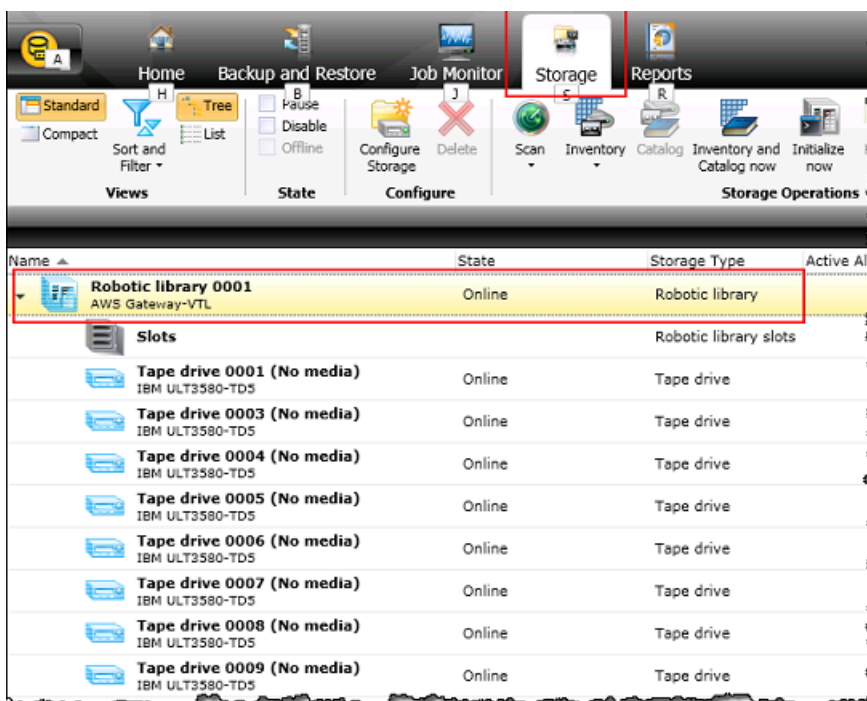
Importing a Tape in Backup Exec

You are now ready to import a tape from your gateway into a slot.

1. Choose the **Storage** tab, and then expand the **Robotic library** tree to display the VTL devices.

Important

Veritas Backup Exec software requires the Tape Gateway medium changer type. If the medium changer type listed under **Robotic library** is not Tape Gateway, you must change it before you configure storage in the backup application. For information about how to select a different medium changer type, see [Selecting a Medium Changer After Gateway Activation](#).




2. Choose the **Slots** icon to display all slots.

Note


When you import tapes into the robotic library, the tapes are stored in slots instead of tape drives. Therefore, the tape drives might have a message that indicates there is no media in the drives (No media). When you initiate a backup or restore job, the tapes are moved into the tape drives.

You must have tapes available in your gateway tape library to import a tape into a storage slot. For instructions on how to create tapes, see [Adding Virtual Tapes](#).


- Open the context (right-click) menu for an empty slot, choose **Import**, and then choose **Import media now**. In the following screen shot, slot number **3** is empty. You can select more than one slot and import multiple tapes in a single import operation.



- In the **Media Request** window that appears, choose **View details**.



5. In the **Action Alert: Media Intervention** window, choose **Respond OK** to insert the media into the slot.



The tape appears in the slot you selected.

Note

Tapes that are imported include empty tapes and tapes that have been retrieved from the archive to the gateway.

Writing Data to a Tape in Backup Exec

You write data to a Tape Gateway virtual tape by using the same procedure and backup policies you do with physical tapes. For detailed information, see the *Backup Exec Administrative Guide* in the documentation section in the Backup Exec software.

Note


If your Tape Gateway restarts for any reason during an ongoing backup job, the backup job might fail. If the backup job fails, the tape status in Veritas Backup Exec changes to **Not Applicable**. You can archive the tape or continue to read data from it. To complete the failed backup job, you must resubmit it on a new tape.

Archiving a Tape Using Backup Exec

When you archive a tape, Tape Gateway moves the tape from your gateway's virtual tape library (VTL) to the offline storage. You begin tape archival by exporting the tape using your Backup Exec software.

To archive your tape

1. Choose the **Storage** menu, choose **Slots**, open the context (right-click) menu for the slot you want to export the tape from, choose **Export media**, and then choose **Export media now**. You can select more than one slot and export multiple tapes in a single export operation.



2. In the **Media Request** pop-up window, choose **View details**, and then choose **Respond OK** in the **Alert: Media Intervention** window.

In the Storage Gateway console, you can verify the status of the tape you are archiving. It might take some time to finish uploading data to AWS. During this time, the exported tape is listed in the Tape Gateway VTL with the status **IN TRANSIT TO VTS**. When the upload is completed and the archiving process begins, the status changes to **ARCHIVING**. When data archiving has completed, the exported tape is no longer listed in the VTL but is archived in S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive.

3. Choose your gateway, and then choose **VTL Tape Cartridges** and verify that the virtual tape is no longer listed in your gateway.
4. On the Navigation pane of the Storage Gateway console, choose **Tapes**. Verify that your tape's status is **ARCHIVED**.

Restoring Data from a Tape Archived in Backup Exec

Restoring your archived data is a two-step process.

To restore data from an archived tape

1. Retrieve the archived tape to a Tape Gateway. For instructions, see [Retrieving Archived Tapes](#).
2. Use Backup Exec to restore the data. This process is the same as restoring data from physical tapes. For instructions, see the *Backup Exec Administrative Guide* in the documentation section in the Backup Exec software.

Deactivating a Tape Drive in Backup Exec

A Tape Gateway provides 10 tape drives, but you might decide to use fewer tape drives. In that case, you deactivate the tape drives you don't use.

1. Open Backup Exec, and choose the **Storage** tab.
2. In the **Robotic library** tree, open the context (right-click) menu for the tape drive you want to deactivate, and then choose **Disable**.

Next Step

[Cleaning Up Resources You Don't Need](#)

Testing Your Setup by Using Veritas NetBackup

You can back up your data to virtual tapes, archive the tapes, and manage your virtual tape library (VTL) devices by using Veritas NetBackup. In this topic, you can find basic documentation on how to configure the NetBackup application for a Tape Gateway and perform a backup and restore operation. To do so, you can use the following versions of NetBackup:

- Veritas NetBackup 7.x
- Veritas NetBackup 8.x

The procedure for using these versions of Backup Exec with a Tape Gateway is similar. For detailed information about how to use NetBackup, see the [Veritas Services and Operations Readiness Tools \(SORT\)](#) on the Veritas website. For Veritas support information on hardware compatibility, see the [NetBackup 7.0 - 7.6.x Hardware Compatibility List](#), [NetBackup 8.0 - 8.1.x Hardware Compatibility List](#), or [NetBackup 8.2 - 8.x.x Hardware Compatibility List](#) on the Veritas website.

For more information about compatible backup applications, see [Supported third-party backup applications for a Tape Gateway](#).

Topics

- [Configuring NetBackup Storage Devices](#)
- [Backing Up Data to a Tape](#)
- [Archiving the Tape](#)
- [Restoring Data from the Tape](#)

Configuring NetBackup Storage Devices


After you have connected the virtual tape library (VTL) devices to the Windows client, you configure Veritas NetBackup storage to recognize your devices. For information about how to connect VTL devices to the Windows client, see [Connecting Your VTL Devices](#).

To configure NetBackup to use storage devices on your Tape Gateway


1. Open the NetBackup Administration Console and run it as an administrator.



2. Choose **Configure Storage Devices** to open the Device Configuration wizard.
3. Choose **Next**. The NetBackup application detects your computer as a device host.
4. In the **Device Hosts** column, select your computer, and then choose **Next**. The NetBackup application scans your computer for devices and discovers all devices.



5. In the **Scanning Hosts** page, choose **Next**, and then choose **Next**. The NetBackup application finds all 10 tape drives and the medium changer on your computer.




6. In the **Backup Devices** window, choose **Next**.
7. In the **Drag and Drop Configuration** window, verify that your medium changer is selected, and then choose **Next**.
8. In the dialog box that appears, choose **Yes** to save the configuration on your computer. The NetBackup application updates the device configuration.


9. When the update is completed, choose **Next** to make the devices available to the NetBackup application.
10. In the **Finished!** window, choose **Finish**.

To verify your devices in the NetBackup application

1. In the NetBackup Administration Console, expand the **Media and Device Management** node, and then expand the **Devices** node. Choose **Drives** to display all the tape drives.

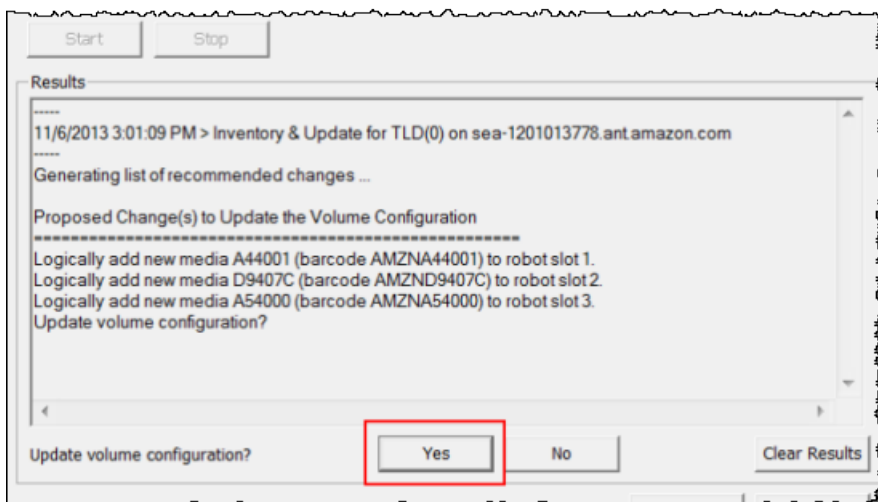


2. In the **Devices** node, choose **Robots** to display all your medium changers. In the NetBackup application, the medium changer is called a *robot*.
3. In the **All Robots** pane, open the context (right-click) menu for **TLD(0)** (that is, your robot), and then choose **Inventory Robot**.
4. In the **Robot Inventory** window, verify that your host is selected from the **Device-Host** list located in the **Select robot** category.
5. Verify that your robot is selected from the **Robot** list.
6. In the **Robot Inventory** window, select **Update volume configuration**, select **Preview changes**, select **Empty media access port prior to update**, and then choose **Start**.




The process then inventories your medium changer and virtual tapes in the NetBackup Enterprise Media Management (EMM) database. NetBackup stores media information, device configuration, and tape status in the EMM.

- In the **Robot Inventory** window, choose **Yes** once the inventory is complete. Choosing **Yes** here updates the configuration and moves virtual tapes found in import/export slots to the virtual tape library.



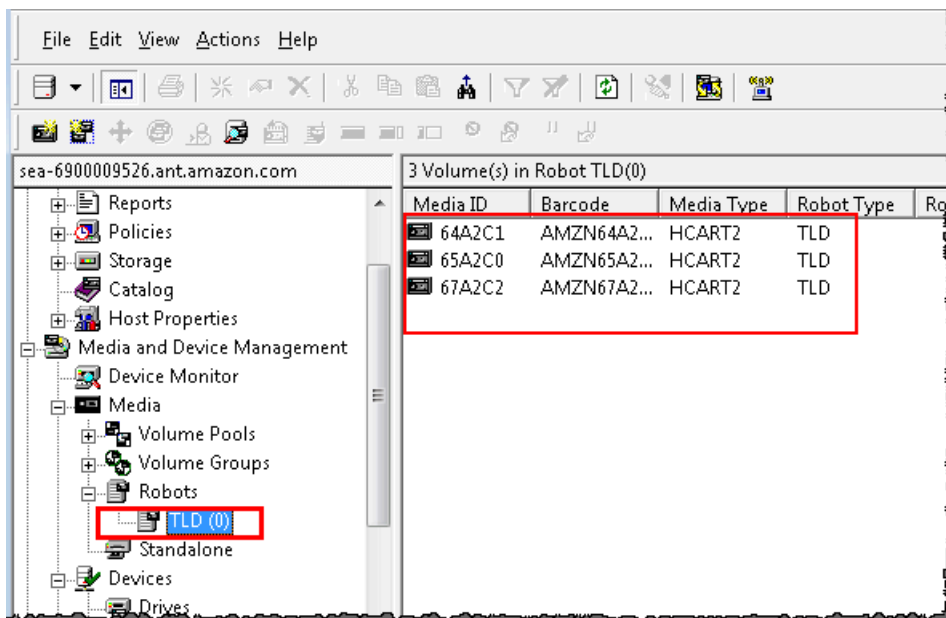
For example, the following screenshot shows three virtual tapes found in the import/export slots.



8. Close the **Robot Inventory** window.
9. In the **Media** node, expand the **Robots** node and choose **TLD(0)** to show all virtual tapes that are available to your robot (medium changer).

Note

If you have previously connected other devices to the NetBackup application, you might have multiple robots. Make sure that you select the right robot.



Now that you have connected your devices and made them available to your backup application, you are ready to test your gateway. To test your gateway, you back up data onto the virtual tapes you created and archive the tapes.

Backing Up Data to a Tape

You test the Tape Gateway setup by backing up data onto your virtual tapes.


Note

- You should back up only a small amount of data for this Getting Started exercise, because there are costs associated with storing, archiving, and retrieving data. For pricing information, see [Pricing](#) on the Storage Gateway detail page.
- If your Tape Gateway restarts for any reason during an ongoing backup job, the backup job will be suspended. The suspended backup job will resume automatically when your gateway finishes restarting.

To create a volume pool

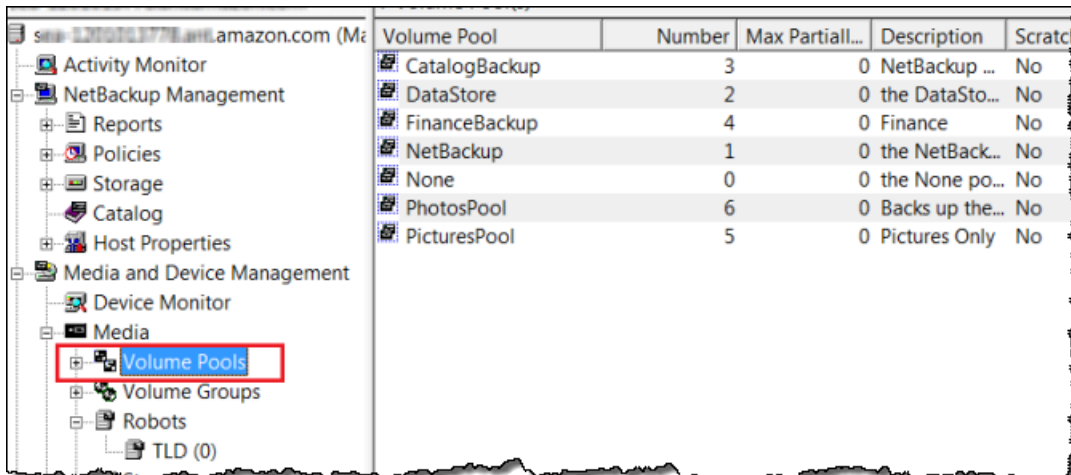
A *volume pool* is a collection of virtual tapes to use for a backup.

1. Start the NetBackup Administration Console.
2. Expand the **Media** node, open the context (right-click) menu for **Volume Pool**, and then choose **New**. The **New Volume Pool** dialog box appears.



3. For **Name**, type a name for your volume pool.
4. For **Description**, type a description for the volume pool, and then choose **OK**. The volume pool you just created is added to the volume pool list.

The following screenshot shows a list of volume pools.




Volume Pool	Number	Max Partial...	Description	Scratch
CatalogBackup	3	0	NetBackup ...	No
DataStore	2	0	the DataSto...	No
FinanceBackup	4	0	Finance	No
NetBackup	1	0	the NetBack...	No
None	0	0	the None po...	No
PhotosPool	6	0	Backs up the...	No
PicturesPool	5	0	Pictures Only	No

To add virtual tapes to a volume pool

1. Expand the **Robots** node, and select the **TLD(0)** robot to display the virtual tapes this robot is aware of.

If you have previously connected a robot, your Tape Gateway robot might have a different name.

2. From the list of virtual tapes, open the context (right-click) menu for the tape you want to add to the volume pool, and choose **Change** to open the **Change Volumes** dialog box. The following screenshot shows the **Change Volumes** dialog box.



3. For **Volume Pool**, choose **New pool**.
4. For **New pool**, select the pool you just created, and then choose **OK**.


You can verify that your volume pool contains the virtual tape that you just added by expanding the **Media** node and choosing your volume pool.

To create a backup policy

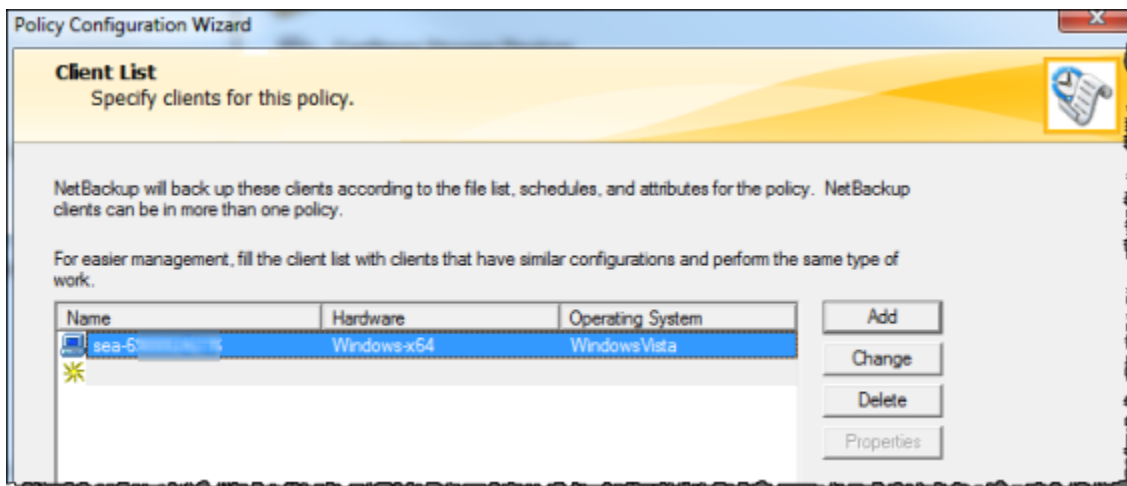
The backup policy specifies what data to back up, when to back it up, and which volume pool to use.

1. Choose your **Master Server** to return to the Veritas NetBackup console.


The following screenshot shows the NetBackup console with **Create a Policy** selected.



2. Choose **Create a Policy** to open the **Policy Configuration Wizard** window.
3. Select **File systems, databases, applications**, and choose **Next**.
4. For **Policy Name**, type a name for your policy and verify that **MS-Windows** is selected from the **Select the policy type** list, and then choose **Next**.
5. In the **Client List** window, choose **Add**, type the host name of your computer in the **Name** column, and then choose **Next**. This step applies the policy you are defining to localhost (your client computer).



6. In the **Files** window, choose **Add**, and then choose the folder icon.




7. In the **Browse** window, browse to the folder or files you want to back up, choose **OK**, and then choose **Next**.
8. In the **Backup Types** window, accept the defaults, and then choose **Next**.


Note

If you want to initiate the backup yourself, select **User Backup**.

9. In the **Frequency and Retention** window, select the frequency and retention policy you want to apply to the backup. For this exercise, you can accept all the defaults and choose **Next**.



10. In the **Start** window, select **Off hours**, and then choose **Next**. This selection specifies that your folder should be backed up during off hours only.




11. In the **Policy Configuration** wizard, choose **Finish**.


The policy runs the backups according to the schedule. You can also perform a manual backup at any time, which we do in the next step.

To perform a manual backup

1. On the navigation pane of the NetBackup console, expand the **NetBackup Management** node.
2. Expand the **Policies** node.
3. Open the context (right-click) menu for your policy, and choose **Manual Backup**.



4. In the **Manual Backup** window, select a schedule, select a client, and then choose **OK**.



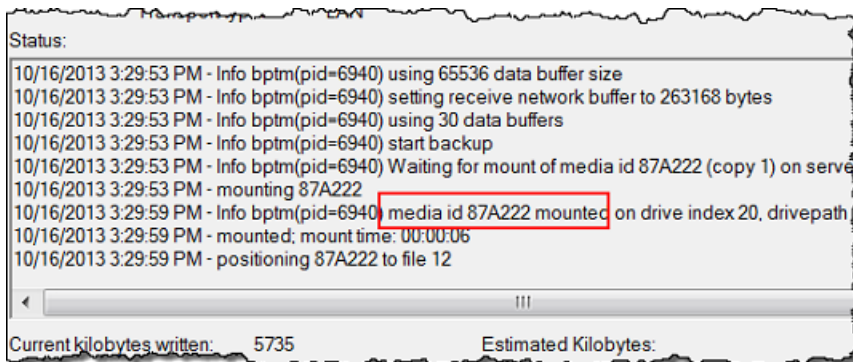
5. In the **Manual Backup Started** dialog box that appears, choose **OK**.
6. On the navigation pane, choose **Activity Monitor** to view the status of your backup in the **Job ID** column.

nbtest: 11 Jobs (0 Queued 0 Active 0 Waiting for Retry 0 Suspended 0 Incomplete 11 Done)								
Job ID	Type	Job State	State Details	Status	Job Policy	Job Schedule	Client	
18	Backup	Done		0	EC2Policy	Full	localhost	
17	Backup	Done		0	EC2Policy	Full	localhost	
14	Backup	Done		0	EC2Policy	Full	localhost	
10	Image Cleanup	Done		1				
11	Image Cleanup	Done		1				

To find the barcode of the virtual tape where NetBackup wrote the file data during the backup, look in the **Job Details** window as described in the following procedure. You need this barcode in the procedure in the next section, where you archive the tape.

To find the barcode of a tape

1. In **Activity Monitor**, open the context (right-click) menu for the identifier of your backup job in the **Job ID** column, and then choose **Details**.
2. In the **Job Details** window, choose the **Detailed Status** tab.
3. In the **Status** box, locate the media ID. For example, in the following screenshot, the media ID is **87A222**. This ID helps you determine which tape you have written data to.



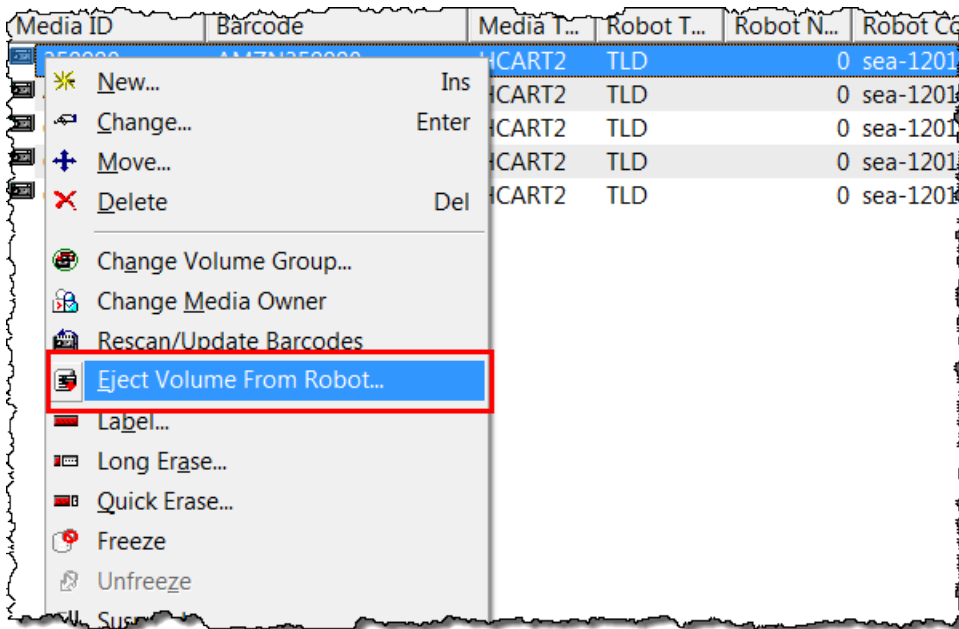
You have now successfully deployed a Tape Gateway, created virtual tapes, and backed up your data. Next, you can archive the virtual tapes and retrieve them from the archive.

Archiving the Tape

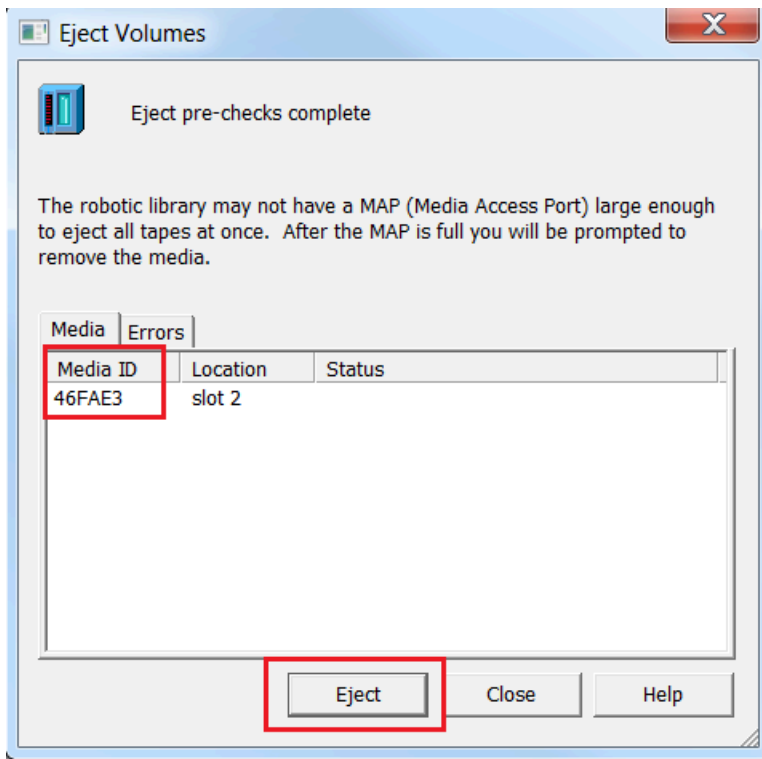
When you archive a tape, Tape Gateway moves the tape from your gateway's virtual tape library (VTL) to the archive, which provides offline storage. You initiate tape archival by ejecting the tape using your backup application.

To archive a virtual tape


1. In the NetBackup Administration console, expand the **Media and Device Management** node, and expand the **Media** node.
2. Expand **Robots** and choose **TLD(0)**.
3. Open the context (right-click) menu for the virtual tape you want to archive, and choose **Eject Volume From Robot**.




4. In the **Eject Volumes** window, make sure the **Media ID** matches the virtual tape you want to eject, and then choose **Eject**.



5. In the dialog box, choose **Yes**. The dialog box is shown following.



When the eject process is completed, the status of the tape in the **Eject Volumes** dialog box indicates that the eject succeeded.



6. Choose **Close** to close the **Eject Volumes** window.
7. In the Storage Gateway console, verify the status of the tape you are archiving in the gateway's VTL. It can take some time to finish uploading data to AWS. During this time, the ejected tape is listed in the gateway's VTL with the status **IN TRANSIT TO VTS**. When archiving starts, the status is **ARCHIVING**. Once data upload has completed, the ejected tape is no longer listed in the VTL but is archived in S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive.
8. To verify that the virtual tape is no longer listed in your gateway, choose your gateway, and then choose **VTL Tape Cartridges**.
9. In the navigation pane of the Storage Gateway console, choose **Tapes**. Verify that your archived tape's status is **ARCHIVED**.

Restoring Data from the Tape

Restoring your archived data is a two-step process.

To restore data from an archived tape

1. Retrieve the archived tape to a Tape Gateway. For instructions, see [Retrieving Archived Tapes](#).

2. Use the Backup, Archive, and Restore software installed with the Veritas NetBackup application. This process is the same as restoring data from physical tapes. For instructions, see [Veritas Services and Operations Readiness Tools \(SORT\)](#) on the Veritas website.

Next Step

[Cleaning Up Resources You Don't Need](#)

Where Do I Go from Here?

After your Tape Gateway is in production, you can perform several maintenance tasks, such as adding and removing tapes, monitoring and optimizing gateway performance, and troubleshooting. For general information about these management tasks, see [Managing Your Gateway](#).

You can perform some of the Tape Gateway maintenance tasks on the AWS Management Console, such as configuring your gateway's bandwidth rate limits and managing gateway software updates. If your Tape Gateway is deployed on-premises, you can perform some maintenance tasks on the gateway's local console. These include routing your Tape Gateway through a proxy and configuring your gateway to use a static IP address. If you are running your gateway as an Amazon EC2 instance, you can perform specific maintenance tasks on the Amazon EC2 console, such as adding and removing Amazon EBS volumes. For more information on maintaining your Tape Gateway, see [Managing Your Tape Gateway](#).

If you plan to deploy your gateway in production, you should take your real workload into consideration in determining the disk sizes. For information on how to determine real-world disk sizes, see [Managing local disks for your Storage Gateway](#). Also, consider cleaning up if you don't plan to continue using your Tape Gateway. Cleaning up lets you avoid incurring charges. For information on cleanup, see [Cleaning Up Resources You Don't Need](#).

Cleaning Up Resources You Don't Need

If you created the gateway as an example exercise or a test, consider cleaning up to avoid incurring unexpected or unnecessary charges.

If you plan to continue using your Tape Gateway, see additional information in [Where Do I Go from Here?](#)

To clean up resources you don't need

1. Delete tapes from both your gateway's virtual tape library (VTL) and archive. For more information, see [Deleting Your Gateway by Using the AWS Storage Gateway Console and Removing Associated Resources](#).
 - a. Archive any tapes that have the **RETRIEVED** status in your gateway's VTL. For instructions, see [Archiving Tapes](#).
 - b. Delete any remaining tapes from your gateway's VTL. For instructions, see [Deleting Tapes](#).
 - c. Delete any tapes you have in the archive. For instructions, see [Deleting Tapes](#).
2. Unless you plan to continue using the Tape Gateway, delete it: For instructions, see [Deleting Your Gateway by Using the AWS Storage Gateway Console 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.

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 Tape 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 Gateway

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

Gateway software releases will periodically include OS updates and security patches that have been validated. These updates are applied as part of the regular gateway update process during a scheduled maintenance window, and are typically released every six months. Note: Users should treat the Storage Gateway appliance as a managed virtual machine, and should not attempt to access or modify the Storage Gateway appliance instance. Attempting to install or update any software packages using other methods (ex: SSM or Hypervisor tools) than the normal gateway update mechanism may result in disruption to the proper functioning of the Gateway.

Topics

- [Managing Your Tape Gateway](#)
- [Moving your data to a new gateway](#)

Managing Your Tape Gateway

Following, you can find information about how to manage your Tape Gateway resources in AWS Storage Gateway.

Topics

- [Editing Basic Gateway Information](#)
- [Adding Virtual Tapes](#)
- [Managing Automatic Tape Creation](#)
- [Archiving Virtual Tapes](#)
- [Moving Your Tape from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive Storage Class](#)
- [Retrieving Archived Tapes](#)
- [Viewing Tape Usage](#)
- [Deleting Tapes](#)
- [Deleting Custom Tape Pools](#)
- [Deactivating Your Tape Gateway](#)

- [Understanding Tape Status](#)

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. Modify the settings you want to change, then choose **Save changes**.

Note

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.

Adding Virtual Tapes

You can add tapes in your Tape Gateway when you need them. For information about how to create tapes, see [Creating Tapes](#).

After your tape is created, you can find information about it on the **Tape overview** page. By default, this list displays up to 1,000 tapes at a time, but the searches that you perform apply to all of your tapes. You can use the search bar to find tapes that match a specific criteria, or to reduce the list to less than 1,000 tapes. When your list contains 1,000 tapes or fewer, you can then sort your tapes in ascending or descending order by various properties. For information about Tape Gateway tape quotas, see [AWS Storage Gateway quotas](#).

Managing Automatic Tape Creation

The Tape Gateway automatically creates new virtual tapes to maintain the minimum number of available tapes that you configure. It then makes these new tapes available for import by the

backup application so that your backup jobs can run without interruption. Automatic tape creation removes the need for custom scripting in addition to the manual process for creating new virtual tapes.

To delete an automatic tape creation policy

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose the **Gateways** tab.
3. Choose the gateway for which you need to manage automatic tape creation.
4. In the **Actions** menu, choose **Configure tape auto-create**.
5. To delete an automatic tape creation policy on a gateway, choose **Remove** to the right of the policy you want to delete.

To stop automatic tape creation on a gateway, delete all of the automatic tape creation policies for that gateway.

Choose **Save changes** to confirm deletion of tape auto-create policies for the selected Tape Gateway.


Note

Deleting a tape auto-creation policy from a gateway cannot be undone.

To change the automatic tape creation policies for a Tape Gateway

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose the **Gateways** tab.
3. Choose the gateway for which you need to manage automatic tape creation.
4. In the **Actions** menu, choose **Configure tape auto-create**, and change the settings on the page that appears.
5. For **Minimum number of tapes**, enter the minimum number of virtual tapes that should be available on the Tape Gateway at all times. The valid range for this value is a minimum of 1 and a maximum of 10.

6. For **Capacity**, enter the size, in bytes of the virtual tape capacity. The valid range for this value is a minimum of 100 GiB and a maximum of 15 TiB.
7. For **Barcode prefix**, enter the prefix that you want to prepend to the barcode of your virtual tapes.

 **Note**

Virtual tapes are uniquely identified by a barcode, and you can add a prefix to the barcode. The prefix is optional, but you can use it to help identify your virtual tapes. The prefix must be uppercase letters (A–Z) and must be one to four characters long.

8. For **Pool**, choose **Glacier Pool** or **Deep Archive Pool**. This pool represents the storage class in which your tapes are stored when they are ejected by your backup software.
 - Choose **Glacier Pool** if you want to archive the tapes in the S3 Glacier Flexible Retrieval storage class. When your backup software ejects the tapes, they are automatically archived in S3 Glacier Flexible Retrieval. You use S3 Glacier Flexible Retrieval for more active archives, where you can retrieve a tape typically within 3–5 hours. For detailed information, see [Storage Classes for Archiving Objects](#) in the *Amazon Simple Storage Service User Guide*.
 - Choose **Deep Archive Pool** if you want to archive the tapes in S3 Glacier Deep Archive. When your backup software ejects the tape, the tape is automatically archived in S3 Glacier Deep Archive. You use S3 Glacier Deep Archive for long-term data retention and digital preservation, where data is accessed once or twice a year. You can retrieve a tape archived in S3 Glacier Deep Archive typically within 12 hours. For detailed information, see [Storage Classes for Archiving Objects](#) in the *Amazon Simple Storage Service User Guide*.

If you archive tapes in S3 Glacier Flexible Retrieval, you can move them to S3 Glacier Deep Archive later. For more information, see [Moving Your Tape from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive Storage Class](#).

9. You can find information about your tapes on the **Tape overview** page. By default, this list displays up to 1,000 tapes at a time, but the searches that you perform apply to all of your tapes. You can use the search bar to find tapes that match a specific criteria, or to reduce the list to less than 1,000 tapes. When your list contains 1,000 tapes or fewer, you can then sort your tapes in ascending or descending order by various properties.

The status of available virtual tapes is initially set to **CREATING** when the tapes are being created. After the tapes are created, their status changes to **AVAILABLE**. For more information, see [Managing Your Tape Gateway](#).

For more information about enabling automatic tape creation, see [Creating Tapes Automatically](#).

Archiving Virtual Tapes

You can archive your tapes to S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive. When you create a tape, you choose the archive pool that you want to use to archive your tape.

You choose **Glacier Pool** if you want to archive the tape in S3 Glacier Flexible Retrieval. When your backup software ejects the tape, it is automatically archived in S3 Glacier Flexible Retrieval. You use S3 Glacier Flexible Retrieval for more active archives where the data is regularly retrieved and needed in minutes. For detailed information, see [Storage Classes for Archiving Objects](#)

You choose **Deep Archive Pool** if you want to archive the tape in S3 Glacier Deep Archive. When your backup software ejects the tape, the tape is automatically archived in S3 Glacier Deep Archive. You use S3 Glacier Deep Archive for long-term data retention and digital preservation at a very low cost. Data in S3 Glacier Deep Archive is not retrieved often or is rarely retrieved. For detailed information, see [Storage Classes for Archiving Objects](#).

Note

Any tape created before March 27, 2019, are archived directly in S3 Glacier Flexible Retrieval when your backup software ejects it.

When your backup software ejects a tape, it is automatically archived in the pool that you chose when you created the tape. The process for ejecting a tape varies depending on your backup software. Some backup software requires that you export tapes after they are ejected before archiving can begin. For information about supported backup software, see [Using Your Backup Software to Test Your Gateway Setup](#).

Moving Your Tape from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive Storage Class

Move your tapes from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive for long-term data retention and digital preservation at a very low cost. You use S3 Glacier Deep Archive for long-term data retention and digital preservation where the data is accessed once or twice a year. For detailed information, see [Storage Classes for Archiving Objects](#).

To move a tape from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive

1. In the navigation pane, choose **Tape Library > Tapes** to see your tapes. By default, this list displays up to 1,000 tapes at a time, but the searches that you perform apply to all of your tapes. You can use the search bar to find tapes that match a specific criteria, or to reduce the list to less than 1,000 tapes. When your list contains 1,000 tapes or fewer, you can then sort your tapes in ascending or descending order by various properties.
2. Select the check boxes for the tapes you want to move to S3 Glacier Deep Archive. You can see the pool that each tape is associated with in the **Pool** column.
3. Choose **Assign to pool**.
4. In the Assign tape to pool dialog box, verify the barcodes for the tapes you are moving and choose **Assign**.

Note

If a tape has been ejected by the backup application and archived in S3 Glacier Deep Archive, you can't move it back to S3 Glacier Flexible Retrieval. There's a charge for moving your tapes from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive. In addition, if you move tapes from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive prior to 90 days, there is an early deletion fee for S3 Glacier Flexible Retrieval.

5. After the tape is moved, you can see the updated status in the **Pool** column on the **Tape overview** page.

Retrieving Archived Tapes

To access data stored on an archived virtual tape, you must first retrieve the tape that you want to your Tape Gateway. Your Tape Gateway provides one virtual tape library (VTL) for each gateway.

If you have more than one Tape Gateway in an AWS Region, you can retrieve a tape to only one gateway.

The retrieved tape is write-protected; you can only read the data on the tape.

⚠ Important

If you archive a tape in S3 Glacier Flexible Retrieval, you can retrieve the tape typically within 3-5 hours. If you archive the tape in S3 Glacier Deep Archive, you can retrieve it typically within 12 hours.

ℹ Note

There is a charge for retrieving tapes from archive. For detailed pricing information, see [Storage Gateway Pricing](#).

To retrieve an archived tape to your gateway

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose **Tape Library > Tapes** to see your tapes. By default, this list displays up to 1,000 tapes at a time, but the searches that you perform apply to all of your tapes. You can use the search bar to find tapes that match a specific criteria, or to reduce the list to less than 1,000 tapes. When your list contains 1,000 tapes or fewer, you can then sort your tapes in ascending or descending order by various properties.
3. Choose the virtual tape you want to retrieve from the **Virtual Tape Shelf** tab, and choose **Retrieve tape**.


ℹ Note

The status of the virtual tape that you want to retrieve must be ARCHIVED.

4. In the **Retrieve tape** dialog box, for **Barcode**, verify that the barcode identifies the virtual tape you want to retrieve.
5. For **Gateway**, choose the gateway that you want to retrieve the archived tape to, and then choose **Retrieve tape**.

The status of the tape changes from ARCHIVED to RETRIEVING. At this point, your data is being moved from the virtual tape shelf (backed by S3 Glacier Flexible Retrieval or S3 Glacier Deep

Archive) to the virtual tape library (backed by Amazon S3). After all the data is moved, the status of the virtual tape in the archive changes to RETRIEVED.

 **Note**


Retrieved virtual tapes are read-only.

Viewing Tape Usage

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

To view the amount of data stored on a tape

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose **Tape Library > Tapes** to see your tapes. By default, this list displays up to 1,000 tapes at a time, but the searches that you perform apply to all of your tapes. You can use the search bar to find tapes that match a specific criteria, or to reduce the list to less than 1,000 tapes. When your list contains 1,000 tapes or fewer, you can then sort your tapes in ascending or descending order by various properties.
3. Choose the tape you are interested in.
4. The page that appears provides various details and information about the tape, including the following:
 - **Size:** The total capacity of the selected tape.
 - **Used:** The size of data written to the tape by your backup application.

 **Note**

This value is not available for tapes created before May 13, 2015.

Deleting Tapes

You can delete virtual tapes from your Tape Gateway by using the Storage Gateway console.

Note

If the tape you want to delete from your Tape Gateway has a status of `RETRIEVED`, you must first eject the tape using your backup application before deleting the tape. For instructions on how to eject a tape using the Symantec NetBackup software, see [Archiving the Tape](#). After the tape is ejected, the tape status changes back to `ARCHIVED`. You can then delete the tape.

Make copies of your data before you delete your tapes. After you delete a tape, you can't get it back.

To delete a virtual tape**Warning**

This procedure permanently deletes the selected virtual tape.

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose **Tape Library > Tapes** to see your tapes. By default, this list displays up to 1,000 tapes at a time, but the searches that you perform apply to all of your tapes. You can use the search bar to find tapes that match a specific criteria, or to reduce the list to less than 1,000 tapes. When your list contains 1,000 tapes or fewer, you can then sort your tapes in ascending or descending order by various properties.
3. Select one or more tapes to delete.
4. For **Actions** choose **Delete tape**. The confirmation dialog box appears.
5. Verify that you want to delete the specified tapes, then type the word *delete* in the confirmation box and choose **Delete**.

After the tape is deleted, it disappears from the Tape Gateway.

Deleting Custom Tape Pools

You can delete a custom tape pool only if there are no archived tapes in the pool, and there are no automatic tape creation policies attached to the pool.

To delete your custom tape pool

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose **Pools** to see the available pools.
3. Select one or more tape pools to delete.

If the **Tape Count** for the tape pools that you want to delete is **0**, and if there are no automatic tape creation policies that reference the custom tape pool, you can delete the pools.

4. Choose **Delete**. The confirmation dialog box appears.
5. Verify that you want to delete the specified tape pools, then type the word *delete* in the confirmation box and choose **Delete**.

Warning

This procedure permanently deletes the selected tape pools and can't be undone.

After the tape pools are deleted, they disappear from the tape library.

Deactivating Your Tape Gateway

You deactivate a Tape Gateway if the Tape Gateway has failed and you want to recover the tapes from the failed gateway to another gateway.

To recover the tapes, you must first deactivate the failed gateway. Deactivating a Tape Gateway locks down the virtual tapes in that gateway. That is, any data that you might write to these tapes after deactivating the gateway isn't sent to AWS. You can only deactivate a gateway on the Storage Gateway console if the gateway is no longer connected to AWS. If the gateway is connected to AWS, you can't deactivate the Tape Gateway.

You deactivate a Tape Gateway as part of data recovery. For more information about recovering tapes, see [You Need to Recover a Virtual Tape from a Malfunctioning Tape Gateway](#).

To deactivate your 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 failed gateway.
3. Choose the **Details** tab for the gateway to display the deactivate gateway message.
4. Choose **Create recovery tapes**.
5. Choose **Disable gateway**.

Understanding Tape Status

Each tape has an associated status that tells you at a glance what the health of the tape is. Most of the time, the status indicates that the tape is functioning normally and that no action is needed on your part. In some cases, the status indicates a problem with the tape that might require action on your part. You can find information following to help you decide when you need to act.


Topics

- [Understanding Tape Status Information in a VTL](#)
- [Determining Tape Status in an Archive](#)

Understanding Tape Status Information in a VTL

A tape's status must be AVAILABLE for you to read or write to the tape. The following table lists and describes possible status values.

Status	Description	Tape Data Is Stored In
CREATING	The virtual tape is being created. The tape can't be loaded into a tape drive, because the tape is being created.	—
AVAILABLE	The virtual tape is created and ready to be loaded into a tape drive.	Amazon S3
IN TRANSIT TO VTS	The virtual tape has been ejected and is being uploaded for archive. At this point, your Tape Gateway is uploading data to AWS. If the amount of data being uploaded is small, this status might not appear. When the upload is completed, the status changes to ARCHIVING.	Amazon S3

Status	Description	Tape Data Is Stored In
ARCHIVING	The virtual tape is being moved by your Tape Gateway to the archive, which is backed by S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive. This process happens after the data upload to AWS is completed.	Data is being moved from Amazon S3 to S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive.
DELETING	The virtual tape is being deleted.	Data is being deleted from Amazon S3
DELETED	The virtual tape has been successfully deleted.	—
RETRIEVING	The virtual tape is being retrieved from the archive to your Tape Gateway. <div data-bbox="324 835 1096 1054" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p> Note The virtual tape can be retrieved only to a Tape Gateway.</p> </div>	Data is being moved from S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive to Amazon S3
RETRIEVED	The virtual tape is retrieved from the archive. The retrieved tape is write-protected.	Amazon S3
RECOVERED	The virtual tape is recovered and is read-only. When your Tape Gateway is not accessible for any reason, you can recover virtual tapes associated with that Tape Gateway to another Tape Gateway. To recover the virtual tapes, first deactivate the inaccessible Tape Gateway.	Amazon S3
IRRECOVERABLE	The virtual tape can't be read from or written to. This status indicates an error in your Tape Gateway.	Amazon S3

Determining Tape Status in an Archive


You can use the following procedure to determine the status of a virtual tape in an archive.

To determine the status of a virtual tape

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose **Tapes**.
3. In the **Status** column of the tape library grid, check the status of the tape.

The tape status also appears in the **Details** tab of each virtual tape.

Following, you can find a description of the possible status values.

Status	Description
ARCHIVED	The virtual tape has been ejected and is uploaded to the archive.
RETRIEVING	The virtual tape is being retrieved from the archive. <div style="border: 1px solid #0070C0; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p> Note</p> <p>The virtual tape can be retrieved only to a Tape Gateway.</p> </div>
RETRIEVED	The virtual tape has been retrieved from the archive. The retrieved tape is read-only.

For additional information about how to work with tapes and VTL devices, see [Working With Tapes](#).

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 virtual tapes to a new Tape Gateway

To move your virtual tape to a new Tape Gateway

1. Use your backup application to back up all your data onto a virtual tape. Wait for the backup to finish successfully.
2. Use your backup application to eject your tape. The tape will be stored in one of the Amazon S3 storage classes. Ejected tapes are archived in S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive, and are read-only.

Before proceeding, confirm that the ejected tapes have been archived:

- a. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
- b. In the navigation pane, choose **Tape Library > Tapes** to see your tapes. By default, this list displays up to 1,000 tapes at a time, but the searches that you perform apply to all of your tapes. You can use the search bar to find tapes that match a specific criteria, or to reduce the list to less than 1,000 tapes. When your list contains 1,000 tapes or fewer, you can then sort your tapes in ascending or descending order by various properties.
- c. In the **Status** column of the list, check the status of the tape.

The tape status also appears in the **Details** tab of each virtual tape.

For more information about determining tape status in an archive, see [Determining Tape Status in an Archive](#).

3. Using your backup application, verify that there are no active backup jobs going to the existing Tape Gateway before you stop it. If there are any active backup jobs, wait for them to finish and eject your tapes (see previous step) before stopping the gateway.
4. Use the following steps to stop the existing Tape Gateway:
 - a. In the navigation pane, choose **Gateways**, and then choose the old Tape 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 old Tape 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.

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

5. Create a new Tape Gateway. For detailed instructions, see [Creating a Gateway](#).
6. Use the following steps to create new tapes:
 - a. In the navigation pane, choose the **Gateways** tab.
 - b. Choose **Create tape** to open the **Create tape** dialog box.
 - c. For **Gateway**, choose a gateway. The tape is created for this gateway.
 - d. For **Number of tapes**, choose the number of tapes that you want to create. For more information about tape limits, see [AWS Storage Gateway quotas](#).

You can also set up automatic tape creation at this point. For more information, see [Creating Tapes Automatically](#).

- e. For **Capacity**, enter the size of the virtual tape that you want to create. Tapes must be larger than 100 GiB. For information about capacity limits, see [AWS Storage Gateway quotas](#).
- f. For **Barcode prefix**, enter the prefix that you want to prepend to the barcode of your virtual tapes.

 **Note**


Virtual tapes are uniquely identified by a barcode. You can add a prefix to the barcode. The prefix is optional, but you can use it to help identify your virtual tapes. The prefix must be uppercase letters (A–Z) and must be one to four characters long.

- g. For **Pool**, choose **Glacier Pool** or **Deep Archive Pool**. This pool represents the storage class in which your tape will be stored when it is ejected by your backup software.

Choose **Glacier Pool** if you want to archive the tape in S3 Glacier Flexible Retrieval. When your backup software ejects the tape, it is automatically archived in S3 Glacier Flexible Retrieval. You use S3 Glacier Flexible Retrieval for more active archives where you can retrieve a tape typically within 3-5 hours. For more information, see [Storage classes for archiving objects](#) in the *Amazon Simple Storage Service User Guide*.

Choose **Deep Archive Pool** if you want to archive the tape in S3 Glacier Deep Archive. When your backup software ejects the tape, the tape is automatically archived in S3 Glacier Deep Archive. You use S3 Glacier Deep Archive for long-term data retention and digital preservation where data is accessed once or twice a year. You can retrieve a tape archived in S3 Glacier Deep Archive typically within 12 hours. For more information, see [Storage classes for archiving objects](#) in the *Amazon Simple Storage Service User Guide*.

If you archive a tape in S3 Glacier Flexible Retrieval, you can move it to S3 Glacier Deep Archive later. For more information, see [Moving Your Tape from S3 Glacier Flexible Retrieval to S3 Glacier Deep Archive Storage Class](#).

 **Note**

Tapes created before March 27, 2019, are archived directly in S3 Glacier Flexible Retrieval when your backup software ejects them.

- h. (Optional) For **Tags**, enter a key and value to add tags to your tape. A tag is a case-sensitive key-value pair that helps you manage, filter, and search for your tapes.
 - i. Choose **Create tapes**.
7. Use your backup application to start a backup job, and back up your data to the new tape.
 8. (Optional) If your tape is archived and you need to restore data from it, retrieve it to the new Tape Gateway. The tape will be in read-only mode. For more information about retrieving archived tapes, see [Retrieving Archived Tapes](#).

 **Note**

Outbound data charges might apply.

- a. In the navigation pane, choose **Tape Library > Tapes** to see your tapes. By default, this list displays up to 1,000 tapes at a time, but the searches that you perform apply to all of your

tapes. You can use the search bar to find tapes that match a specific criteria, or to reduce the list to less than 1,000 tapes. When your list contains 1,000 tapes or fewer, you can then sort your tapes in ascending or descending order by various properties.

- b. Choose the virtual tape that you want to retrieve. For **Actions**, choose **Retrieve Tape**.

 **Note**

The status of the virtual tape that you want to retrieve must be ARCHIVED.

- c. In the **Retrieve tape** dialog box, for **Barcode**, verify that the barcode identifies the virtual tape you want to retrieve.
- d. For **Gateway**, choose the new Tape Gateway that you want to retrieve the archived tape to, and then choose **Retrieve tape**.

When you have confirmed that your new Tape Gateway is working correctly, you can delete the old Tape Gateway.

 **Important**

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

9. Use the following steps to delete the old Tape 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 gateway that you want to delete.
- b. For **Actions**, choose **Delete gateway**.

In the confirmation dialog box that appears, make sure that the gateway ID listed specifies the old Tape Gateway that you want to delete, enter **delete** in the confirmation field, and then choose **Delete**.

- c. Delete the VM. For more information about deleting a VM, see the documentation for your hypervisor.

Monitoring Storage Gateway

This section describes how to monitor a 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](#).

Topics

- [Understanding gateway metrics](#)
- [Dimensions for Storage 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 Tape 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
AvailabilityNotifi cations	<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>
CacheUsed	<p>The total number of bytes being used in the gateway's cache storage. The sample</p>

Metric	Description	
	<p>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>	
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>	

Metric	Description
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>

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</p>

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

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 Tape Gateway works \(architecture\)](#).

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.

The following image shows a typical email notification.



- 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 Tape Gateway works \(architecture\)](#).

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

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 Tape Gateway, see [Understanding gateway metrics](#) and [Understanding Virtual Tape 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 .

State	Description
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 Tape 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 Tape Gateway

This section describes how to monitor your Tape Gateway, virtual tapes associated with your Tape Gateway, cache storage, and the upload buffer. You use the AWS Management Console to view metrics for your Tape Gateway. With metrics, you can track the health of your Tape Gateway and set up alarms to notify you when one or more metrics are 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 of how your Tape Gateway and virtual tapes are performing. For detailed information about CloudWatch, see the [Amazon CloudWatch User Guide](#).

Topics

- [Getting Tape Gateway Health Logs with CloudWatch Log Groups](#)
- [Using Amazon CloudWatch Metrics](#)
- [Understanding Virtual Tape Metrics](#)
- [Measuring Performance Between Your Tape Gateway and AWS](#)

Getting Tape Gateway Health Logs with CloudWatch Log Groups

You can use Amazon CloudWatch Logs to get information about the health of your Tape Gateway and related resources. You can use the 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 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 Tape 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 virtual tape 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 File 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 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 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 CloudWatch console.

Following is an example of a Tape Gateway event message that is sent to CloudWatch. This example shows a `TapeStatusTransition` message.

```
{
  "severity": "INFO",
  "source": "FZTT16FCF5",
  "type": "TapeStatusTransition",
  "gateway": "sgw-C51DFEAC",
  "timestamp": "1581553463831",
  "newStatus": "RETRIEVED"
```

}

Using Amazon CloudWatch Metrics

You can get monitoring data for your Tape Gateway by 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. The CloudWatch API can also be used through one of the [Amazon AWS Software Development Kits \(SDKs\)](#) or the [Amazon CloudWatch API](#) tools. Depending on your needs, you might prefer to use either the graphs displayed in the console or retrieved from the API.

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

- 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` and `GatewayName`. In the CloudWatch console, you can use the `Gateway Metrics` view to easily select gateway-specific and tape-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 are available to you.

Amazon CloudWatch Namespace	Dimension	Description
AWS/StorageGateway	<code>GatewayId</code> , <code>GatewayName</code>	These dimensions filter for metric data that describes aspects of the Tape Gateway. You can identify a Tape Gateway to work with by specifying both the <code>GatewayId</code> and the <code>GatewayName</code> dimensions. Throughput and latency data of a Tape Gateway is based on all the virtual tapes in the Tape Gateway.

Amazon CloudWatch Namespace	Dimension	Description
		Data is available automatically in 5-minute periods at no charge.

Working with gateway and tape 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](#)

Understanding Virtual Tape Metrics

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

Some tape-specific metrics might have the same name as certain gateway-specific metrics. These metrics represent the same kinds of measurements but are scoped to a tape instead of a gateway. Before starting work, specify whether you want to work with a gateway metric or a tape metric. When working with tape metrics, specify the tape ID for the tape 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 tapes.

Metric	Description
CachePercentDirty	<p>The tape'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>
CloudTraffic	<p>The amount of bytes uploaded and downloaded from the cloud to the tape.</p> <p>Units: bytes</p>
IoWaitPercent	<p>The percentage of allocated IoWait units that are currently used by the tape.</p> <p>Units: Percent</p>
HealthNotification	<p>The number of health notifications sent by the tape.</p> <p>Units: count</p>
MemUsedBytes	<p>The percentage of allocated memory that is currently used by the tape.</p> <p>Units: Bytes</p>
MemTotalBytes	<p>The percentage of total memory that is currently used by the tape.</p> <p>Units: Bytes</p>

Metric	Description
ReadBytes	<p>The total number of bytes read from your on-premises applications in the reporting period for a file share.</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>
UserCpuPercent	<p>The percentage of allocated CPU compute units for the user that are currently used by the tape.</p> <p>Units: Percent</p>
WriteBytes	<p>The total number of bytes written to your on-premises applications in the reporting period.</p> <p>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.</p> <p>Units: Bytes</p>

Measuring Performance Between Your Tape Gateway and AWS

Data throughput, data latency, and operations per second are measures that you can use to understand how your application storage that is using your Tape Gateway is performing. When you use the correct aggregation statistic, these values can be measured by using the Storage Gateway metrics that are provided for you.

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), 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 the corresponding statistic you can use to measure the throughput, latency, and IOPS between your Tape Gateway and AWS.


Item of Interest	How to Measure
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.
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 Tape Gateway as a rate in 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 Tape Gateway to AWS

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose the **Metrics** tab.
3. Choose the **StorageGateway: Gateway Metrics** dimension, and find the Tape Gateway 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.

The following image shows the `CloudBytesUploaded` metric for a gateway tape with the `Sum` statistic. In the image, placing the cursor over a data point displays information about the data

point, including its value and the number of bytes uploaded. Divide this value by the **Period** value (5 minutes) to get the throughput at that sample point. For the point highlighted, the throughput from the Tape Gateway to AWS is 555,544,576 bytes divided by 300 seconds, which is 1.7 megabytes per second.



To measure the data latency from a Tape Gateway to AWS

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose the **Metrics** tab.
3. Choose the **StorageGateway: GatewayMetrics** dimension, and find the Tape Gateway 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 Tape Gateway's throughput to 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 Tape Gateway that you want to work with.
4. Choose the `CloudBytesUploaded` metric.
5. Define the alarm by defining 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 megabytes for 60 minutes.
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**.

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 Tape 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**.

Maintaining Your Gateway

Maintaining your gateway includes tasks such as configuring cache storage and upload buffer space, and doing general maintenance your gateway's performance. These tasks are common to all gateway types. If you haven't created a gateway, see [Creating Your Gateway](#).

Topics

- [Shutting Down Your Gateway VM](#)
- [Managing local disks for your Storage Gateway](#)
- [Managing Bandwidth for Your Tape Gateway](#)
- [Managing gateway updates using the AWS Storage Gateway console](#)
- [Performing Maintenance Tasks on the Local Console](#)
- [Deleting Your Gateway by Using the AWS Storage Gateway Console and Removing Associated Resources](#)

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. For File Gateway, you just shutdown your VM. 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 Local Console Using Default Credentials](#).
- Storage Gateway API—see [ShutdownGateway](#)

For File Gateway, you simply shutdown your VM. You don't shutdown the gateway.

Starting and Stopping a Tape Gateway

To stop a Tape 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 Tape 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**.

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](#)
- [Optimizing Gateway Performance](#)
- [Determining the size of upload buffer to allocate](#)
- [Determining the size of cache storage to allocate](#)
- [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 (see [Plan your Storage Gateway deployment](#)), the gateway requires the following additional storage:

- Tape 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.

Local storage	Description	
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.

Optimizing Gateway Performance

To achieve optimal performance use high throughput SSD disks for both cache and upload buffer

- Use different disks for cache and upload buffer. If using RAID, ensure that cache and upload buffer disks use separate RAID controllers at the hardware level.
- Add at least 2 different upload buffer disks.
- Use a RAID 0 striped raid configuration for cache + upload buffer devices to improve throughput. This is especially critical for the cache disk.

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 Tape Gateways, when the upload buffer reaches its capacity, your applications can continue to read from and write data to your storage volumes. However, the Tape Gateway does not write any of your volume data to its upload buffer and does not upload any of this data to AWS until Storage Gateway synchronizes the data stored locally with the copy of the data stored in AWS. This synchronization occurs when the volumes are in BOOTSTRAPPING status.

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.

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

Managing Bandwidth for Your Tape 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.AWSSStorageGatewayClient;
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 using the AWS Storage Gateway console

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. AWS updates the cloud services component as needed without causing disruption to deployed gateways.

Your deployed gateway appliances also require periodic updates. 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.

Gateway appliances receive updates on a recurring basis, and applying these 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 software appliance updates and restarts. You can minimize the chance of any disruption to your applications due to the gateway restart by increasing the timeouts of your iSCSI initiator. For more information about increasing iSCSI initiator timeouts for Windows and Linux, see [Customizing Your Windows iSCSI Settings](#) and [Customizing Your Linux iSCSI Settings](#).

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 gateway update mechanism (for example, SSM or hypervisor tools) may cause the gateway to malfunction.

You can configure your gateways to apply maintenance updates automatically during scheduled maintenance windows, or wait until you initiate them manually. 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. 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.

Note

Certain important 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 that software update notifications are sent, go to the [Managing an AWS account](#) page and update the alternate contact for 'operations'.

The procedures in the following sections describe how to manage gateway updates using the Storage Gateway console. To manage these settings programmatically using the API, see [UpdateMaintenanceStartTime](#) in the *Storage Gateway API Reference*.

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 [Modify the gateway maintenance window schedule](#).

If maintenance updates are turned off, the gateway will not apply these updates automatically, but you can still apply them manually using the Storage Gateway console, API, or CLI. Certain important updates will sometimes be applied during your configured maintenance window, regardless of this setting.

To turn maintenance updates on or off for a specific gateway:

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 new software updates according to the maintenance window schedule. Certain important updates will sometimes be applied during your configured maintenance window, regardless of the maintenance updates setting.

To modify the maintenance window schedule for a specific gateway:

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.

Performing Maintenance Tasks on the Local Console

You can perform the following maintenance tasks using the host's local console. Local console tasks can be performed on the VM host or the Amazon EC2 instance. Many of the tasks are common among the different hosts, but there are also some differences.

Performing Tasks on the VM Local Console

For a gateway deployed on-premises, you can perform the following maintenance tasks using the VM host's local console. These tasks are common to VMware, Hyper-V, and Linux Kernel-based Virtual Machine (KVM) hosts.

Topics

- [Logging in to the Local Console Using Default Credentials](#)
- [Setting the Local Console Password from the Storage Gateway Console](#)
- [Routing Your On-Premises Gateway Through a Proxy](#)
- [Configuring Your Gateway Network](#)
- [Testing Your Gateway Connection to the Internet](#)
- [Synchronizing Your Gateway VM Time](#)
- [Running Storage Gateway Commands on the Local Console](#)
- [Viewing Your Gateway System Resource Status](#)

- [Configuring Network Adapters for Your Gateway](#)

Logging in to the Local Console Using Default Credentials

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

1. 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 on the Local Console](#). 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.

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	See This Topic
Configure a SOCKS proxy for your gateway	Routing Your On-Premises Gateway Through a Proxy.
Configure your network	Configuring Your Gateway Network.
Test network connectivity	Testing Your Gateway Connection to the Internet.
Manage VM time	Synchronizing Your Gateway VM Time.
Run Storage Gateway console commands	Running Storage Gateway Commands on the Local Console.
View system resource check	Viewing Your Gateway System Resource Status.

To shut down the gateway, enter **0**.

To exit the configuration session, enter **X**.

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.

Note

The password can consist of any character on the keyboard and can be 1 to 512 characters long.

Routing Your On-Premises Gateway Through a Proxy

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

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 **SOCKS Proxy Configuration**.
3. 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	<p>Enter the corresponding numeral to select Configure SOCKS Proxy.</p> <p>You will need to supply a host name and port to complete configuration.</p>
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>

4. 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


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 **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

To Perform This Task	Do This
	<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>Enter the corresponding numeral to select Configure Static IP.</p> <p>You are prompted to type the following information to configure a static IP:</p> <ul style="list-style-type: none">• Network adapter name• IP address• Netmask• Default gateway address• Primary Domain Name Service (DNS) address• Secondary DNS address <div data-bbox="829 1209 1511 1619" style="border: 1px solid #f08080; border-radius: 10px; padding: 10px;"><p> Important</p><p>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>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"><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.

Synchronizing Your Gateway VM Time

After your gateway is deployed and running, in some scenarios the gateway VM's time can drift. For example, if there is a prolonged network outage and your hypervisor host and gateway do not get time updates, then the gateway VM's time will be different from the true time. When there is a time drift, a discrepancy occurs between the stated times when operations such as snapshots occur and the actual times that the operations occur.

For a gateway deployed on VMware ESXi, setting the hypervisor host time and synchronizing the VM time to the host is sufficient to avoid time drift. For more information, see [Synchronizing VM Time with Host Time](#).

For a gateway deployed on Microsoft Hyper-V, you should periodically check your VM's time. For more information, see [Synchronizing Your Gateway VM Time](#).

Running Storage Gateway Commands on the Local Console

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. <div data-bbox="834 619 1507 1024"><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>
ip	Show / manipulate routing, devices, and tunnels. <div data-bbox="834 1192 1507 1598"><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.

Command	Function
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.
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.

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**.

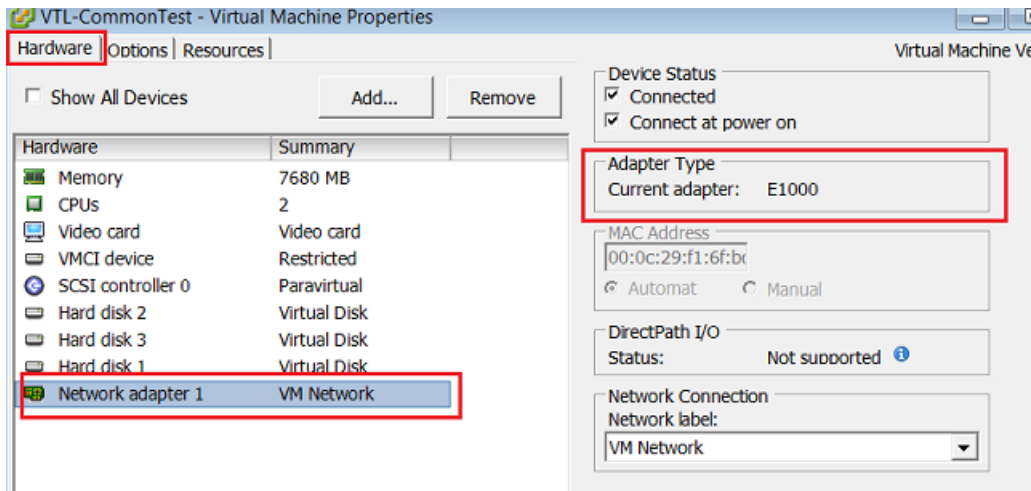
Following are the steps you take to configure your gateway to use the VMXNET3 adapter:

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.



- 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.

- Choose **Add** to open the Add Hardware wizard.
- Choose **Ethernet Adapter**, and then choose **Next**.
- In the Network Type wizard, select **VMXNET3** for **Adapter Type**, and then choose **Next**.
- In the Virtual Machine properties wizard, verify in the **Adapter Type** section that **Current Adapter** is set to **VMXNET3**, and then choose **OK**.
- In the VMware VSphere client, shut down your gateway.
- 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

- 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 Your Gateway for Multiple NICs in a VMware ESXi Host](#)

- [Configuring Your Gateway for Multiple NICs in Microsoft Hyper-V Host](#)

Performing Tasks on the Amazon EC2 Local Console

Some maintenance tasks require that you log in to the local console when running a gateway deployed on an Amazon EC2 instance. This section describes how to log in to the local console and perform maintenance tasks.

Topics

- [Logging In to Your Amazon EC2 Gateway Local Console](#)
- [Routing your gateway deployed on EC2 through an HTTP proxy](#)
- [Testing your gateway's network connectivity](#)
- [Viewing your gateway system resource status](#)
- [Running Storage Gateway commands on the local console](#)

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	See This Topic
Configure a SOCKS proxy for your gateway	Routing your gateway deployed on EC2 through an HTTP proxy
Test network connectivity	Testing your gateway's network connectivity

To Learn About This Task	See This Topic
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 your gateway's 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.
[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.
	<div data-bbox="836 1354 1510 1669" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px;"> <p> Note</p> <p>We recommend configuring network or IP settings using the Storage Gateway console or the dedicated local console menu option.</p> </div>
ip	Show / manipulate routing, devices, and tunnels.

Command	Function
	<p> Note</p> <p>We recommend configuring network or IP settings using the Storage Gateway console or the dedicated local console menu option.</p>
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`.

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
```

You can choose available VMs by Id.

```
[[root@localhost vms]# virsh list
 Id   Name           State
-----
 7    SGW_KVM        running

[[root@localhost vms]# virsh console 7
```

2. Use the following command to access the local console.

```
# virsh console VM_Id
```

```
[[root@localhost vms]# virsh console 7
Connected to domain SGW_KVM
Escape character is ^]

AWS Appliance

Login to change your network configuration and other settings.
localhost login: _
```

3. To get default credentials to log in to the local console, see [Logging in to the Local Console Using Default Credentials](#).
4. After you have logged in, you can activate and configure your gateway.

```
AWS Appliance Activation - Configuration

#####
## Currently connected network adapters:
##
## eth0: 10.0.3.32
#####

1: HTTP/SOCKS Proxy Configuration
2: Network Configuration
3: Test Network Connectivity
4: View System Resource Check (0 Errors)
5: System Time Management
6: License Information
7: Command Prompt

0: Get activation key

Press "x" to exit session

Enter command: _
```


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 is turned on.

Note

If your gateway VM is turned on, a green arrow icon appears with the VM icon, as shown in the following screenshot. If your gateway VM is not turned on, you can turn it on by choosing the green **Power On** icon on the **Toolbar** menu.



3. Choose the **Console** tab.



After a few moments, the VM is ready for you to log in.

Note

To release the cursor from the console window, press **Ctrl+Alt**.

```
AWS Storage Gateway

Login to change your network configuration and other gateway settings.

For more information, please see:
https://docs.aws.amazon.com/console/storagegateway/LocalConsole

localhost login: _
```

4. To log in using the default credentials, continue to the procedure [Logging in to the Local Console Using Default Credentials](#).


Access the Gateway Local Console with Microsoft Hyper-V

To access your gateway's local console (Microsoft Hyper-V)

1. In the **Virtual Machines** list of the Microsoft Hyper-V Manager, select your gateway VM.
2. Make sure that the gateway is turned on.

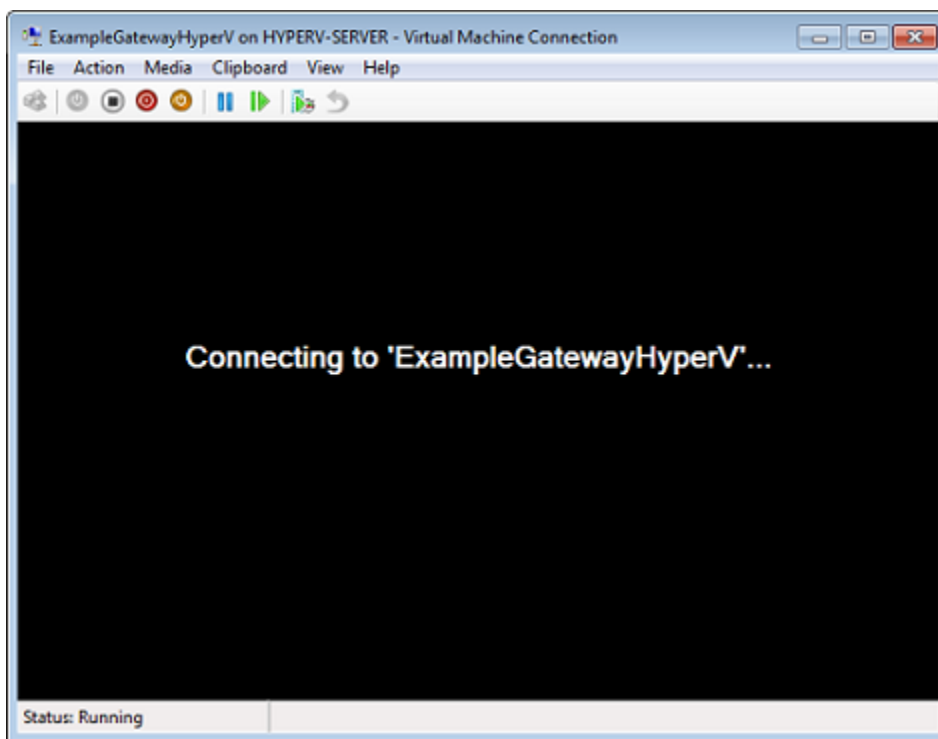
Note

If your gateway VM is turned on, **Running** is displayed as the **State** of the VM, as shown in the following screenshot. If your gateway VM is not turned on, you can turn it on by choosing **Start** in the **Actions** pane.



3. In the **Actions** pane, choose **Connect**.

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 VM is ready for you to log in.

```
AWS Storage Gateway

Login to change your network configuration and other gateway settings.

For more information, please see:
https://docs.aws.amazon.com/console/storagegateway/LocalConsole

localhost login: _
```

4. To log in using the default credentials, continue to the procedure [Logging in to the Local Console Using Default Credentials](#).

Configuring Network Adapters for Your Gateway

In this section you can find information about how configure multiple network adapters for your gateway.

Topics

- [Configuring Your Gateway for Multiple NICs in a VMware ESXi Host](#)
- [Configuring Your Gateway for Multiple NICs in Microsoft Hyper-V Host](#)

Configuring Your Gateway for Multiple NICs in 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 Your Gateway for Multiple NICs in 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. For instructions, see [To stop a Tape Gateway](#).
2. In the Microsoft Hyper-V Manager, select your gateway VM.

3. If the VM isn't turned off already, open the context (right-click) menu for your gateway and choose **Turn Off**.
4. In the client, open the context menu for your gateway VM and choose **Settings**.




5. In the **Settings** dialog box for the VM, for **Hardware**, choose **Add Hardware**.
6. In the **Add Hardware** pane, choose **Network Adapter**, and then choose **Add** to add a device.



7. Configure the network adapter, and then choose **Apply** to apply settings.

In the following example, **Virtual Network 2** is selected for the new adapter.



8. In the **Settings** dialog box, for **Hardware**, confirm that the second adapter was added, and then choose **OK**.
9. On the Storage Gateway console, turn on the gateway. For instructions, see [To start a Tape Gateway](#).
10. In the **Navigation** pane choose **Gateways**, then select the gateway to which you added the adapter. Confirm that the second IP address is listed in the **Details** tab.

Note

The example mounting commands provided on the info page for a file share in the Storage Gateway console will always include the IP address of the network adapter that was most recently added to the file share's associated gateway.

For information about local console tasks common to VMware, Hyper-V, and KVM hosts, see [Performing Tasks on the VM Local Console](#)

Deleting Your Gateway by Using the AWS Storage Gateway Console 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.

Note

When you delete a Tape Gateway, any tapes that are currently in the AVAILABLE status are also deleted, and any data on those tapes is lost. If you want to retain data from tapes that are being used by a gateway that you want to delete, you must archive the tapes before you delete the gateway. For more information, see [Archiving Virtual Tapes](#).

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 Tape Gateway Deployed on a VM

When you delete a gateway–virtual tape library (VTL), you perform additional cleanup steps before and after you delete the gateway. These additional steps help you remove resources you don't need so you don't continue to pay for them.

If the Tape Gateway you want to delete is deployed on a virtual machine (VM), we suggest that you take the following actions to clean up resources.

Important

Before you delete a Tape Gateway, you must cancel all tape retrieval operations and eject all retrieved tapes.

After you have deleted the Tape Gateway, you must remove any resources associated with the Tape Gateway that you don't need to avoid paying for those resources.

When you delete a Tape Gateway, you can encounter one of two scenarios.

- **The Tape Gateway is connected to AWS** – If the Tape Gateway is connected to AWS and you delete the gateway, the iSCSI targets associated with the gateway (that is, the virtual tape drives and media changer) will no longer be available.
- **The Tape Gateway is not connected to AWS** – If the Tape Gateway is not connected to AWS, for example if the underlying VM is turned off or your network is down, then you cannot delete the gateway. If you attempt to do so, after your environment is back up and running you might have a Tape Gateway running on-premises with available iSCSI targets. However, no Tape Gateway data will be uploaded to, or downloaded from, AWS.

If the Tape Gateway you want to delete is not functioning, you must first deactivate it before you delete it, as described following:

- To delete tapes that have the RETRIEVED status from the library, eject the tape using your backup software. For instructions, see [Archiving the Tape](#).

After deactivating the Tape Gateway and deleting tapes, you can delete the Tape Gateway. For instructions on how to delete a gateway, see [Deleting Your Gateway by Using the Storage Gateway Console](#).

If you have tapes archived, those tapes remain and you continue to pay for storage until you delete them. For instruction on how to delete tapes from a archive. see [Deleting Tapes](#).

Important

You are charged for a minimum of 90 days storage for virtual tapes in a archive. If you retrieve a virtual tape that has been stored in the archive for less than 90 days, you are still charged for 90 days storage.

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 Tape Gateway Deployed on Amazon EC2

If you deployed a Tape Gateway, we suggest that you take the following actions to delete your gateway and clean up its resources:

1. Delete all virtual tapes that you have retrieved to your Tape Gateway. For more information, see [Deleting Tapes](#).
2. Delete all virtual tapes from the tape library. For more information, see [Deleting Tapes](#).
3. Delete the Tape Gateway. For more information, see [Deleting Your Gateway by Using the Storage Gateway Console](#).

4. Terminate all Amazon EC2 instances, and delete all Amazon EBS volumes. For more information, see [Clean Up Your Instance and Volume](#) in the *Amazon EC2 User Guide*.
5. Delete all archived virtual tapes. For more information, see [Deleting Tapes](#).

 Important

You are charged for a minimum of 90 days storage for virtual tapes in the archive. If you retrieve a virtual tape that has been stored in the archive for less than 90 days, you are still charged for 90 days storage.

Performance

This section describes Storage Gateway performance.

Topics

- [Performance guidance for Tape Gateways](#)
- [Optimizing Gateway Performance](#)
- [Using VMware vSphere High Availability with Storage Gateway](#)

Performance guidance for Tape Gateways

In this section, you can find configuration guidance for provisioning hardware for your Tape Gateway VM. The Amazon EC2 instance sizes and types that are listed in the table are examples, and are provided for reference.

Configuration	Write Throughput Gbps	Read from Cache Throughput Gbps	Read from Amazon Web Services Cloud Throughput Gbps
Host Platform: Amazon EC2 instance— c5.4xlarge CPU: 16 vCPU RAM: 32 GB Root disk: 80 GB, io1 SSD, 4,000 IOPS Cache disk: striped RAID (2 x 500 GB, io1 EBS SSD, 25000 IOPS) Upload buffer disk: 450 GB, io1 SSD, 2000 IOPS Network bandwidth to cloud: 10 Gbps	2.3	4.0	2.2

Configuration	Write Throughput Gbps	Read from Cache Throughput Gbps	Read from Amazon Web Services Cloud Throughput Gbps
Host platform: Storage Gateway Hardware Appliance Cache disk: 2.5 TB Upload buffer disk: 2 TB Network bandwidth to cloud: 10 Gbps	2.3	8.8	3.8
Host platform: Amazon EC2instance— c5d.9xlarge CPU: 36 vCPU RAM: 72 GB Root disk: 80 GB, io1 SSD, 4,000 IOPS Cache disk: 900 GB NVMe disk Upload buffer disk: 900 GB NVMe disk Network bandwidth to cloud: 10 Gbps	5.2	11.6	5.2

Configuration	Write Throughput Gbps	Read from Cache Throughput Gbps	Read from Amazon Web Services Cloud Throughput Gbps
Host platform: Amazon EC2 instance— c5d.metal CPU: 96 vCPU RAM: 192 GB Root disk: 80 GB, io1 SSD, 4,000 IOPS Cache disk: striped RAID (2 x 900 GB NVMe disk) Upload buffer disk: 900 GB NVMe disk Network bandwidth to cloud: 10 Gbps	5.2	11.6	7.2

Note

This performance was achieved by using a 1 MB block size and ten tape drives simultaneously.

The EC2 configurations in the above table are only intended to be representative of the performance you might attain on your own physical servers with similar resources. For example, the EC2 configurations using a striped RAID were done through a special mechanism that is not generally supported by our gateway on EC2. To achieve similar performance, you should instead use a hardware RAID controller attached to the on-premise server running your gateway.

Your performance might vary based on your host platform configuration and network bandwidth.

To improve write and read throughput performance of your Tape Gateway, see [Optimize iSCSI Settings](#), [Use a Larger Block Size for Tape Drives](#), and [Optimize the Performance of Virtual Tape Drives in the Backup Software](#).

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 64 dedicated physical CPU cores
- For Tape 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

Note

For optimal gateway performance, you must provision at least 32 GiB of RAM.

- Disk 1, to be used as the gateway cache as follows:
 - Striped RAID (redundant array of independent disks) consisting of NVMe SSDs.
- Disk 2, to be used as the gateway upload buffer as follows:
 - Striped RAID consisting of NVMe SSDs.
- Disk 3, to be used as the gateway upload buffer as follows:
 - Striped RAID consisting of NVMe SSDs.
- 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 Tape 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. For more information about gateway metrics, see [Measuring Performance Between Your Tape Gateway and AWS](#).

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.

Network bandwidth to and from AWS defines the *theoretical maximum* average performance of your Tape Gateway during sustained workloads.

- The average rate at which you can write data to your Tape Gateway over long intervals will not exceed your upload bandwidth to AWS.
- The average rate at which you can read data from your Tape Gateway over long intervals will not exceed your download bandwidth to AWS.

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.

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.

Use a Larger Block Size for Tape Drives

For a Tape Gateway, the default block size for a tape drive is 64 KB. However, you can increase the block size up to 1 MB to improve I/O performance.

The block size that you choose depends on the maximum block size that your backup software supports. We recommend that you set the block size of the tape drives in your backup software to a size that is as large as possible. However, this block size must not be greater than the 1 MB maximum size that the gateway supports.

Tape Gateways negotiate the block size for virtual tape drives to automatically match what is set on the backup software. When you increase the block size on the backup software, we recommend that you also check the settings to ensure that the host initiator supports the new block size. For more information, see the documentation for your backup software. For more information about specific gateway performance guidance, see [Performance](#).

Optimize the Performance of Virtual Tape Drives in the Backup Software

Your backup software can back up data on up to 10 virtual tape drives on a Tape Gateway at the same time. We recommend that you configure backup jobs in your backup software to use at least 4 virtual tape drives simultaneously on the Tape Gateway. You can achieve better write throughput when the backup software is backing up data to more than one virtual tape at the same time.

As a general rule, you can achieve a higher maximum throughput by operating on (reading or writing from) more virtual tapes at the same time. By using more tape drives, you allow your gateway to service more requests concurrently, potentially improving performance.

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.
- Using as many tape drives concurrently as possible. iSCSI does not support queuing multiple requests for the same target, meaning that the more tape drives you use, the more requests that your gateway can service concurrently. This will allow you to more fully utilize the bandwidth between your gateway and initiator, increasing your gateway's apparent throughput.

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 more information about these metrics, see [Measuring Performance Between Your Tape Gateway and AWS](#).

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.

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 VMware HA with Storage Gateway, take the steps listed following.

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.

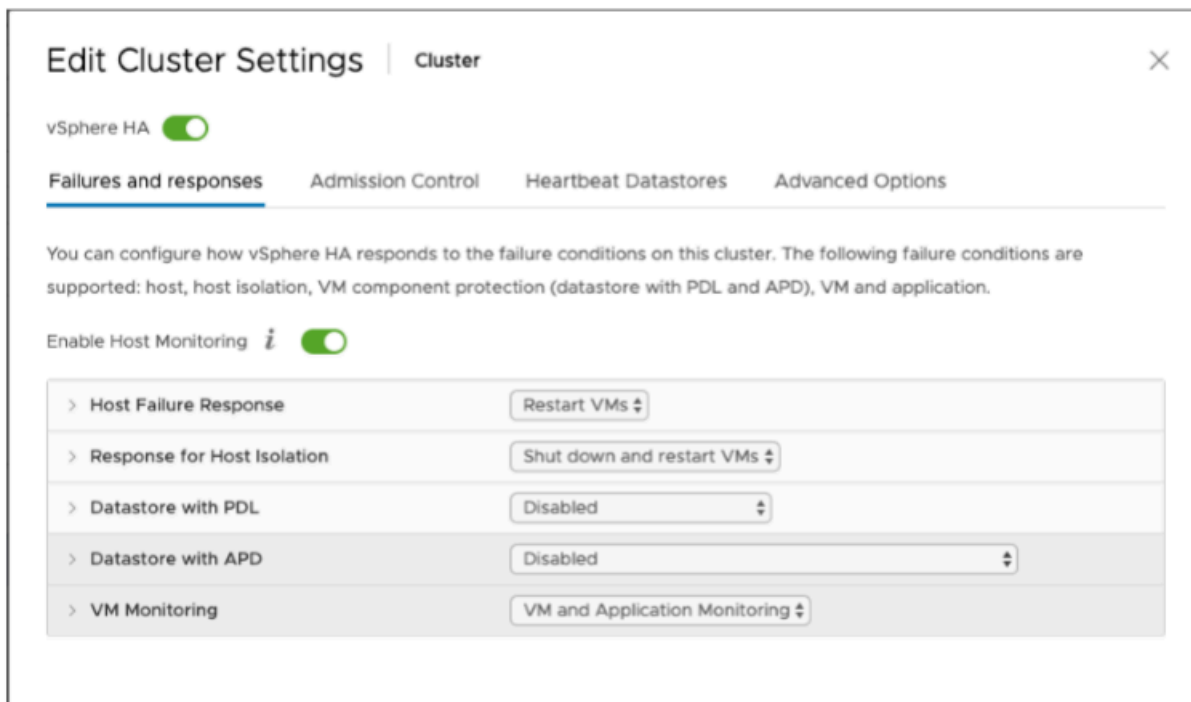
Next, configure your VMware cluster to work with Storage Gateway.

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 options as listed:

- **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**

For an example, see the following screenshot.



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 Tape 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.

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.

For override options, see the following screenshot.



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 Tape Gateway to AWS](#)
 - b. [Review settings and activate your Tape Gateway](#)
 - c. [Configure your Tape 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 Tape Gateway Health Logs with CloudWatch Log Groups](#).

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](#)
- [Logging and Monitoring in AWS Storage Gateway](#)
- [Compliance validation for AWS Storage Gateway](#)
- [Resilience in AWS Storage Gateway](#)
- [Infrastructure Security in AWS Storage Gateway](#)
- [AWS Security Best Practices](#)

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.
- 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 140-2 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 \(FIPS\) 140-2](#).

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?](#)

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 AWS 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 AWS 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 [How SCPs work](#) 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

IAM feature	AWS SGW support
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 AWS 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 AWS 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 AWS 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 AWS 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 AWS 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*.

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": "AIDAI15AUEPBH2M7JTNVC",
    "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/cloudtrailgatewayvt1"
        },
        "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": "AIDAI5AUEPBH2M7JTNVC",
      "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"  
  }  
}
```

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. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

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](#).
- Archive virtual tapes in S3 Glacier Flexible Retrieval. For more information, see [Archiving Virtual Tapes](#).

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.

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](#).

Troubleshooting your gateway

Following, you can find information about troubleshooting issues related to gateways, file shares, volumes, virtual tapes, and snapshots. The on-premises gateway troubleshooting information covers gateways deployed on both the VMware ESXi and Microsoft Hyper-V clients. The troubleshooting information for file shares applies to the File Gateway type. The troubleshooting information for volumes applies to the Volume Gateway type. The troubleshooting information for tapes applies to the Tape Gateway type. The troubleshooting information for gateway issues applies to using CloudWatch metrics. The troubleshooting information for high availability issues covers gateways running on VMware vSphere High Availability (HA) platform.

Topics

- [Troubleshooting: gateway status shows offline in the AWS Storage Gateway console](#)
- [Troubleshooting: internal error received during Storage Gateway activation](#)
- [Troubleshooting on-premises gateway issues](#)
- [Troubleshooting Microsoft Hyper-V setup](#)
- [Troubleshooting Amazon EC2 gateway issues](#)
- [Troubleshooting hardware appliance issues](#)
- [Troubleshooting virtual tape issues](#)
- [Troubleshooting high availability issues](#)
- [Best practices for recovering your data](#)

Troubleshooting: gateway status shows offline in the AWS Storage Gateway console

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.

For more information on troubleshooting a corrupted cache disk for a tape gateway, see [You need to recover a virtual tape from a malfunctioning cache disk](#).

Troubleshooting: internal error received during Storage 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 Routing Your On-Premises Gateway Through a Proxy.

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 Synchronizing Your Gateway VM 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.
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.	For instructions about removing a disk allocated as upload buffer space, see Removing Disks from Your Gateway .

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 Tape Gateway and AWS. Improving your internet connectivity helps to ensure that your upload buffer does not fill up.</p>

Issue	Action to Take
Throughput to or from your gateway drops to zero.	<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 Tape Gateway and AWS.</p>
You are having trouble importing (deploying) Storage Gateway on Microsoft Hyper-V.	See Troubleshooting Microsoft Hyper-V setup , which discusses some of the common issues of deploying a gateway on Microsoft Hyper-V.

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

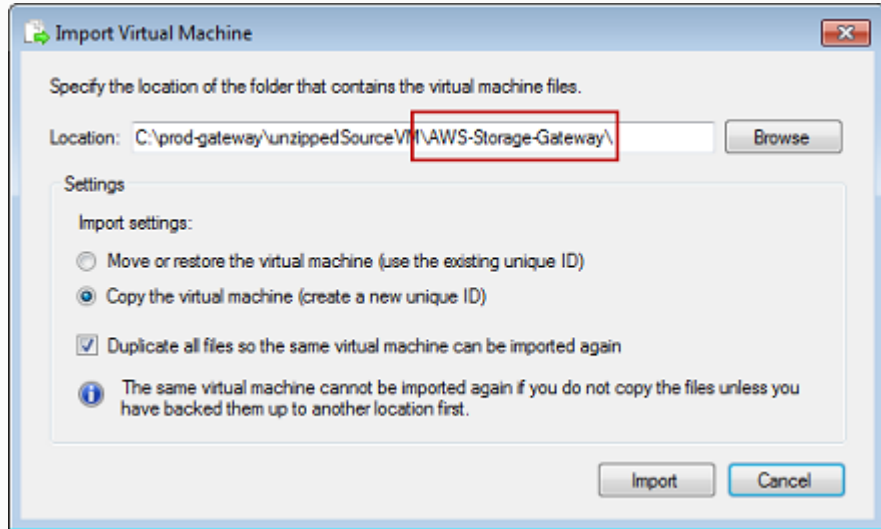
The following table lists typical issues that you might encounter when deploying Storage Gateway on the Microsoft Hyper-V platform.

Issue	Action to Take
You try to import a gateway and receive the error message: "Import failed. Unable to find virtual machine import file under location ...".	<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 , as the following example shows:

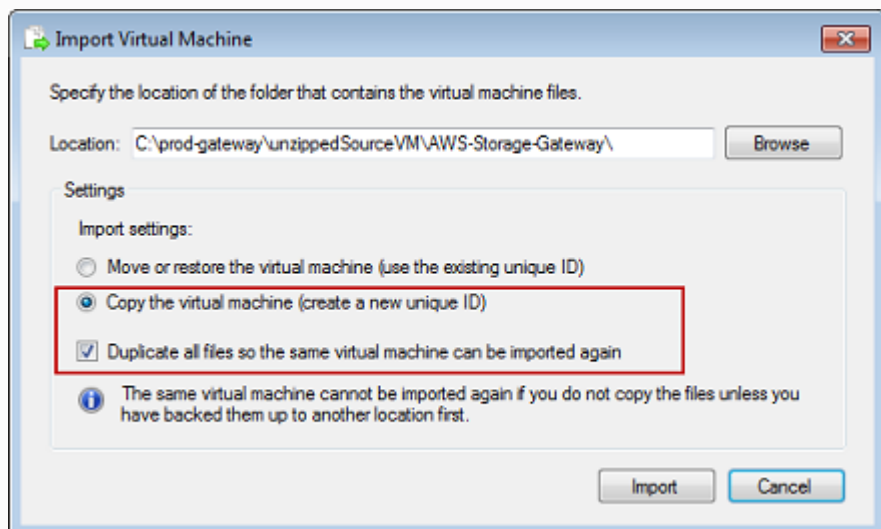
Issue

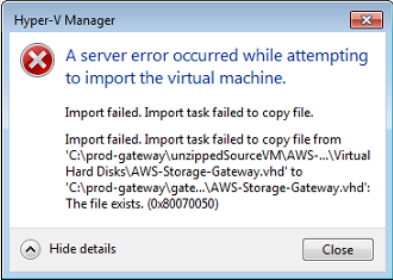
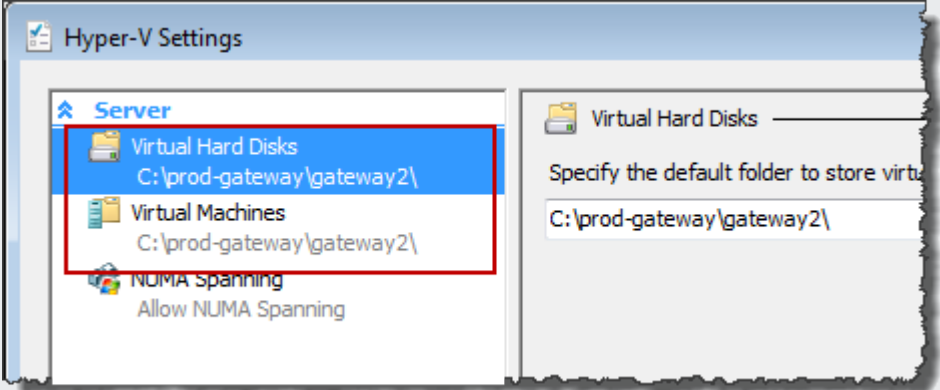


Action to Take




- 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. The following example shows the options that you must check if you plan on creating multiple gateways from one unzipped source files location.




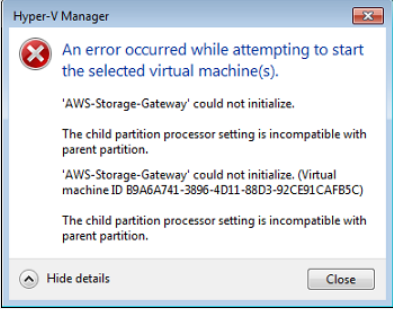
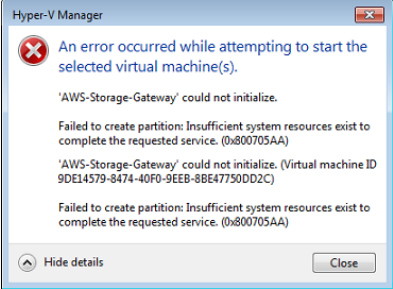
Issue	Action to Take
<p>You try to import a gateway and receive the error message: "Import failed. Import task failed to copy file."</p>  <p>The screenshot shows an error dialog box titled "Hyper-V Manager" with a red 'X' icon. The text reads: "A server error occurred while attempting to import the virtual machine. Import failed. Import task failed to copy file. Import failed. Import task failed to copy file from 'C:\prod-gateway\unzippedSourceVM\AWS-...Virtual Hard Disks\AWS-Storage-Gateway.vhd' to 'C:\prod-gateway\gate...\AWS-Storage-Gateway.vhd: 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 in the Hyper-V Settings dialog box.</p>  <p>The screenshot shows the "Hyper-V Settings" dialog box. The "Server" tab is selected, and the "Virtual Hard Disks" and "Virtual Machines" settings are highlighted with a red box. Both are set to "C:\prod-gateway\gateway2\". The "Virtual Hard Disks" section also shows a field for "Specify the default folder to store virtual hard disks" with the same path.</p>

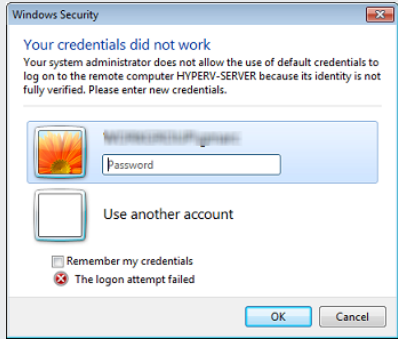
You try to import a gateway and receive an error message: "Import failed. Import failed because the virtual machine must have a new identifier. Select a new identifier and try the import again."



When you import the gateway make sure you select the **Copy the virtual machine** option and check the **Duplicate all files** option in the **Import Virtual Machine** dialog box to create a new unique ID for the VM. The following example shows the options in the **Import Virtual Machine** dialog box that you should use.



Issue	Action to Take
<p>You try to start a gateway VM and receive an error message "The child partition processor setting is incompatible with parent partition."</p> 	<p>This error is likely caused by a CPU discrepancy between the required CPUs for the gateway and the available CPUs on the host. Ensure that the VM CPU count is supported by the underlying hypervisor.</p> <p>For more information about the requirements for Storage Gateway, see Requirements.</p>
<p>You try to start a gateway VM and receive an error message "Failed to create partition: Insufficient resources exist to complete the requested service."</p> 	<p>This error is likely caused by a RAM discrepancy between the required RAM for the gateway and the available RAM on the host.</p> <p>For more information about the requirements for Storage Gateway, see Requirements.</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 Synchronizing Your Gateway VM Time.</p>

Issue	Action to Take
<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>
<p>You may notice poor network performance if you activate virtual machine queue (VMQ) on a Hyper-V host that's using a Broadcom network adapter.</p>	<p>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 activated.</p>

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 [Deploying an Amazon EC2 instance to host your Tape 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 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 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 [Deploying an Amazon EC2 instance to host your Tape 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 Tape 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-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 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

To find the serial number of the hardware appliance, go to the **Hardware appliance overview** page in the Storage Gateway console, as shown following.

Storage Gateway console hardware tab with appliance selected and details shown.

The screenshot shows the AWS Storage Gateway console interface. At the top, a green notification banner states "Successfully launched File Gateway on praksuji-bh". Below this, there are buttons for "Order appliance", "Quotes and orders", "Activate appliance", and an "Actions" dropdown menu. A search filter is present: "Filter by hardware appliance name, ID or launched gateway type." A table lists two hardware appliances:

	Hardware Appliance Name	Hardware Appliance ID	Model	Launched Gateway
<input checked="" type="checkbox"/>	praksuji-bh	v15loueix9yotyn5	Dell PowerEdge R640	File Gateway
<input type="checkbox"/>	praksuji-hw-pdx	wlyd0dgh6j7kg4no	Dell PowerEdge R640	File Gateway

Below the table, the "Details" tab is active, showing the following information:

Name	praksuji-bh	Vendor	Dell
ID	v15loueix9yotyn5	Model	Dell PowerEdge R640
Time Zone	GMT	Serial Number	5Q8Y0M2
		RAID Volume Manager	ZFS

Storage Gateway console hardware tab with appliance selected and details shown.

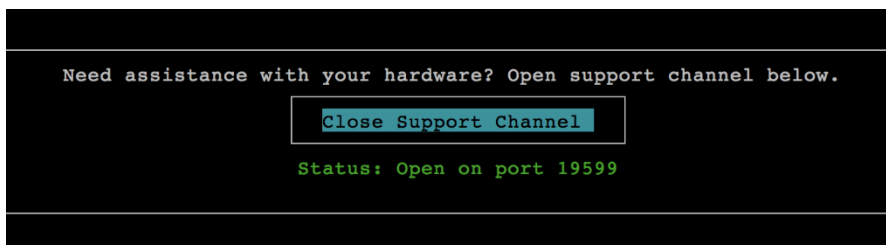
Where to obtain hardware appliance support

To contact the Storage Gateway Hardware Appliance support, 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** as shown following.
hardware appliance console with support channel status shown.



hardware appliance console with support channel status shown.

The assigned port number should appear within 30 seconds, if there are no network connectivity or firewall issues.

3. Note the port number and provide it to AWS Support.

Troubleshooting virtual tape issues

You can find information following about actions to take if you experience unexpected issues with your virtual tapes.

Topics

- [Recovering a Virtual Tape From An Unrecoverable Gateway](#)
- [Troubleshooting Irrecoverable Tapes](#)
- [High Availability Health Notifications](#)

Recovering a Virtual Tape From An Unrecoverable Gateway

Although it is rare, your Tape Gateway might encounter an unrecoverable failure. Such a failure can occur in your hypervisor host, the gateway itself, or the cache disks. If a failure occurs, you can recover your tapes by following the troubleshooting instructions in this section.

Topics

- [You Need to Recover a Virtual Tape from a Malfunctioning Tape Gateway](#)
- [You Need to Recover a Virtual Tape from a Malfunctioning Cache Disk](#)

You Need to Recover a Virtual Tape from a Malfunctioning Tape Gateway

If your Tape Gateway or the hypervisor host encounters an unrecoverable failure, you can recover any data that has already been uploaded to AWS to another Tape Gateway.

Note that the data written to a tape might not be completely uploaded until that tape has been successfully archived into VTS. The data on tapes recovered to another gateway in this manner may be incomplete or empty. We recommend performing an inventory on all recovered tapes to ensure they contain the expected content.

To recover a tape to another Tape Gateway

1. Identify an existing functioning Tape Gateway to serve as your recovery target gateway. If you don't have a Tape Gateway to recover your tapes to, create a new Tape Gateway. For information about how to create a gateway, see [Creating a Gateway](#).
2. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
3. In the navigation pane, choose **Gateways**, and then choose the Tape Gateway you want to recover tapes from.
4. Choose the **Details** tab. A tape recovery message is displayed in the tab.
5. Choose **Create recovery tapes** to deactivate the gateway.
6. In the dialog box that appears, choose **Disable gateway**.

This process permanently halts normal function of your Tape Gateway and exposes any available recovery points. For instructions, see [Deactivating your Tape Gateway](#).

7. From the tapes that the deactivated gateway displays, choose the virtual tape and the recovery point you want to recover. A virtual tape can have multiple recovery points.
8. To begin recovering any tapes you need to the target Tape Gateway, choose **Create recovery tape**.
9. In the **Create recovery tape** dialog box, verify the barcode of the virtual tape you want to recover.
10. For **Gateway**, choose the Tape Gateway you want to recover the virtual tape to.
11. Choose **Create recovery tape**.
12. Delete the failed Tape Gateway so you don't get charged. For instructions, see [Deleting Your Gateway by Using the AWS Storage Gateway Console and Removing Associated Resources](#).

Storage Gateway moves the tape from the failed Tape Gateway to the Tape Gateway you specified. The Tape Gateway marks the tape status as RECOVERED.

You Need to Recover a Virtual Tape from a Malfunctioning Cache Disk

If your cache disk encounters an error, the gateway prevents read and write operations on virtual tapes in the gateway. For example, an error can occur when a disk is corrupted or removed from the gateway. The Storage Gateway console displays a message about the error.

In the error message, Storage Gateway prompts you to take one of two actions that can recover your tapes:

- **Shut Down and Re-Add Disks** – Take this approach if the disk has intact data and has been removed. For example, if the error occurred because a disk was removed from your host by accident but the disk and the data is intact, you can re-add the disk. To do this, see the procedure later in this topic.
- **Reset Cache Disk** – Take this approach if the cache disk is corrupted or not accessible. If the disk error causes the cache disk to be inaccessible, unusable, or corrupted, you can reset the disk. If you reset the cache disk, tapes that have clean data (that is, tapes for which data in the cache disk and Amazon S3 are synchronized) will continue to be available for you to use. However, tapes that have data that is not synchronized with Amazon S3 are automatically recovered. The status of these tapes is set to RECOVERED, but the tapes will be read-only. For information about how to remove a disk from your host, see [Determining the size of upload buffer to allocate](#).

Important

If the cache disk you are resetting contains data that has not been uploaded to Amazon S3 yet, that data can be lost. After you reset cache disks, no configured cache disks will be left in the gateway, so you must configure at least one new cache disk for your gateway to function properly.

To reset the cache disk, see the procedure later in this topic.

To shut down and re-add a disk

1. Shut down the gateway. For information about how to shut down a gateway, see [Shutting Down Your Gateway VM](#).
2. Add the disk back to your host, and make sure the disk node number of the disk has not changed. For information about how to add a disk, see [Determining the size of upload buffer to allocate](#).
3. Restart the gateway. For information about how to restart a gateway, see [Shutting Down Your Gateway VM](#).

After the gateway restarts, you can verify the status of the cache disks. The status of a disk can be one of the following:

- **present** – The disk is available to use.
- **missing** – The disk is no longer connected to the gateway.
- **mismatch** – The disk node is occupied by a disk that has incorrect metadata, or the disk content is corrupted.

To reset and reconfigure a cache disk

1. In the **A disk error has occurred** error message illustrated preceding, choose **Reset Cache Disk**.
2. On the **Configure gateway** page, configure the disk for cache storage. For information about how to do so, see [Configure your Tape Gateway](#).
3. After you have configured cache storage, shut down and restart the gateway as described in the previous procedure.

The gateway should recover after the restart. You can then verify the status of the cache disk.

To verify the status of a cache disk

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose **Gateways**, and then choose your gateway.
3. For **Actions**, choose **Configure Local Storage** to display the **Configure Local Storage** dialog box. This dialog box shows all local disks in the gateway.

The cache disk node status is displayed next to the disk.

Note

If you don't complete the recovery process, the gateway displays a banner that prompts you to configure local storage.

Troubleshooting Irrecoverable Tapes

If your virtual tape fails unexpectedly, Storage Gateway sets the status of the failed virtual tape to IRRECOVERABLE. The action you take depends on the circumstances. You can find information following on some issues you might find, and how to troubleshoot them.

You Need to Recover Data From an IRRECOVERABLE Tape

If you have a virtual tape with the status IRRECOVERABLE, and you need to work with it, try one of the following:

- Activate a new Tape Gateway if you don't have one activated. For more information, see [Creating a Gateway](#).
- Deactivate the Tape Gateway that contains the irrecoverable tape, and recover the tape from a recovery point to the new Tape Gateway. For more information, see [You Need to Recover a Virtual Tape from a Malfunctioning Tape Gateway](#).

Note

You have to reconfigure your iSCSI initiator and backup application to use the new Tape Gateway. For more information, see [Connecting Your VTL Devices](#).

You Don't Need an IRRECOVERABLE Tape That Isn't Archived

If you have a virtual tape with the status IRRECOVERABLE, you don't need it, and the tape has never been archived, you should delete the tape. For more information, see [Deleting Tapes](#).

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.

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 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 tape](#)
- [Recovering your data from a malfunctioning cache disk](#)

- [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 tapes setups, when your gateway becomes reachable, your tapes 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 Tape Status](#).
- 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 Tape Gateway or the hypervisor host encounters an unrecoverable failure, you can use the following steps to recover the tapes from the malfunctioning Tape Gateway to another Tape Gateway:

1. Identify the Tape Gateway that you want to use as the recovery target, or create a new one.
2. Deactivate the malfunctioning gateway.
3. Create recovery tapes for each tape that you want to recover and specify the target Tape Gateway.
4. Delete the malfunctioning Tape Gateway.

For detailed information on how to recover the tapes from a malfunctioning Tape Gateway to another Tape Gateway, see [You Need to Recover a Virtual Tape from a Malfunctioning Tape Gateway](#).

Recovering your data from an irrecoverable tape

If your tape encounters a failure and the status of the tape is `IRRECOVERABLE`, we recommend you use one of the following options to recover your data or resolve the failure depending on your situation:

- If you need the data on the irrecoverable tape, you can recover the tape to a new gateway.
- If you don't need the data on the tape, and the tape has never been archived, you can simply delete the tape from your Tape Gateway.

For detailed information about how to recover your data or resolve the failure if your tape is `IRRECOVERABLE`, see [Troubleshooting Irrecoverable Tapes](#).

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.

For detailed information, see [You Need to Recover a Virtual Tape from a Malfunctioning Cache Disk](#).

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 Tape Gateway in an inaccessible data center

1. Create and activate a new Tape Gateway on an Amazon EC2 host. For more information, see [Deploying an Amazon EC2 instance to host your Tape Gateway](#).

2. Recover the tapes from the source gateway in the data center to the new gateway you created on Amazon EC2 For more information, see [Recovering a Virtual Tape From An Unrecoverable Gateway](#).

Your tapes should be covered to the new Amazon EC2 gateway.

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

- [Host Setup](#)
- [Tape Gateway](#)
- [Getting an activation key for your gateway](#)
- [Connecting iSCSI Initiators](#)
- [Using AWS Direct Connect with Storage Gateway](#)
- [Port Requirements](#)
- [Connecting to Your Gateway](#)
- [Understanding Storage Gateway Resources and Resource IDs](#)
- [Tagging Storage Gateway Resources](#)
- [Working with Open-Source Components for AWS Storage Gateway](#)
- [AWS Storage Gateway quotas](#)

Host Setup

Topics

- [Configuring VMware for Storage Gateway](#)
- [Synchronizing Your Gateway VM Time](#)
- [Deploying an Amazon EC2 instance to host your Tape Gateway](#)
- [Deploy an Amazon EC2 with Default Settings](#)
- [Modify Amazon EC2 instance metadata options](#)

Configuring VMware for Storage Gateway

When configuring VMware for Storage Gateway, make sure to synchronize your VM time with your host time, configure VM to use paravirtualized disk controllers when provisioning storage and provide protection from failures in the infrastructure layer supporting a gateway VM.

Topics

- [Synchronizing VM Time with Host Time](#)
- [Configuring the AWS Storage Gateway VM to Use Paravirtualized Disk Controllers](#)
- [Using Storage Gateway with VMware High Availability](#)

Synchronizing VM Time with 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, open the context (right-click) menu for your gateway VM, and choose **Edit Settings**.

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



- b. Choose the **Options** tab, and choose **VMware Tools** in the options list.
- c. Check the **Synchronize guest time with host** option, and then choose **OK**.

The VM synchronizes its time with the host.

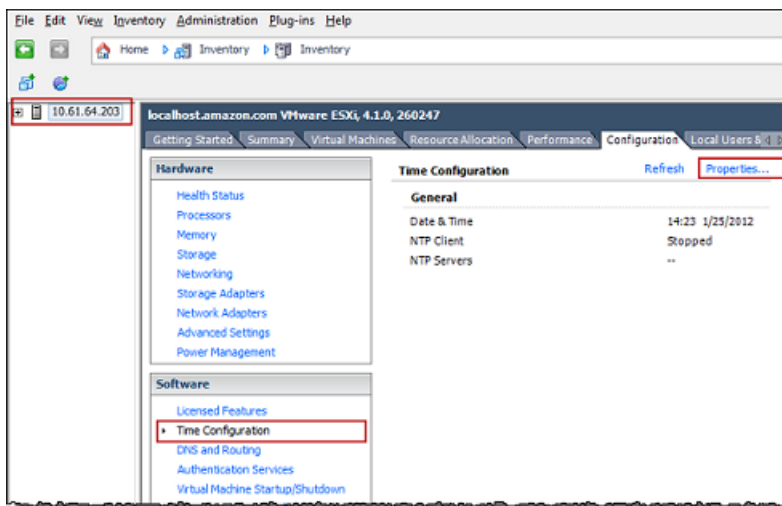


2. Configure the host time.


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 pane, 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. In the **Date and Time** panel, set the date and time.




- 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 pane.




- 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 shown in the following example.



- iv. In the **NTP Daemon (ntpd) Options** dialog box, choose **General** in the left pane.
- v. In the **Service Commands** pane, 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 the AWS Storage Gateway VM to Use Paravirtualized Disk Controllers


In this task, you set the iSCSI controller so that the VM uses paravirtualization. *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.

Note


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

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**.



Using Storage Gateway with VMware High Availability

VMware High Availability (HA) is a component of vSphere that can provide protection from failures in the infrastructure layer supporting a gateway VM. VMware HA does this by using multiple hosts configured as a cluster so that if a host running a gateway VM fails, the gateway VM can

be restarted automatically on another host within the cluster. For more information about VMware HA, see [VMware HA: Concepts and Best Practices](#) on the VMware website.

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.

Synchronizing Your Gateway VM Time

For a gateway deployed on VMware ESXi, setting the hypervisor host time and synchronizing the VM time to the host is sufficient to avoid time drift. For more information, see [Synchronizing VM Time with Host Time](#). For a gateway deployed on Microsoft Hyper-V, you should periodically check your VM's time using the procedure described following.

To view and synchronize the time of a hypervisor gateway VM to a Network Time Protocol (NTP) server

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 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, enter **4** for **System Time Management**.

```
AWS Storage Gateway Configuration
#####
## Currently connected network adapters:
##
## eth0: 10.0.0.45
#####

1: SOCKS Proxy Configuration
2: Network Configuration
3: Test Network Connectivity
4: System Time Management
5: Gateway Console
6: View System Resource Check (0 Errors)

0: Stop AWS Storage Gateway

Press "x" to exit session

Enter command: _
```

3. On the **System Time Management** menu, enter **1** for **View and Synchronize System Time**.

```
System Time Management

1: View and Synchronize System Time

Press "x" to exit

Enter command: _
```

4. If the result indicates that you should synchronize your VM's time to the NTP time, enter **y**. Otherwise, enter **n**.

If you enter **y** to synchronize, the synchronization might take a few moments.

The following screenshot shows a VM that doesn't require time synchronization.


```
System Time Management
1: View and Synchronize System Time
Press "x" to exit
Enter command: 1
Current System Time: Sat Aug 22 00:33:41 UTC 2015
Determining current NTP time (this may take a few seconds ...)
Your Storage Gateway VM system time differs from NTP time
by 0.217617 seconds
A sync is recommended if the time differs by more than 60 seconds
Do you want to sync Storage Gateway VM system time with
NTP time? [y/n]: _
```

The following screenshot shows a VM that does require time synchronization.

```
System Time Management
1: View and Synchronize System Time
Press "x" to exit
Enter command: 1
Current System Time: Sat Aug 22 00:33:41 UTC 2015
Determining current NTP time (this may take a few seconds ...)
Your Storage Gateway VM system time differs from NTP time
by 61.217617 seconds
A sync is recommended if the time differs by more than 60 seconds
Do you want to sync Storage Gateway VM system time with
NTP time? [y/n]: _
```

Deploying an Amazon EC2 instance to host your Tape Gateway

You can deploy and activate a Tape 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.

Tape 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 Tape Gateway

1. Start setting up a new gateway using the Storage Gateway console. For instructions, see [Set up a Tape 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 Tape Gateway.
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 Tape 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 Tape 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**
- Tape 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 Tape Gateway.

You can determine the AMI ID to use for launching a Tape 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 Tape 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/VTL/latest`. 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/VTL/
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/VTL/
latest",
    "Name": "/aws/service/storagegateway/ami/VTL/latest",
    "Value": "ami-123c45dd67d891000"
  }
}
```

Deploy an Amazon EC2 with Default Settings

This topic lists the steps to deploy an Amazon EC2 host using the default specifications.

You can deploy and activate a Tape 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 Tape 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			
Delete on termination	Yes	Yes	Yes			
Encrypted	No	No	No			
Throughput	125	125	125			

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.

Tape Gateway

Topics

- [Removing Disks from Your Gateway](#)
- [Adding and Removing Amazon EBS Volumes for Your Gateway Hosted on Amazon EC2](#)
- [Working with VTL Devices](#)
- [Working With Tapes](#)

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

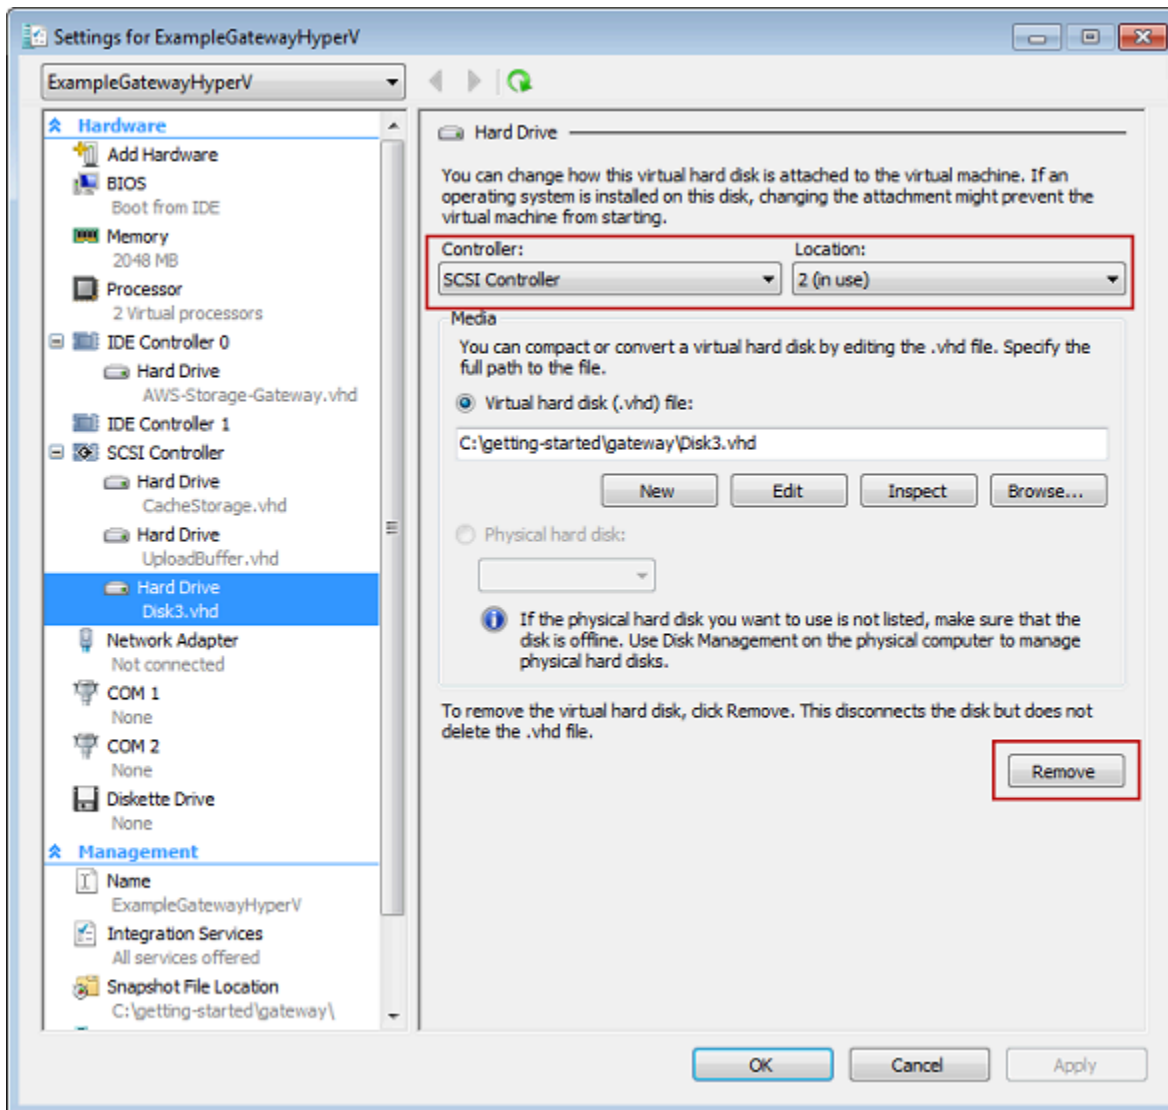
Using the following procedure, you can remove a disk from your gateway hosted on a Microsoft Hyper-V hypervisor.

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.

Adding and Removing Amazon EBS Volumes for Your Gateway Hosted on Amazon EC2

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. If you remove an Amazon EBS volume that is allocated as cache storage from a Tape Gateway, virtual tapes on the gateway will have the IRRECOVERABLE status, and you risk data loss. For more information on the IRRECOVERABLE status, see [Understanding Tape Status Information in a VTL](#).

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.

Working with VTL Devices

Your Tape Gateway setup provides the following SCSI devices, which you select when activating your gateway.

Topics

- [Selecting a Medium Changer After Gateway Activation](#)
- [Updating the Device Driver for Your Medium Changer](#)
- [Displaying Barcodes for Tapes in Microsoft System Center DPM](#)

For medium changers, AWS Storage Gateway works with the following:

- AWS-Gateway-VTL – This device is provided with the gateway.
- STK-L700 – This device emulation is provided with the gateway.

When activating your Tape Gateway, you select your backup application from the list and Storage Gateway uses the appropriate medium changer. If your backup application is not listed, you choose **Other** and then choose the medium changer that works with backup application.

The type of medium changer you choose depends on the backup application you plan to use. The following table lists third-party backup applications that have been tested and found to be compatible with Tape Gateways. This table includes the medium changer type recommended for each backup application.

Backup Application	Medium Changer Type
Arcserve Backup	AWS-Gateway-VTL
Bacula Enterprise V10.x	AWS-Gateway-VTL or STK-L700
Commvault V11	STK-L700
Dell EMC NetWorker 19.5	AWS-Gateway-VTL
IBM Spectrum Protect v8.1.10	IBM-03584L32-0402
Micro Focus (HPE) Data Protector 9 or 11.x	AWS-Gateway-VTL
Microsoft System Center 2012 R2 or 2016 Data Protection Manager	STK-L700
NovaStor DataCenter/Network 6.4 or 7.1	STK-L700
Quest NetVault Backup 12.4 or 13.x	STK-L700
Veeam Backup & Replication 11A	AWS-Gateway-VTL
Veritas Backup Exec 2014 or 15 or 16 or 20 or 22.x	AWS-Gateway-VTL

Backup Application	Medium Changer Type
Veritas Backup Exec 2012	STK-L700
<div data-bbox="147 304 808 527"><p>Note</p><p>Veritas has ended support for Backup Exec 2012.</p></div>	
Veritas NetBackup Version 7.x or 8.x	AWS-Gateway-VTL

Important

We highly recommend that you choose the medium changer that's listed for your backup application. Other medium changers might not function properly. You can choose a different medium changer after the gateway is activated. For more information, see [Selecting a Medium Changer After Gateway Activation](#).

For tape drives, Storage Gateway works with the following:

- IBM-ULT3580-TD5—This device emulation is provided with the gateway.

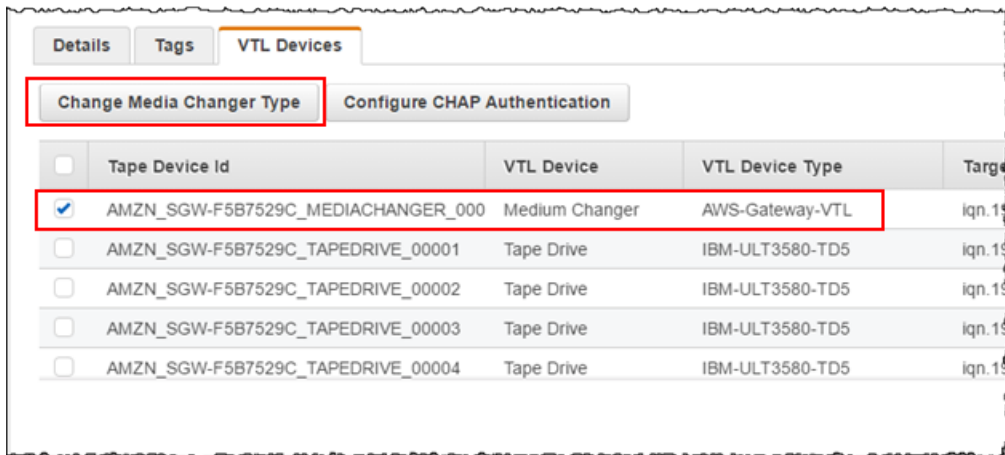
Selecting a Medium Changer After Gateway Activation

After your gateway is activated, you can choose to select a different medium changer type.

To select a different medium changer type after gateway activation

1. Stop any related jobs that are running in your backup software.
2. On the Windows server, open the iSCSI initiator properties window.
3. Choose the **Targets** tab to display the discovered targets.
4. On the Discovered targets pane, choose the medium changer you want to change, choose **Disconnect**, and then choose **OK**.
5. On the Storage Gateway console, choose **Gateways** from the navigation pane, and then choose the gateway whose medium changer you want to change.

- Choose the **VTL Devices** tab, select the medium changer you want to change, and then choose **Change Media Changer**.



- In the Change Media Changer Type dialog box that appears, select the media changer you want from the drop-down list box and then choose **Save**.

Updating the Device Driver for Your Medium Changer

- Open Device Manager on your Windows server, and expand the **Medium Changer devices** tree.
- Open the context (right-click) menu for **Unknown Medium Changer**, and choose **Update Driver Software** to open the **Update Driver Software-unknown Medium Changer** window.
- In the **How do you want to search for driver software?** section, choose **Browse my computer for driver software**.
- Choose **Let me pick from a list of device drivers on my computer**.

Note

We recommend using the Sony TSL-A500C Autoloader driver with the Veeam Backup & Replication 11A and Microsoft System Center Data Protection Manager backup software. This Sony driver has been tested with these types of backup software up to and including Windows Server 2019.

- In the **Select the device driver you want to install for this hardware** section, clear the **Show compatible hardware** check box, choose **Sony** in the **Manufacturer** list, choose **Sony - TSL-A500C Autoloader** in the **Model** list, and then choose **Next**.
- In the warning box that appears, choose **Yes**. If the driver is successfully installed, close the **Update drive software** window.

Displaying Barcodes for Tapes in Microsoft System Center DPM

If you use the media changer driver for Sony TSL-A500C Autoloader, Microsoft System Center Data Protection Manager doesn't automatically display barcodes for virtual tapes created in Storage Gateway. To display barcodes correctly for your tapes, change the media changer driver to Sun/StorageTek Library.

To display barcodes

1. Ensure that all backup jobs have completed and that there are no tasks pending or in progress.
2. Eject and move the tapes to offline storage (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive) and exit the DPM Administrator console. For information about how to eject a tape in DPM, see [Archiving a Tape by Using DPM](#).
3. In **Administrative Tools**, choose **Services** and open the context (right-click) menu for **DPM Service** in the **Detail** pane, and then choose **Properties**.
4. On the **General** tab, ensure that the **Startup type** is set to **Automatic** and choose **Stop** to stop the DPM service.
5. Get the StorageTek drivers from [Microsoft Update Catalog](#) on the Microsoft website.

Note

Take note of the different drivers for the different sizes.

For **Size 18K**, choose **x86 drivers**.

For **Size 19K**, choose **x64 drivers**.

6. On your Windows server, open Device Manager, and expand the **Medium Changer Devices** tree.
7. Open the context (right-click) menu for **Unknown Medium Changer**, and choose **Update Driver Software** to open the **Update Driver Software-unknown Medium Changer** window.
8. Browse to the path of the new driver location and install. The driver appears as **Sun/StorageTek Library**. The tape drives remain as an IBM ULT3580-TD5 SCSI sequential device.
9. Reboot the DPM server.
10. In the Storage Gateway console, create new tapes.

11. Open the DPM Administrator console, choose **Management**, then choose **Rescan for new tape libraries** . You should see the **Sun/StorageTek library**.
12. Choose the library and choose **Inventory**.
13. Choose **Add Tapes** to add the new tapes into DPM. The new tapes should now display their barcodes.

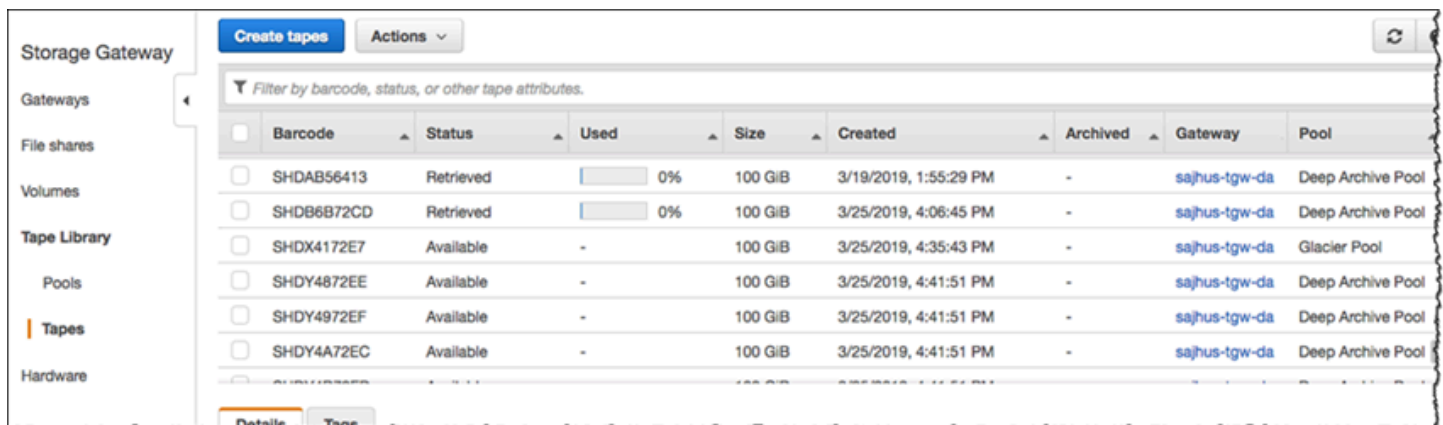
Working With Tapes

Storage Gateway provides one virtual tape library (VTL) for each Tape Gateway you activate. Initially, the library contains no tapes, but you can create tapes whenever you need to. Your application can read and write to any tapes available on your Tape Gateway. A tape's status must be AVAILABLE for you to write to the tape. These tapes are backed by Amazon Simple Storage Service (Amazon S3)—that is, when you write to these tapes, the Tape Gateway stores data in Amazon S3. For more information, see [Understanding Tape Status Information in a VTL](#).

Topics

- [Archiving Tapes](#)
- [Canceling Tape Archival](#)

The tape library shows tapes in your Tape Gateway. The library shows the tape barcode, status, and size, amount of the tape used, and the gateway the tape is associated with.



	Barcode	Status	Used	Size	Created	Archived	Gateway	Pool
<input type="checkbox"/>	SHDAB56413	Retrieved	0%	100 GiB	3/19/2019, 1:55:29 PM	-	sajhus-tgw-da	Deep Archive Pool
<input type="checkbox"/>	SHDB6B72CD	Retrieved	0%	100 GiB	3/25/2019, 4:06:45 PM	-	sajhus-tgw-da	Deep Archive Pool
<input type="checkbox"/>	SHDX4172E7	Available	-	100 GiB	3/25/2019, 4:35:43 PM	-	sajhus-tgw-da	Glacier Pool
<input type="checkbox"/>	SHDY4872EE	Available	-	100 GiB	3/25/2019, 4:41:51 PM	-	sajhus-tgw-da	Deep Archive Pool
<input type="checkbox"/>	SHDY4972EF	Available	-	100 GiB	3/25/2019, 4:41:51 PM	-	sajhus-tgw-da	Deep Archive Pool
<input type="checkbox"/>	SHDY4A72EC	Available	-	100 GiB	3/25/2019, 4:41:51 PM	-	sajhus-tgw-da	Deep Archive Pool

When you have a large number of tapes in the library, the console supports searching for tapes by barcode, by status, or by both. When you search by barcode, you can filter by status and gateway.

To search by barcode, status, and gateway

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the navigation pane, choose **Tapes**, and then type a value in the search box. The value can be the barcode, status, or gateway. By default, Storage Gateway searches for all virtual tapes. However, you can also filter your search by status.

If you filter for status, tapes that match your criteria appear in the library in the Storage Gateway console.

If you filter for gateway, tapes that are associated with that gateway appear in the library in the Storage Gateway console.

Note

By default, Storage Gateway displays all tapes regardless of status.

Archiving Tapes

You can archive the virtual tapes that are in your Tape Gateway. When you archive a tape, Storage Gateway moves the tape to the archive.

To archive a tape, you use your backup software. Tape archival process consists of three stages, seen as the tape statuses **IN TRANSIT TO VTS**, **ARCHIVING**, and **ARCHIVED**:

- To archive a tape, use the command provided by your backup application. When the archival process begins the tape status changes to **IN TRANSIT TO VTS** and the tape is no longer accessible to your backup application. In this stage, your Tape Gateway is uploading data to AWS. If needed, you can cancel the archival in progress. For more information about canceling archival, see [Canceling Tape Archival](#).

Note

The steps for archiving a tape depend on your backup application. For detailed instructions, see the documentation for your backup application.

- After the data upload to AWS completes, the tape status changes to **ARCHIVING** and Storage Gateway begins moving the tape to the archive. You cannot cancel the archival process at this point.
- After the tape is moved to the archive, its status changes to **ARCHIVED** and you can retrieve the tape to any of your gateways. For more information about tape retrieval, see [Retrieving Archived Tapes](#).

The steps involved in archiving a tape depend on your backup software. For instructions on how to archive a tape by using Symantec NetBackup software, see [Archiving the Tape](#).

Canceling Tape Archival

After you start archiving a tape, you might decide you need your tape back. For example, you might want to cancel the archival process, get the tape back because the archival process is taking too long, or read data from the tape. A tape that is being archived goes through three statuses, as shown following:

- IN TRANSIT TO VTS: Your Tape Gateway is uploading data to AWS.
- ARCHIVING: Data upload is complete and the Tape Gateway is moving the tape to the archive.
- ARCHIVED: The tape is moved and the archive and is available for retrieval.

You can cancel archival only when the tape's status is IN TRANSIT TO VTS. Depending on factors such as upload bandwidth and the amount of data being uploaded, this status might or might not be visible in the Storage Gateway console. To cancel a tape archival, use the [CancelRetrieval](#) action in the API reference.

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  
}
```

```
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 Your VTL Devices to a Windows client](#)
- [Connecting Your Volumes or VTL Devices 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 Your VTL Devices to a Windows client

A Tape Gateway exposes several tape drives and a media changer, referred to collectively as VTL devices, as iSCSI targets. For more information, see [Requirements](#).

Note

You connect only one application to each iSCSI target.

The following diagram highlights the iSCSI target in the larger picture of the Storage Gateway architecture. For more information on Storage Gateway architecture, see [How Tape Gateway works \(architecture\)](#).



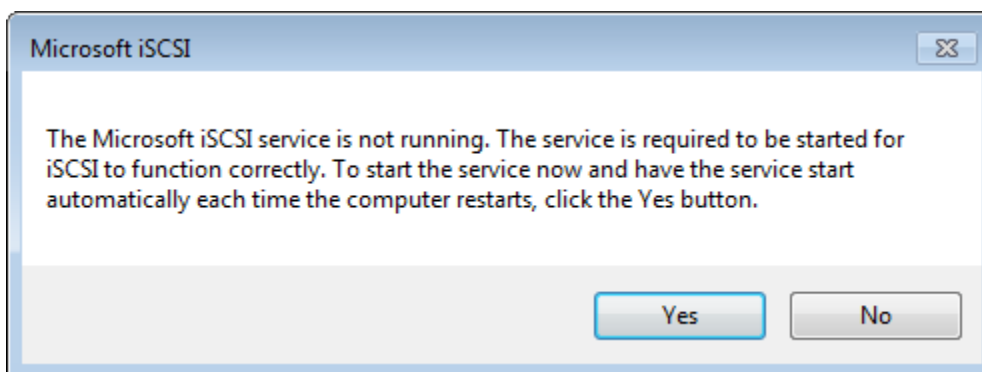
To connect your Windows client to the VTL devices

1. On the **Start** menu of your Windows client computer, enter **iscsicpl.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.

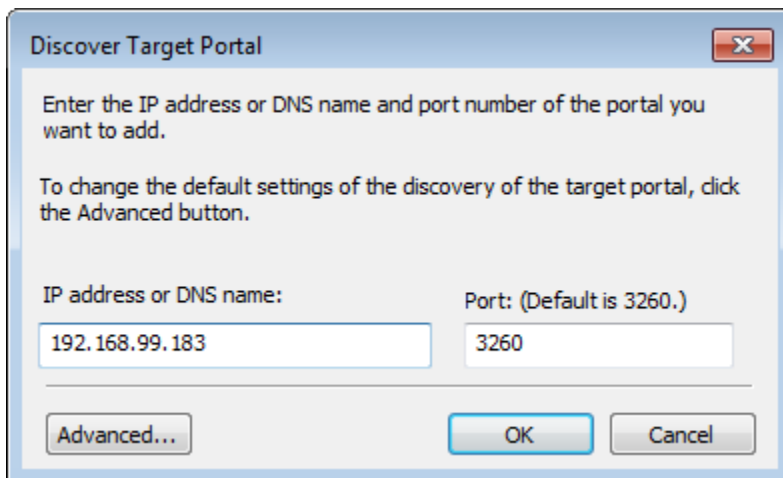
2. If prompted, choose **Yes** to start the Microsoft iSCSI initiator service.



3. In the **iSCSI Initiator Properties** dialog box, choose the **Discovery** tab, and then choose **Discover Portal**.



- In the **Discover Target Portal** dialog box, enter the IP address of your Tape Gateway 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.




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. Choose the **Targets** tab, and then choose **Refresh**. All 10 tape drives and the media changer appear in the **Discovered targets** box. The status for the targets is **Inactive**.

The following screenshot shows the discovered targets.

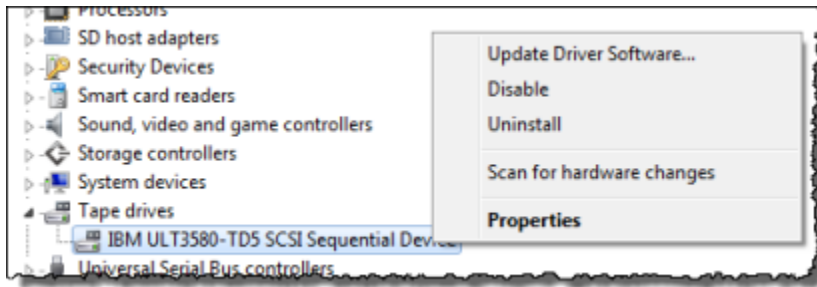


6. Select the first device and choose **Connect**. You connect the devices one at a time.
7. In the **Connect to Target** dialog box, choose **OK**.
8. Repeat steps 6 and 7 for each of the devices to connect all of them, and then choose **OK** in the **iSCSI Initiator Properties** dialog box.

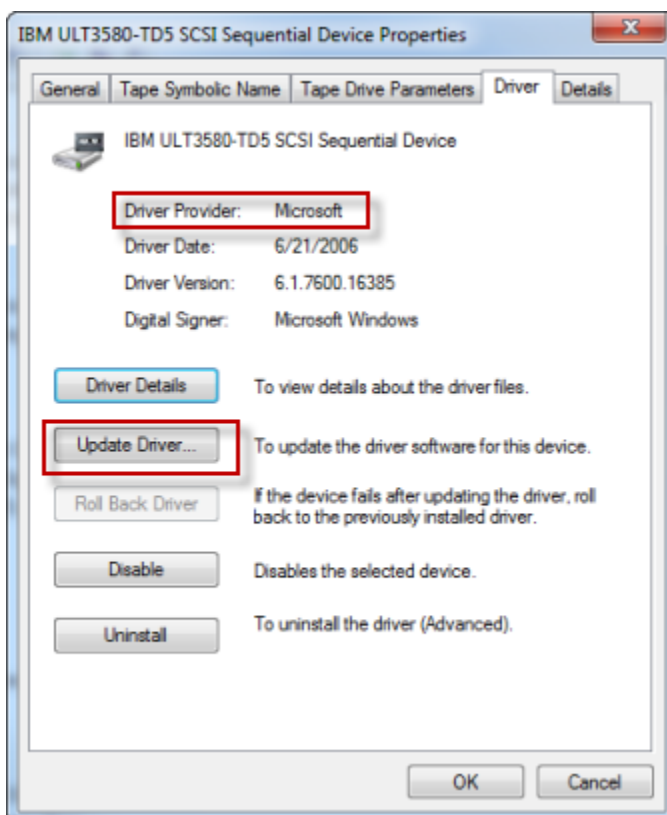
On a Windows client, the driver provider for the tape drive must be Microsoft. Use the following procedure to verify the driver provider, and update the driver and provider if necessary.

To verify the driver provider and (if necessary) update the provider and driver on a Windows client

1. On your Windows client, start Device Manager.
2. Expand **Tape drives**, choose the context (right-click) menu for a tape drive, and choose **Properties**.

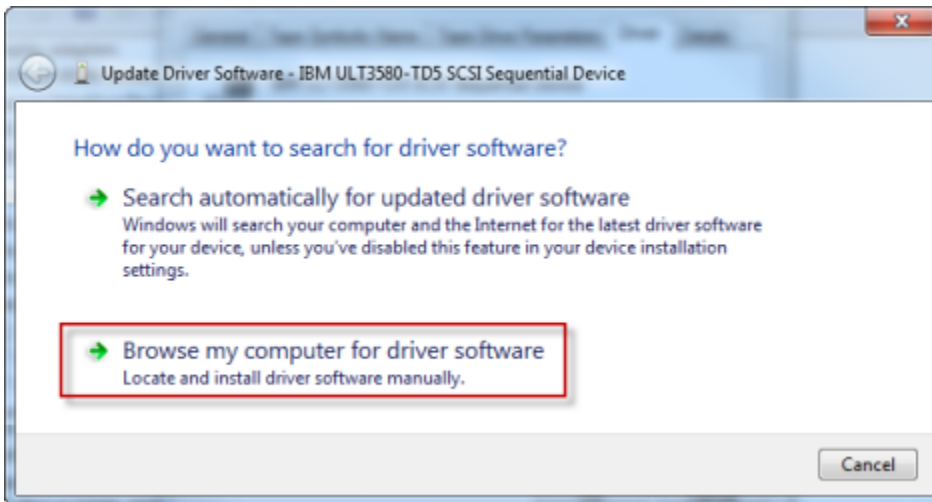


3. In the **Driver** tab of the **Device Properties** dialog box, verify that **Driver Provider** is **Microsoft**.

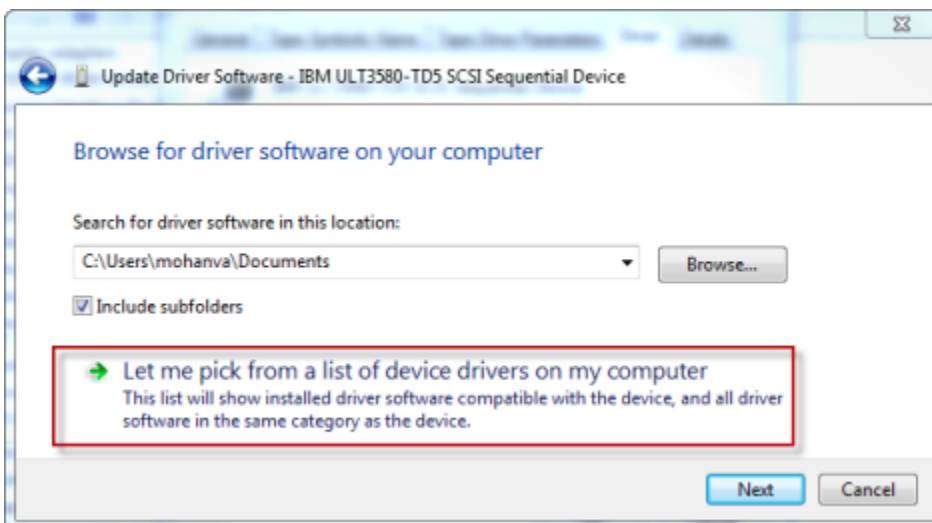


4. If **Driver Provider** is not **Microsoft**, set the value as follows:
 - a. Choose **Update Driver**.


- b. In the **Update Driver Software** dialog box, choose **Browse my computer for driver software**.



- c. In the **Update Driver Software** dialog box, choose **Let me pick from a list of device drivers on my computer**.



- d. Select **LTO Tape drive** and choose **Next**.



- e. Choose **Close** to close the **Update Driver Software** window, and verify that the **Driver Provider** value is now set to **Microsoft**.
5. Repeat steps 4.1 through 4.5 to update all the tape drives.

Connecting Your Volumes or VTL Devices 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

For a Tape Gateway setup, connecting to your VTL devices by using a Microsoft iSCSI initiator is a two-step process:

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.

The Getting Started example setup provides instructions for both these steps. It uses the Symantec NetBackup backup application. For more information, see [Connecting Your VTL Devices](#) and [Configuring NetBackup Storage 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`, as shown in the following example.

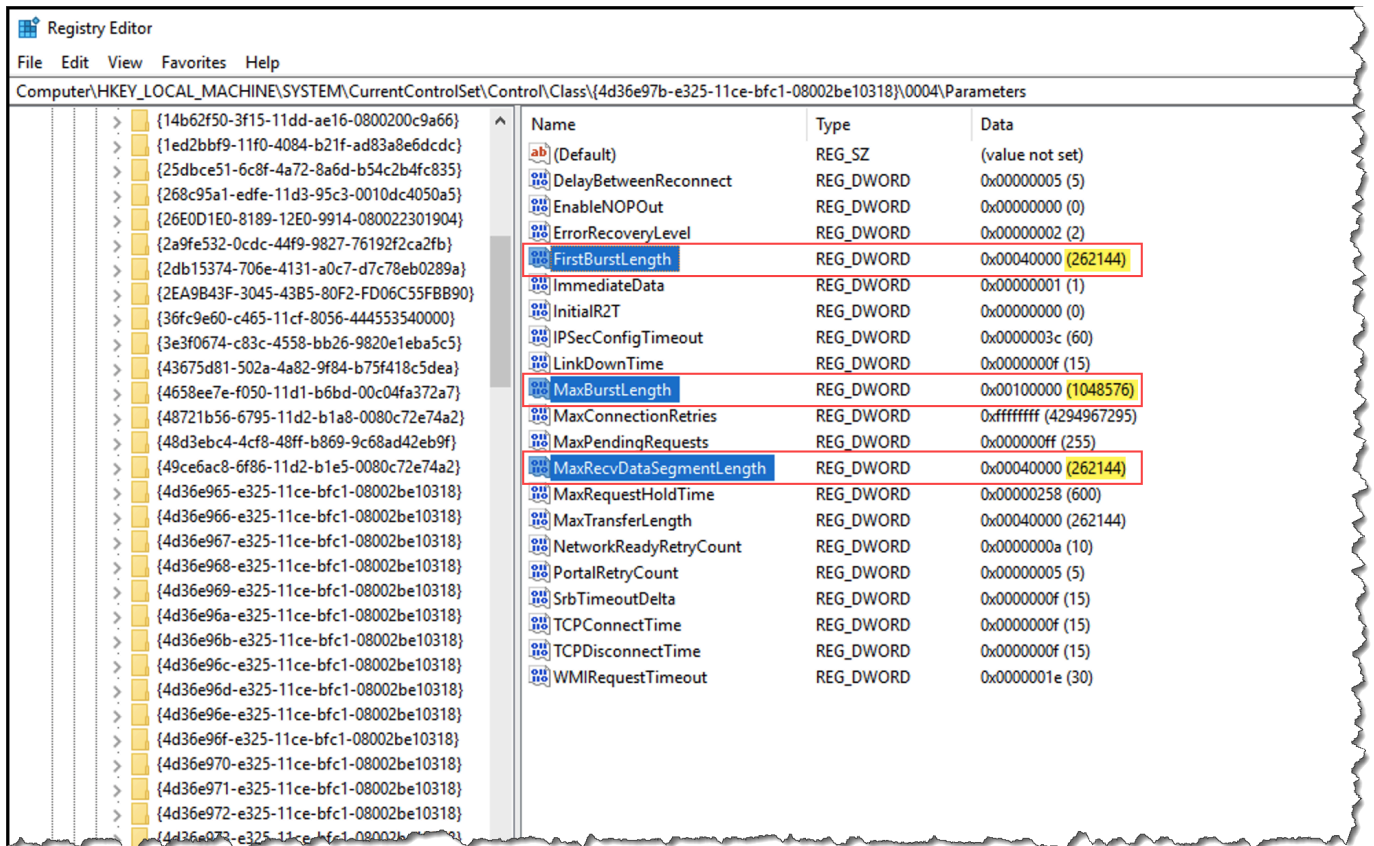


- d. To show the iSCSI settings, choose the **Parameters** subkey.
- e. Open the context (right-click) menu for the **MaxRequestHoldTime** DWORD (32-bit) value, choose **Modify**, and then change the value to **600**.

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, as shown in the following example.



- 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, as shown in the following example.



- 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

- Increase the maximum time for which requests are queued.
 - 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 VTL device 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 VTL device target on the Storage Gateway console

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

1. In the navigation pane, choose **Gateways**.
2. Choose your gateway, and then choose the **VTL Devices** tab to display all your VTL devices.
3. Choose the device that you want to configure CHAP for.
4. 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-tape-device-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**.
5. On the **VTL Devices** tab, confirm that the iSCSI CHAP authentication field 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.

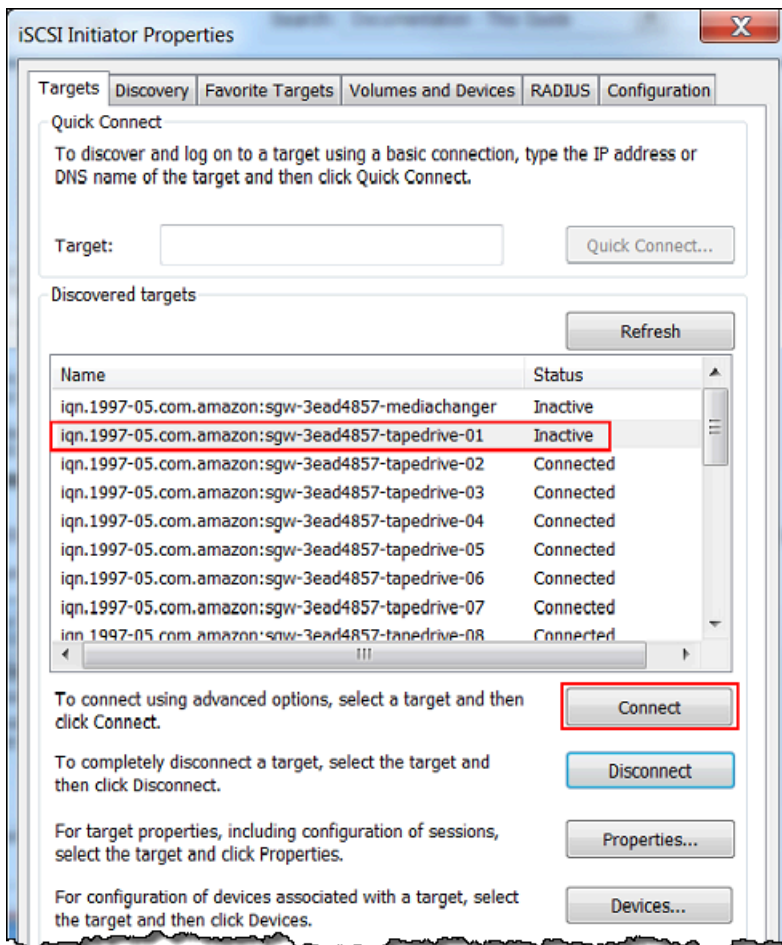
Choose **OK**, and then enter the key again.




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 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 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*.

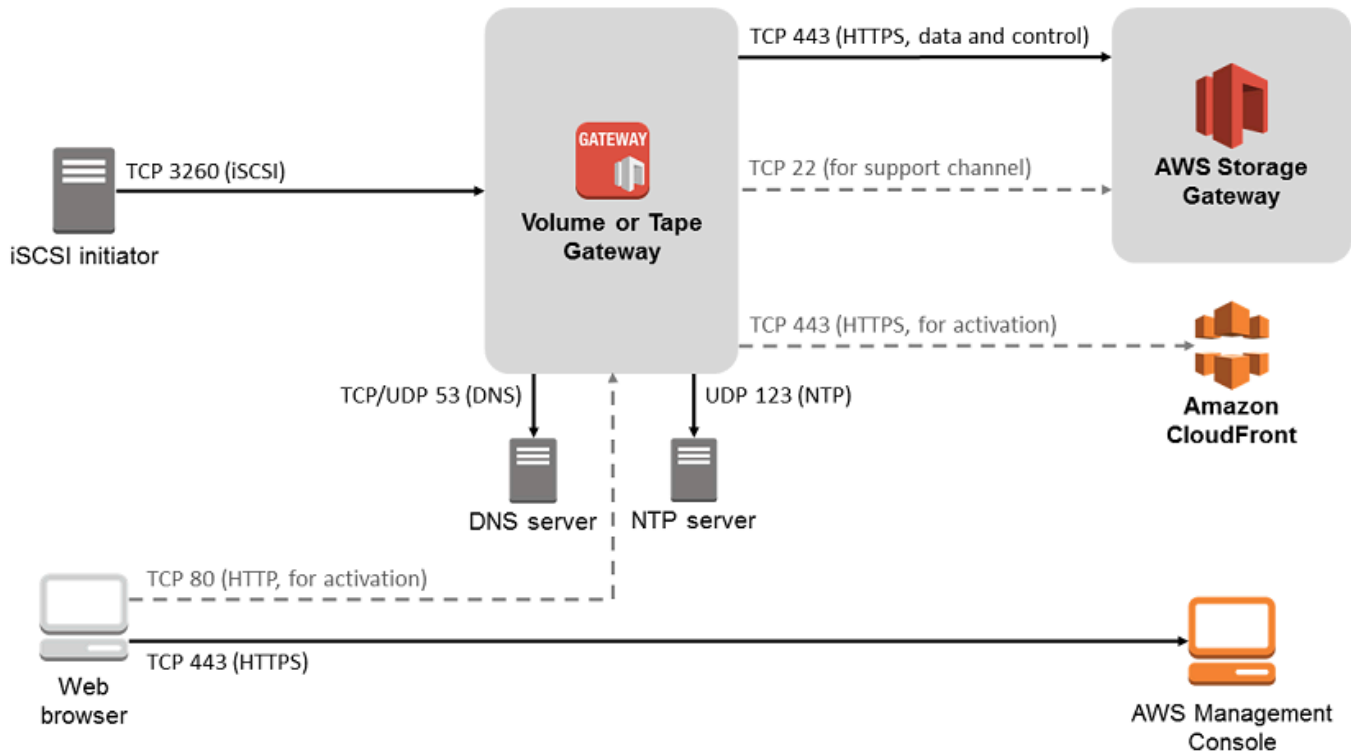
For details about AWS Direct Connect, see [What is AWS Direct Connect?](#) in the *AWS Direct Connect User Guide*.

Port Requirements

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 Tape Gateway.

Tape Gateway

The following illustration shows all of the ports you need to open for Tape 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

From	To	Protocol	Port	How Used	
				AWS Storage Gateway 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, Tape 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.

Connecting to Your Gateway

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>
Tape ARN	arn:aws:storagegateway: <i>region:account-id</i> :tape/ <i>tapebarcode</i>
Target ARN (iSCSI target)	arn:aws:storagegateway: <i>region:account-id</i> :gateway/ <i>gateway-id</i> /target/ <i>iSCSItarget</i>
VTL Device ARN	arn:aws:storagegateway: <i>region:account-id</i> :gateway/ <i>gateway-id</i> /device/ <i>vtldevice</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 AWS 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 tapes](#)
- [Recommended local disk sizes for your gateway](#)

Quotas for tapes

The following table lists quotas for tapes.

Description	Tape Gateway
Minimum size of a virtual tape	100 GiB

Description	Tape Gateway
Maximum size of a virtual tape	15 TiB
Maximum number of virtual tapes assigned to a gateway	1,500
Total size of all tapes assigned to a gateway	1 PiB
Maximum number of virtual tapes in archive	No limit
Total size of all tapes in archive	No limit

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
Tape gateway	150 GiB	64 TiB	150 GiB	2 TiB	—

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
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
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
<code>ActivationKeyExpired</code>	The specified activation key has expired.	ActivateGateway
<code>ActivationKeyInvalid</code>	The specified activation key is not valid.	ActivateGateway
<code>ActivationKeyNotFound</code>	The specified activation key was not found.	ActivateGateway
<code>BandwidthThrottleScheduleNotFound</code>	The specified bandwidth throttle was not found.	DeleteBandwidthRateLimit
<code>CannotExportSnapshot</code>	The specified snapshot cannot be exported.	CreateCachediSCSIVolume CreateStorediSCSIVolume
<code>InitiatorNotFound</code>	The specified initiator was not found.	DeleteChapCredentials
<code>DiskAlreadyAllocated</code>	The specified disk is already allocated.	AddCache AddUploadBuffer AddWorkingStorage CreateStorediSCSIVolume
<code>DiskDoesNotExist</code>	The specified disk does not exist.	AddCache AddUploadBuffer

Operation Error Code	Message	Operations That Return this Error Code
		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
		<ul style="list-style-type: none"><u>ListVolumes</u><u>ListVolumeRecoveryPoints</u><u>ShutdownGateway</u><u>StartGateway</u><u>UpdateBandwidthRateLimit</u><u>UpdateChapCredentials</u><u>UpdateMaintenanceStartTime</u><u>UpdateGatewaySoftwareNow</u><u>UpdateSnapshotSchedule</u>

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
		<ul style="list-style-type: none">ListVolumesListVolumeRecoveryPointsShutdownGatewayStartGatewayUpdateBandwidthRateLimitUpdateChapCredentialsUpdateMaintenanceStartTimeUpdateGatewaySoftwareNowUpdateSnapshotSchedule

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 Tape 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
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 gateway updates using the AWS Storage Gateway console .	June 6, 2024
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	October 24, 2023

behavior if your gateway restarts during an ongoing backup job. For more information, see [Using Your Backup Software to Test Your Gateway Setup](#).

[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

[Increased maximum tape size to 15 TiB for Tape Gateways](#)

For Tape Gateways, the maximum size of a virtual tape is now increased from 5 TiB to 15 TiB. For more information, see [Quotas for Tapes](#) in the *Storage Gateway User Guide*.

October 4, 2022

[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 .	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.</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 .	May 02, 2017
Updates to file share settings Support for cache refresh for file shares	<p>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.</p> <p>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.</p>	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 .	December 13, 2016
New Region	Storage Gateway is now available in the Canada (Central) Region. For detailed information, see AWS Regions .	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 .	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
<p>New Region</p> <p>Support for storage up to 512 TiB in size for stored volumes</p> <p>Other gateway updates and enhancements to the Storage Gateway local console</p>	<p>Tape Gateway is now available in the Asia Pacific (Seoul) Region. For more information, see AWS Regions.</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 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 .	
	<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 using the AWS Storage Gateway console. • 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.</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 Deploying an Amazon EC2 instance to host your Tape 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 Tape Gateway works (architecture).</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">• 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>.• 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.	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 Tape Gateway Appliance Software

These release notes describe the new and updated features, improvements, and fixes that are included with each version of the Tape 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 Tape 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
2024-05-28	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

Release Date	Software Version	Release Notes
2024-05-08	2.8.3	<ul style="list-style-type: none">• Addressed cloud connectivity issue when using SOCKS5 proxy• Addressed upload performance degradation issue under certain conditions (e.g., high number of tape erasure operations)
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• Improved performance for concurrent Backup and Restore workloads
2023-12-19	2.7.0	<ul style="list-style-type: none">• Operating system updates for new gateways
2023-12-14	2.6.6	<ul style="list-style-type: none">• Fixed an issue with relative positioning on larger than 5TiB tapes

Release Date	Software Version	Release Notes
2023-10-19	2.6.5	<ul style="list-style-type: none">• Added safeguards against tape overwrites by clients after a gateway restart