AWS Storage Gateway

User Guide

API Version 2013-06-30

AWS Storage Gateway: User Guide

Table of Contents

What Is AWS Storage Gateway?

Topics

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 Amazon Web Services (AWS) storage infrastructure. You can use the service to store data in the AWS Cloud for scalable and cost-effective storage that helps maintain data security. AWS Storage Gateway offers file-based, volume-based and tape-based storage solutions:

- **File Gateway** – File Gateway is a type of AWS Storage Gateway that supports a file interface into Amazon S3 and that adds to the current block-based volume and VTL storage. File Gateway combines a service and virtual software appliance, enabling you to store and retrieve objects in Amazon S3 using industry-standard file protocols such as Network File System (NFS). The software appliance, or gateway, is deployed into your on-premises environment as a virtual machine (VM) running on VMware ESXi. The gateway provides access to objects in S3 as files on a NFS mount point.

  With file gateway, you can do the following:

  - You can now store and retrieve files directly using NFS 3 or 4.1 protocol.
  - You can access your data directly in S3 from any cloud application or service.
  - You can manage your data directly in Amazon S3 using lifecycle policies, cross-region replication, and versioning. We refer to this new capability as file gateway. You can think of this as an NFS mount on S3.

  File gateway simplifies file storage in Amazon S3, integrates to existing applications through industry standard file system protocols and provides a cost-effective alternative to on-premises storage. It also provides low-latency access to data through transparent local caching. File gateway manages data transfer to and from AWS, buffers applications from network congestion, optimizes and streams data in parallel, and manages bandwidth consumption. File gateway integrates with the AWS platform. For example, it integrates AWS Identity and Access Management (IAM) to provide common access management, encryption using AWS Key Management Service (AWS KMS), monitoring using Amazon CloudWatch (CloudWatch), audit using AWS CloudTrail (CloudTrail), operations using the AWS Management Console and AWS Command Line Interface (AWS CLI), billing and cost management, and data management using S3 lifecycle policies and cross-region replication.

- **Volume Gateway** – Volume Gateway provides cloud-backed storage volumes that you can mount as Internet Small Computer System Interface (iSCSI) devices from your on-premises application servers. The gateway supports the following volume configurations:
  - **Cached volumes** – You store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally. cached volumes offer a substantial cost savings on primary storage and minimize the need to scale your storage on-premises. You also retain low-latency access to your frequently accessed data.
  - **Stored volumes** – If you need low-latency access to your entire data set, you can configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3. This configuration provides durable and inexpensive off-site backups that you can recover to your local data center or Amazon EC2. For example, if you need replacement capacity for disaster recovery, you can recover the backups to Amazon EC2.
- **Tape Gateway** – You can cost-effectively and durably archive backup data in Amazon Glacier. 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 choose to run AWS Storage Gateway either on-premises as a virtual machine (VM) appliance, or in AWS as an EC2 instance. You deploy your gateway on an EC2 instance to provision iSCSI storage volumes in AWS. Gateways hosted on EC2 instances can be used for disaster recovery, data mirroring, and providing storage for applications hosted on Amazon EC2.

For an architectural overview, see How AWS Storage Gateway Works (Architecture) (p. 3).

AWS Storage Gateway enables a wide range of use cases. For more information, see the AWS Storage Gateway detail page.

Are You a First-Time AWS Storage Gateway User?

The preceding section summarizes the storage offerings provided by the AWS Storage Gateway. For a detailed architectural overview of these offerings, see How AWS Storage Gateway Works (Architecture) (p. 3).

This documentation provides 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 a gateway. The management section shows you how to manage your gateway and resources:

- Creating a File Gateway (p. 18) provides instructions on how to create and use a file gateway. It shows you how to create a file share, map your drive to an Amazon S3 bucket and upload files and folders from your to Amazon S3.

- Creating a Volume Gateway (p. 24) provides instructions on how to create and use a volume gateway. It shows you how to create storage volumes and back up data to the volumes.

- Creating a Tape Gateway (p. 48) 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 (p. 99) provides instructions on how to perform management tasks for all gateways types and resources.

The instructions in this guide primarily show the gateway operations by using the AWS Management Console. If you want to perform these operations programmatically, see the AWS Storage Gateway API Reference for information about the supported operations.

How AWS Storage Gateway Works (Architecture)

This section provides an architectural overview of the available AWS Storage Gateway solutions.

Topics

File Gateway

To use File Gateway storage, you download a virtual machine image for the file storage gateway and activate it from the AWS Management Console or the storage gateway API. Once activated, you configure the S3 bucket(s) that the gateway will expose as file system(s) through NFS v3 or v4.1. Files written to NFS become objects in Amazon S3, with the path as the key. There is a one-to-one mapping between files and objects, and the gateway asynchronously updates the objects in Amazon S3 as you make changes to the files. Existing objects in the bucket appear as files in the filesystem and the key becomes the path. Objects are encrypted with server-side encryption with Amazon S3-managed encryption keys(SSE-S3) and all data transfer is done through HTTPS. The service optimizes data transfer between the gateway and AWS using multi-part parallel uploads or byte-range downloads, to better utilize the available bandwidth. Similar to cached volumes, a local cache is maintained to provide low latency access to the recently accessed data and reduce data egress charges. CloudWatch metrics provide insight into utilization of resources on the VM, data transfer to and from AWS, and CloudTrail tracks all API calls.

Gateway-file storage targets use cases such as ingest to S3 for cloud workloads, backup and archive, and storage tiering to the AWS Cloud. The following diagram provides an overview of the AWS Storage Gateway file storage deployment.

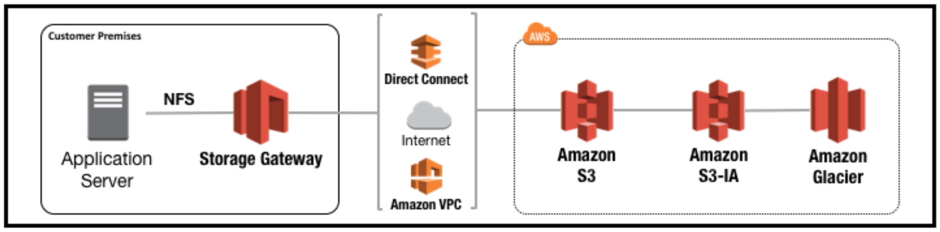

Volume Gateways

Topics

- Stored Volume Architecture (p. 5)

Cached Volume Architecture

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

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

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

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

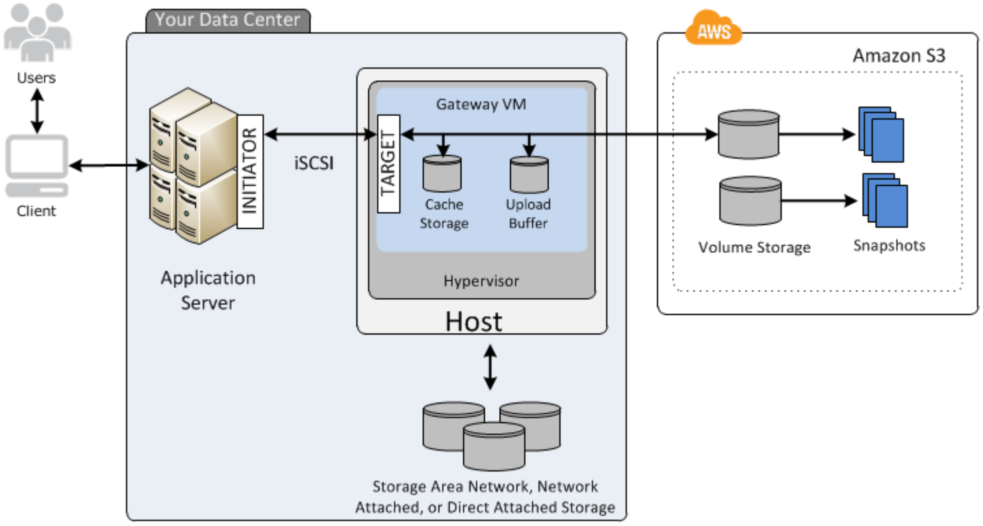

After you've installed the AWS Storage Gateway software appliance—the virtual machine (VM)—on a host in your data center and activated it, you can use the AWS Management Console to provision storage volumes backed by Amazon S3. You can also provision storage volumes programmatically using the AWS Storage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

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

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

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

You can use the following guidelines to determine the amount of disk space to allocate for cache storage. Generally, you should allocate at least 20 percent of your existing file store size as cache storage. Cache storage should also be larger than the upload buffer. This latter guideline helps ensure cache storage is large enough to persistently hold all data in the upload buffer that has not yet been uploaded to Amazon S3.

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

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

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

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

Stored Volume Architecture

Stored volumes let you store your primary data locally, while asynchronously backing up that data to AWS. Stored volumes provide your on-premises applications with low-latency access to their entire datasets, while providing durable, off-site backups. You can create storage volumes and mount them as iSCSI devices from your on-premises application servers. Data written to your stored volumes is stored on your on-premises storage hardware. This data is asynchronously backed up to Amazon Simple Storage Service (Amazon S3) in the form of Amazon Elastic Block Store (Amazon EBS) snapshots.

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

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

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

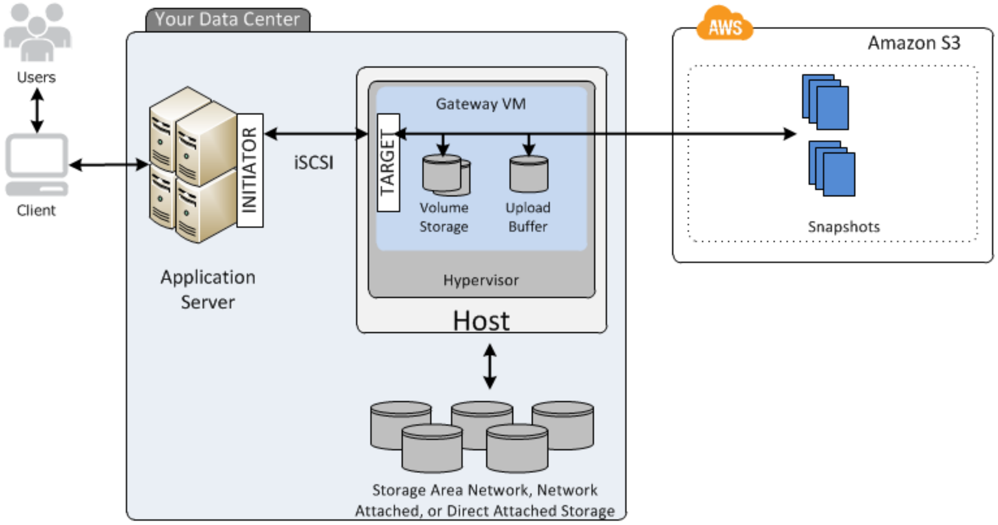

After you've installed the AWS Storage Gateway software appliance—the virtual machine (VM)—on a host in your data center and activated it, you can create gateway *storage volumes* and map them to on-premises direct-attached storage (DAS) or storage area network (SAN) disks. You can start with either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

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

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

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

Tape Gateway

Tape Gateway offers a durable, cost-effective solution to archive your data in the AWS Cloud. The VTL interface it provides lets you leverage your existing tape-based backup application 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, which 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 the Tape Gateway deployment.

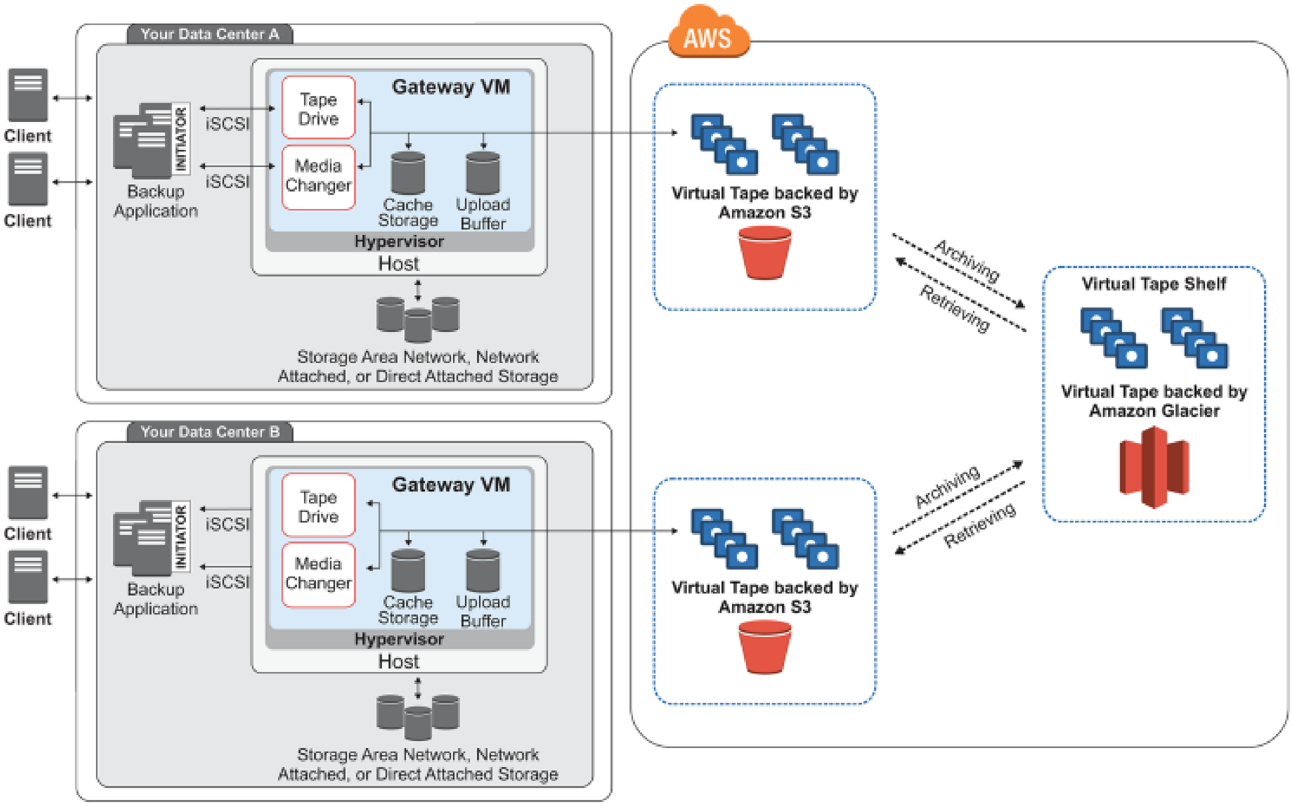

The diagram identifies the following tape gateway components:

- **Virtual tape** – Virtual tape is analogous to a physical tape cartridge. However, virtual tape data is stored in the AWS 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 AWS Storage Gateway console or programmatically by using the AWS Storage Gateway API. Each gateway can contain up to 1500 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 2.5 TiB.
- **Virtual tape library (VTL)** – A VTL is analogous to a physical tape library available on-premises with robotic arms and tape drives, including the collection of virtual tapes stored within the library. 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 Simple Storage Service (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 off-site tape holding facility. You can archive tapes from your gateway's VTL to the archive and, if needed, 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 Amazon Glacier, an extremely low-cost storage service for data archiving and backup. For more information, see Amazon Glacier.

- **Retrieving tapes** – Archived tapes cannot be read directly. To read an archived tape, you must first retrieve it to your tape gateway either by using the AWS Storage Gateway console or by using the AWS Storage Gateway API. A retrieved tape will be available in your VTL in about 24 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 and then 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 will need 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 (p. 202).

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

For guidelines to determine the amount of disk space you should allocate for the cache storage and upload buffer, see Deciding the Amount of Local Disk Storage (p. 218).

AWS Storage Gateway Pricing

For current information about pricing, see the AWS Storage Gateway pricing page.

Plan Your AWS Storage Gateway Deployment

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

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

1. **Storage solution** – Depending on your need, you can choose from one of the following storage solutions:

   - **File Gateway** – The primary uses of File Gateway include file ingest to S3 for use by object-based workloads, cost-effective storage for traditional backup applications, and tiering of on-premises file storage to S3. (If you want additional use cases covered, let us know.) You can cost-effectively and durably store and retrieve your on-premises objects in Amazon S3 using industry standard file protocols. File storage is a new addition to the set of interfaces on AWS Storage Gateway, alongside the current block-based volume and virtual tape library (VTL) storage.

- **Volume Gateway** – Volume gateways let you create storage volumes in the AWS Cloud that your on-premises applications can access as Internet Small Computer System Interface (iSCSI) targets. There are two options—cached or stored volumes.

  With cached volumes, you store volume data in AWS, with a small portion of recently accessed data in the cache on-premises. This approach enables low-latency access to your frequently accessed dataset and also provides seamless access to your entire dataset stored in AWS. This type of data access lets you scale your storage resource without having to provision additional hardware.

  With stored volumes, you store the entire set of volume data on-premises and store periodic point-in-time backups (snapshots) in AWS. In this model, your on-premises storage is primary, delivering low-latency access to your entire dataset, and AWS storage is the backup that you can restore in the event of a disaster in your data center.

  For an architectural overview of volume gateways, see Cached Volume Architecture (p. 4) and Stored Volume Architecture (p. 5).
- **Tape Gateway** – If you are looking for a cost-effective, durable, long-term, off-site alternative for data archiving, you can deploy the tape gateway solution. The virtual tape library (VTL) interface it provides lets you leverage your existing tape-based backup software infrastructure to store data on virtual tape cartridges that you create on the gateway—for more information, see Compatible Third-Party Backup Software for Tape Gateway (p. 17). When you archive tapes, you don't worry about managing tapes on your premises and arranging shipments of tapes off-site. For an architectural overview, see Tape Gateway (p. 6).

2. **Hosting option** – You can choose to run AWS Storage Gateway either on-premises, as a virtual machine (VM) appliance, or in AWS, as an Amazon EC2 instance. For more information, see Requirements (p. 11). If your data center goes offline and you don't have an available host, you can deploy a gateway on an EC2 instance. AWS 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 will need to allocate sufficient storage for the gateway VM.

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

1. For a gateway deployed on-premises, you have decided the type of host you want to set up (VMware ESXi Hypervisor or Microsoft Hyper-V) and set it up. For more information, see Requirements (p. 11). If you deploy the gateway behind a firewall, you must make sure certain ports are accessible to the gateway VM. For more information, see Requirements (p. 11). The following topics provide steps for configuring the host:
   - Configuring a VMware ESXi Host for AWS Storage Gateway (p. 208)
   - Configuring a Hyper-V Host for AWS Storage Gateway (p. 227)
2. For a tape gateway, you have installed client backup software. For more information, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Recommended Reading

Before you get started, we recommend that you read the following sections in this guide:

- What Is AWS Storage Gateway? (p. 1)
- How AWS Storage Gateway Works (Architecture) (p. 3)

Getting Started

In this section, you can find instructions about how to get started with AWS 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

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

To sign up for AWS account

1. Open http://aws.amazon.com/, and then choose **Create an AWS Account**.
2. Follow the online instructions.

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

For information about pricing, see AWS Storage Gateway Pricing on the AWS Storage Gateway detail page.

Regions

AWS Storage Gateway stores file share, volume, snapshot, and tape data in the AWS Region in which the gateway is activated. You select a region at the upper right of the AWS Storage Gateway console before you start deploying your gateway. Unless otherwise noted, the following are the available regions for storage gateway.

Note
Only cached volumes and tape gateway are available on Amazon EC2.

Region Name	Region String	File Gateway	Volume Gateway	Tape Gateway
US East (N. Virginia)	`us-east-1`	yes	yes	yes
US East (Ohio)	`us-east-2`	yes	yes	yes
US West (N. California)	`us-west-1`	yes	yes	yes
US West (Oregon)	`us-west-2`	yes	yes	yes
Canada (Central)	`ca-central-1`	yes	yes	yes
EU (Ireland)	`eu-west-1`	yes	yes	yes
EU (Frankfurt)	`eu-central-1`	yes	yes	yes
Asia Pacific (Tokyo)	`ap-northeast-1`	yes	yes	yes
Asia Pacific (Seoul)	`ap-northeast-2`	yes	yes	yes
Asia Pacific (Singapore)	`ap-southeast-1`	yes	yes	no
Asia Pacific (Sydney)	`ap-southeast-2`	yes	yes	yes
South America (São Paulo)	`sa-east-1`	yes	yes	no

Requirements

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

Topics

Hardware Requirements

When deploying your gateway on-premises, you must make sure that the underlying hardware on which you are deploying the gateway VM is able to dedicate the following resources:

- Four or eight virtual processors assigned to the VM.
- 7.5 GB of RAM assigned to the VM

- 75 GB of disk space for installation of VM image and system data

For more information, see Optimizing Gateway Performance (p. 149)

For information about how your hardware affects the performance of the gateway VM, see AWS Storage Gateway Limits (p. 296).

When deploying your gateway on Amazon EC2, you must use the **m3**, **i2**, **c3**, **c4**, **r3**, **d2**, and **m4** instance types and the instance size must be at least size **xlarge**. You must select one of these instance types for the gateway to function. For more information, see AWS Storage Gateway in AWS Marketplace.

Storage Requirements

In addition to 75 GB disk space for the VM, you will also need additional disks for the gateway:

- For file gateway, you will need storage for the local cache.
- For cached volumes, you will need storage for the local cache and an upload buffer.
- For stored volumes, you will need storage for your entire dataset and an upload buffer.
- For tape gateway, you will need storage for the local cache and an upload buffer.

For more information about how to add disks, see Provisioning Local Disk Storage for the Gateway VM (VMWare) (p. 218).

For information about gateway limits, see AWS Storage Gateway Limits (p. 296).

Network and Firewall Requirements

Your locally deployed 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.

Topics

Port Requirements

AWS Storage Gateway requires the following ports for its operation.

Protocol	Port	Source	Destination	How Used
TCP	443	Storage Gateway	Internet	For communication from AWS Storage Gateway to the AWS service endpoint. For information about service endpoints, see Allowing AWS Storage Gateway

Protocol	Port	Source	Destination	How Used
				Access through Firewalls and Routers (p. 14).
TCP	80	Local networks	Storage Gateway	By local systems to obtain the storage gateway activation key. Port 80 is only used during activation of the Storage Gateway appliance. **Note** AWS 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 AWS Storage Gateway Management Console, the host from which you connect to the console must have access to your gateway's port 80.

Protocol	Port	Source	Destination	How Used
TCP	3260	Local networks	Storage Gateway	By local systems to connect to iSCSI targets exposed by the gateway.
UDP	53	Storage Gateway	Domain Name Service (DNS) server	For communication between AWS Storage Gateway and the DNS server.
TCP	22	Storage Gateway	Internet	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.

Allowing AWS Storage Gateway Access through Firewalls and Routers

Your locally deployed gateway requires access to the following 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.

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

The following table provides a list of region strings for the available regions.

Note
Only cached volumes and tape gateway are available on Amazon EC2.

Region Name	Region String	File Gateway	Volume Gateway	Tape Gateway
US East (N. Virginia)	`us-east-1`	yes	yes	yes
US East (Ohio)	`us-east-2`	yes	yes	yes
US West (N. California)	`us-west-1`	yes	yes	yes

Region Name	Region String	File Gateway	Volume Gateway	Tape Gateway
US West (Oregon)	`us-west-2`	yes	yes	yes
Canada (Central)	`ca-central-1`	yes	yes	yes
EU (Ireland)	`eu-west-1`	yes	yes	yes
EU (Frankfurt)	`eu-central-1`	yes	yes	yes
Asia Pacific (Tokyo)	`ap-northeast-1`	yes	yes	yes
Asia Pacific (Seoul)	`ap-northeast-2`	yes	yes	yes
Asia Pacific (Singapore)	`ap-southeast-1`	yes	yes	no
Asia Pacific (Sydney)	`ap-southeast-2`	yes	yes	yes
South America (São Paulo)	`sa-east-1`	yes	yes	no

Depending on your gateway's region, you replace *`region`* in the endpoint with the corresponding region string. For example, if you create a gateway in the US West (Oregon) region, the endpoint looks like this: `storagegateway.us-west-2.amazonaws.com:443`.

Configuring Security Groups for Your Amazon EC2 Gateway Instance

A security group controls traffic to your Amazon EC2 gateway instance. When you create an instance from the Amazon Machine Image (AMI) for AWS Storage Gateway from AWS Marketplace, you have two choices for launching the instance. To launch the instance by using the **1-Click Launch** feature of AWS Marketplace, follow the steps in Provisioning an Amazon EC2 Host (p. 53) . We recommend using this **Manual Launch** feature.

You can also launch an instance by using the **1-Click Launch** feature in AWS Marketplace. In this case, an autogenerated security group that is named `AWS Storage Gateway-1-0-AutogenByAWSMP-` is created. This security group has the correct rule for port 80 to activate your gateway. For more information about security groups, see Security Group Concepts in the *Amazon EC2 User Guide for Linux Instances*.

Regardless of the security group that you use, 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 you allow connections only on ports 3260 (for iSCSI connections) and 80 (for activation).
- If you want to activate your gateway from an 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 Enabling AWS Support to Access a Gateway Hosted on an Amazon EC2 Instance (p. 196).

If you are using an Amazon EC2 instance as an initiator (that is, to connect to the iSCSI targets on the gateway you deployed on Amazon EC2), then we recommend a two-step approach:

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

Supported Hypervisors and Host Requirements

You can choose to run AWS Storage Gateway either on-premises as a virtual machine (VM) appliance, or in AWS as an Amazon Elastic Compute Cloud (Amazon EC2) instance.

AWS Storage Gateway supports the following hypervisor versions and hosts:

- VMware ESXi Hypervisor (version 4.1, 5.0, 5.1, 5.5 or 6.0)—A free version of VMware is available on the VMware website. You will also need a VMware vSphere client to connect to the host.

  Note
  Currently, file gateway supports only the VMware ESXi Hypervisor.
- Microsoft Hyper-V Hypervisor (version 2008 R2, 2012, or 2012 R2)—A free, stand-alone version of Hyper-V is available at the Microsoft Download Center. You will need a Microsoft Hyper-V Manager on a Microsoft Windows client computer to connect to the host.
- EC2 instance—AWS Storage Gateway provides an Amazon Machine Image (AMI) that contains the gateway VM image. Only gateways created with cached volumes and tape gateways can be deployed on Amazon EC2. For information about how to deploy a gateway on Amazon EC2, see Provisioning an Amazon EC2 Host (p. 53).

Supported Protocols For File Gateway

File gateway supports the following protocols:

- Clients that connect to the gateway using NFS V3 and V4.1 over TCP.

  Note
  Windows clients support NFS V3 and we are evaluating options to enhance Windows integration.

Supported File System Operations For File Gateway

Your NFS client can write, read, delete, and truncate files. Writes are sent to S3 through optimized multi-part uploads through a write-back cache. Reads are first served through the local cache. If data is not available, it is fetched through S3 as a read through cache. Writes and reads are optimized in that only the parts that are changed or requested are transferred through the Gateway. Deletes remove objects from S3. Directories are managed as folder objects in S3, using the same syntax as the S3 console.

Supported iSCSI Initiators

When you deploy a gateway-cached or gateway-stored volume gateway, you can create iSCSI storage volumes on your gateway. When you deploy a tape gateway, the gateway is preconfigured with one media changer and ten 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, AWS Storage Gateway supports the following iSCSI initiators:

- Windows Server 2012 and Windows Server 2012 R2

- Windows Server 2008 and Windows Server 2008 R2
- Windows 7
- Red Hat Enterprise Linux 5
- Red Hat Enterprise Linux 6
- Red Hat Enterprise Linux 7
- VMware ESX Initiator (Provides an alternative to using initiators in the guest operating systems of your VMs.)

Important
Storage Gateway does not support Microsoft Multipath I/O (MPIO) from Windows clients. Although AWS Storage Gateway enables applications that are clustered using Windows Server Failover Clustering (WSFC) to use the iSCSI initiator to access your gateway's volumes, connecting multiple hosts to the same iSCSI target is not supported.

Compatible Third-Party Backup Software for Tape Gateway

You use backup software to read, write, and manage tapes with a tape gateway. A tape gateway setup is compatible with many third-party backup software packages, including the following:

Backup Software	Version
Backup Exec	2012, 2014, 15 and 16
Dell NetVault Backup	10.0
EMC NetWorker	8.x
HPE Data Protector	9.x
Microsoft System Center Data Protection Manager	2012 R2
Symantec NetBackup	7.x
Veeam Backup & Replication	V7, V8 and V9

Accessing AWS Storage Gateway

You can use the AWS Storage Gateway 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. You can find the AWS Management Console at AWS Storage Gateway console.

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 AWS Storage Gateway (p. 360).

You can also use the AWS SDKs to develop applications that interact with AWS Storage Gateway. The AWS SDKs for Java, .NET, and PHP wrap the underlying AWS Storage Gateway API to simplify your programming tasks. For information about downloading the SDK libraries, see Sample Code Libraries.

Creating Your Gateway

To create your gateway, open the AWS Storage Gateway console and choose the AWS Region you want to create your gateway in. If you haven't created a gateway in this region, the AWS Storage Gateway page is displayed.

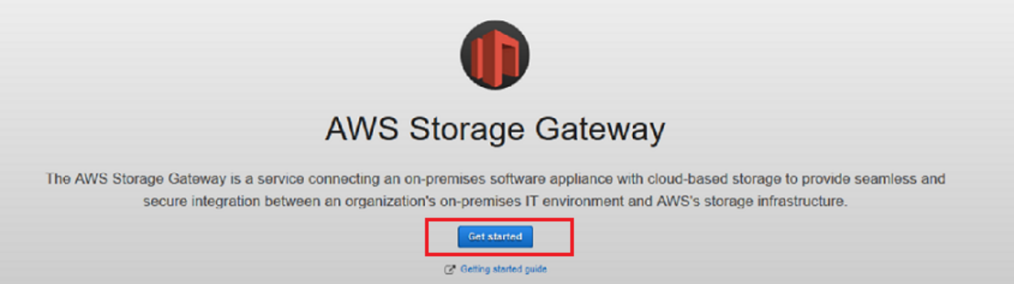

AWS Storage Gateway

The AWS Storage Gateway is a service connecting an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and AWS's storage infrastructure.

Choose **Get started** to open the **Select gateway type** page. On the **Create gateway** page, you select a gateway type. If you have a gateway in this region, the console shows your gateway in the console.

Topics

- Creating a File Gateway (p. 18)
- Creating a Volume Gateway (p. 24)
- Creating a Tape Gateway (p. 48)

Creating a File Gateway

In this section, you can find instructions about how create and use a file gateway.

Topics

- Creating a File Gateway (p. 18)
- Creating a File Share (p. 22)
- Using Your File Share (p. 23)

Creating a File Gateway

In this section, you can find instructions about how to download, deploy, and activate a file gateway.

Topics

Selecting a Gateway Type

With File Gateway (p. 3), you store and retrieve objects in Amazon S3 with a local cache for low latency access to your most recently used data.

To create a gateway

1. Open the AWS Management Console at http://console.aws.amazon.com/storagegateway/home, and choose the AWS Region you want to create your gateway in.

   If you have a gateway in this region, the console shows your gateway in the console.
2. In the navigation pane, choose **Gateways**, and then choose **Create gateway**.
3. On the **Select gateway type** page, choose **File gateway**, and then choose **Next**.

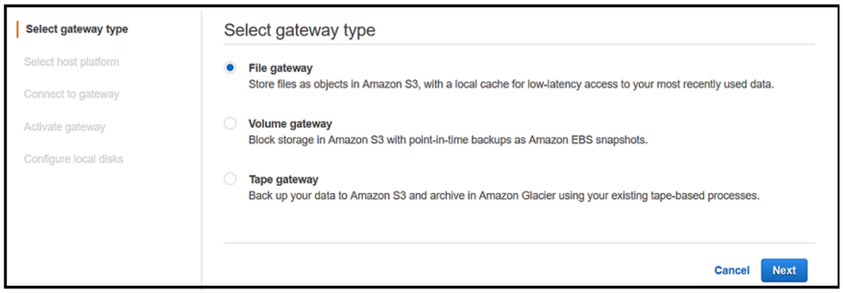

Select gateway type

● **File gateway**
Store files as objects in Amazon S3, with a local cache for low-latency access to your most recently used data.

○ **Volume gateway**
Block storage in Amazon S3 with point-in-time backups as Amazon EBS snapshots.

○ **Tape gateway**
Back up your data to Amazon S3 and archive in Amazon Glacier using your existing tape-based processes.

Choosing a Host Platform

You run file gateway on-premises as a virtual machine (VM) appliance on VMware ESXi hypervisor. For information about supported host platforms, see Supported Hypervisors and Host Requirements (p. 16).

To choose a host platform

1. On the **Select host platform** page, choose the VMware EsXi platform.

   Note
   File gateway currently supports only VMware ESXi
2. Choose **Download image** next to your **VMware ESXi** to download a .zip file that contains the .ova file for your virtualization platform.

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

Provisioning a VMware Host

Following, you can find how to provision an on-premises VMware host and deploy your gateway. You perform this task in the VMware vSphere client and return to the AWS Storage Gateway console to continue your gateway setup.

To provision a VMware host

1. Review the minimum host requirements in Requirements (p. 11).
2. Provision a host in your data center with the VMware ESXi hypervisor. For a minimum set of instructions to install the hypervisor, see Configuring a VMware ESXi Host for AWS Storage Gateway (p. 208).

   Note
   If you plan to deploy AWS Storage Gateway using VMware High Availability (HA) for failover protection, see Using AWS Storage Gateway with VMware High Availability (p. 225). In this tutorial exercise, you deploy your AWS Storage Gateway VM on a single host with no clustering or failover.

To deploy your gateway on a VMware host

1. Connect to your gateway host's hypervisor by using your VMware vSphere client.
2. Deploy the OVF template package that you downloaded.
3. Choose **Next** through the following three screens. You might be prompted to select a data store on which to store the `.ova` package.
4. Allocate disks with **Thick provisioned format**.
5. Synchronize the time of your gateway VM to match your gateway host's time.

   You must also ensure that the clock on your gateway host is synchronized with a Network Time Protocol (NTP) server.

This procedure provides minimal instructions to help you deploy your gateway quickly. If you are not familiar with using the VMware vSphere client software, see Deploying the AWS Storage Gateway VM to Your VMWare Host (p. 212) for detailed instructions.

Next, you allocate local disks for your deployed gateway VM. After you activate your gateway, you configure the disks as local storage for the gateway's use.

To allocate a local disk from direct-attached storage

Important
File gateway requires at least a disk use as cache storage.

1. Decide the number and size of disks to allocate for your gateway.
2. Start your VMware vSphere client, and then connect to your host.
3. In the client, follow the instructions provided by your host's client and add two virtual disks.
4. Configure the disks according to the sizes you decided for your gateway.

   For your disk configuration, we recommend thick provisioning. We also strongly recommend that you choose the **Specify a datastore or datastore cluster** option and select a data store for these disks different from the one that is used for the gateway VM.
5. Configure your VM to use paravirtualized controllers.

This procedure provides minimal instructions to help you to quickly allocate disks for your gateway. If you are not familiar with using the VMware vSphere client software, see Provisioning Local Disk Storage for the Gateway VM (VMWare) (p. 218) for detailed instructions.

Connecting to Your Gateway

You get the IP address of your gateway VM and use it to activate your gateway. For gateways deployed and activated on an on-premises host, you can get the IP address from your gateway VM local console or your hypervisor client. For gateways deployed and activated on an Amazon EC2 instance, you can get the IP address from the Amazon EC2 console. The activation process associates your gateway with your AWS account. You gateway VM must be running for activation to succeed.

To get the IP address for you gateway VM from the local console

1. Log on to your gateway VM local console. For detailed instructions, see Logging In to Your AWS Storage Gateway Local Console (p. 157).
2. Get the IP address from the top of the menu page.
3. Take note of the gateway IP address.
4. Type the IP address of your gateway for **IP address**, and then choose the **Connect gateway** button.

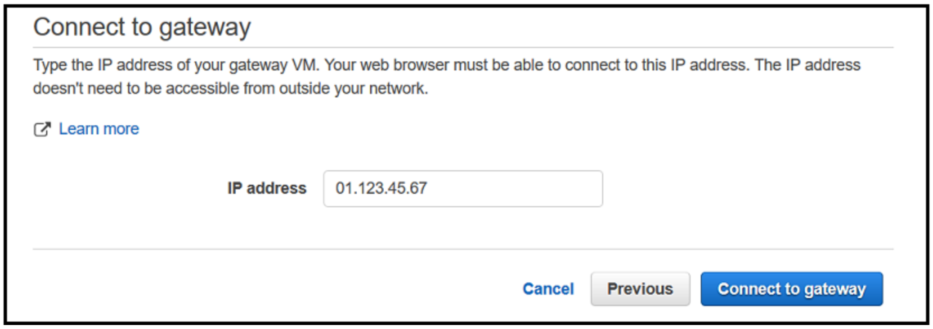

Connect to gateway

Type the IP address of your gateway VM. Your web browser must be able to connect to this IP address. The IP address doesn't need to be accessible from outside your network.

IP address: 01.123.45.67

Activating Your Gateway

To activate your gateway

1. To complete the activation process, provide information on the activation page to configure your gateway setting:

   - **Gateway Time Zone** specifies the time zone to use for your gateway.
   - **Gateway Name** identifies your gateway. You use this name to manage your gateway in the console; you can change it after the gateway is activated. This name must be unique to your account.

     The following screenshot shows the activation page for file gateway.

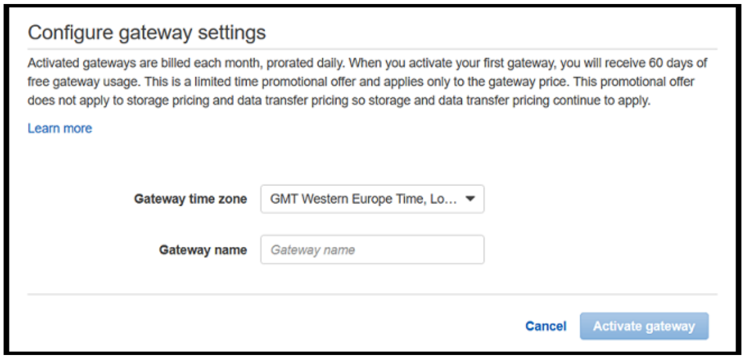

Configure gateway settings

Activated gateways are billed each month, prorated daily. When you activate your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies only to the gateway price. This promotional offer does not apply to storage pricing and data transfer pricing so storage and data transfer pricing continue to apply.

Gateway time zone: GMT Western Europe Time, Lo...

Gateway name: Gateway name

2. Choose **Activate gateway**.
3. If activation is not successful, see Troubleshooting Your Gateway (p. 191) for possible solutions.

Configuring Local Disks

When you deployed the VM, you allocated local disks for your gateway. Now, you configure your gateway to use these disks.

To configure local disks

1. On the **Configure local disks** page, identify the disks you allocated and decide which ones you want to use for cached storage. For information about disk size limits, see Configuration and Performance Limits (p. 297).

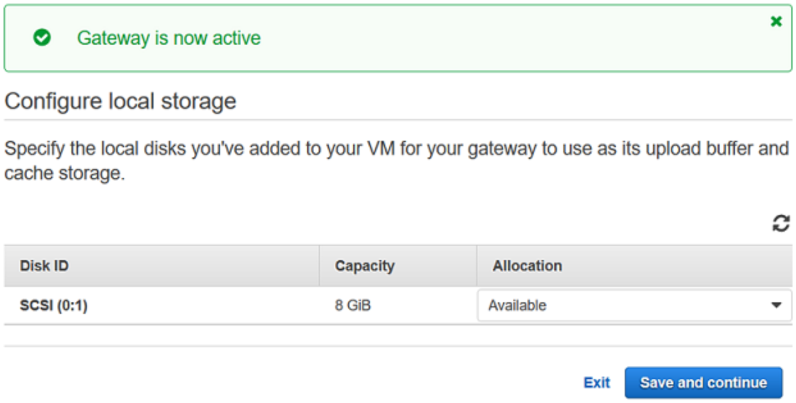

Gateway is now active

Configure local storage

Specify the local disks you've added to your VM for your gateway to use as its upload buffer and cache storage.

Disk ID	Capacity	Allocation
SCSI (0:1)	8 GiB	Available

2. Choose **Cache** for the disk you want to configure as a cache storage.

   If you don't see your disks, choose **Refresh**.
3. Choose **Save and continue** to save your configuration settings.

Creating a File Share

In the following section, you can find instructions on how to create a file share.

To create a file share

1. On the navigation pane, choose **File shares**, and then choose **Create file share**.

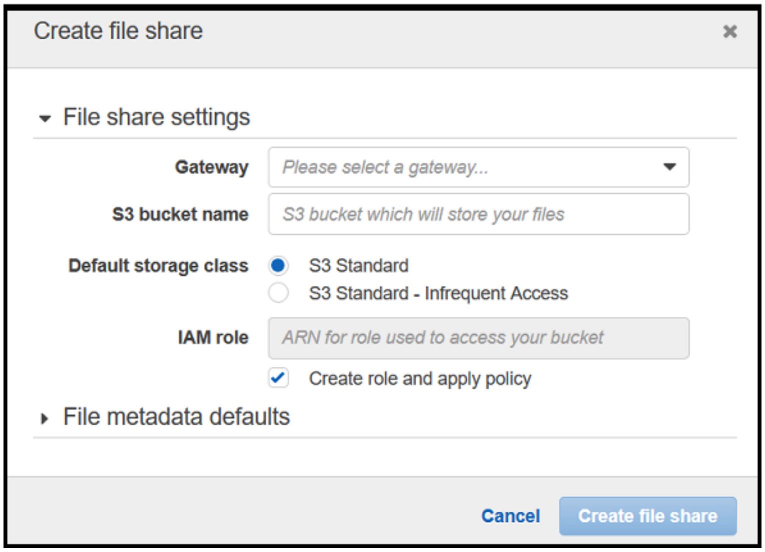

Create file share

File share settings

Gateway: Please select a gateway...

S3 bucket name: S3 bucket which will store your files

Default storage class: ● S3 Standard ○ S3 Standard - Infrequent Access

IAM role: ARN for role used to access your bucket

☑ Create role and apply policy

File metadata defaults

2. In the **Create file share** dialog box, select your gateway from the **Gateway** list.
3. In the **S3 bucket name** box, provide the name of the Amazon S3 bucket you want your gateway to upload your files into.

   Note
   The bucket name you specify must exist and you must have permissions to access the bucket.
4. Select the type of **Default storage class** you to use for your S3 bucket.
5. Choose the **Create role and apply policy** checked box if it is not selected. Storage Gateway creates an IAM access policy and role on your behalf. This role and policy allows the gateway to upload files to your Amazon S3 bucket. If you don't want the gateway to create a role on your behalf, you can specify your own role in the **IAM role** box.
6. By default, files and directories in your S3 bucket are assigned metadata values if you don't explicitly assign the metadata. You can change the metadata values.

   To change the values, expand **File metadata defaults** and provide the new values.
7. Choose **Create file share**.
8. In the **Confirm IAM role creation** dialog box, select the check box, and then choose the **Create role and apply policy** button to allow the gateway to create a role on your behalf. For more information, see Granting Access to an Amazon S3 Destination (File Gateway Only) (p. 101).

Using Your File Share

Following, you can find instructions about how to connect your file share to your Microsoft Windows client, use your share, and test your file gateway.

Topics

- Connecting Your File Share to Microsoft Your Windows Client (p. 23)
- Testing Your File Gateway (p. 24)

Connecting Your File Share to Microsoft Your Windows Client

Now, you mount the file share on a drive on your Microsoft Windows client and map it to your Amazon S3 bucket.

To mount a file share and map it to a Amazon S3 bucket

1. Turn on Services for NFS in your Windows client.
2. At your Windows client's command prompt ,type the following command:

   ```
   mount –o nolock Your gateway VM IP address:/ S3 bucket name Drive letter
   on your windows client:
   ```

   For example, if your VM's IP address is 123.456.1.2, and your Amazon S3 bucket name is test-bucket and you want to map to drive T, your command looks like the following:

   ```
   mount –o nolock 123.456.1.2:/test-bucket T:
   ```

   Note
   If a folder and an object exist in an Amazon S3 bucket and they have the same name and the object name doesn't contain a trailing slash, only the folder is visible in file gateway. For example, if bucket contains and object named test, test/ and a folder named test/test1, only test/ and test/test1 will be visible in file gateway.
   You may need to re-mount your file share after a reboot of your client.

Testing Your File Gateway

You can copy files and folders to your mapped drive and the files automatically upload to your Amazon S3 bucket.

To upload files from your windows client to Amazon S3

1. On your Windows client, navigate to your drive letter you mounted your file share on. The name of your drive is preceded by the name of your S3 bucket.
2. Copy some files or a folder to the drive.
3. On the Amazon S3 management console, navigate to your mapped bucket. You should see your files and folders you copied in the Amazon S3 bucket you specified.

   You can see the file share you created in the **File shares** tab in the AWS Storage Gateway Management Console.

Your NFS client can write, read, delete, rename, and truncate files. Reads are served from a read-through cache, that is, if data is not available, it is fetched from S3 and added to the cache. Writes are sent to S3 through optimized multi-part uploads through a write-back cache. Read and writes are optimized so that only the parts that are requested or changed are transferred over the network. Deletes remove objects from S3. Directories are managed as folder objects in S3, using the same syntax as the Amazon S3 console. You can rename empty directories. Recursive file system operation performance (for example ls –l) will depend on the number of objects in your bucket.

Creating a Volume Gateway

In this section, you can find instructions about how to create and use a volume gateway.

Topics

Creating a Gateway

In this section, you can find instructions about how to download, deploy, and activate a volume gateway.

Topics

Selecting a Gateway Type

After you sign up for AWS, the first thing you decide is the type of gateway you want to create:

- Cached volumes (p. 4)—Store your data in AWS and retain a copy of frequently accessed data subsets locally.
- Stored volumes (p. 5)—Store all your data locally and asynchronously back up point-in-time snapshots to AWS.

To create a gateway

1. Open the AWS Console at http://console.aws.amazon.com/storagegateway/home, and choose the AWS Region you want to create your gateway in.

   If you have a gateway in this region, the console shows your gateway in the console.
2. In the navigation pane, choose **Gateways**, and then choose **Create gateway**.
3. On the **Select gateway type page**, choose **Volume gateway**, choose the type of volume, and then choose **Next**.

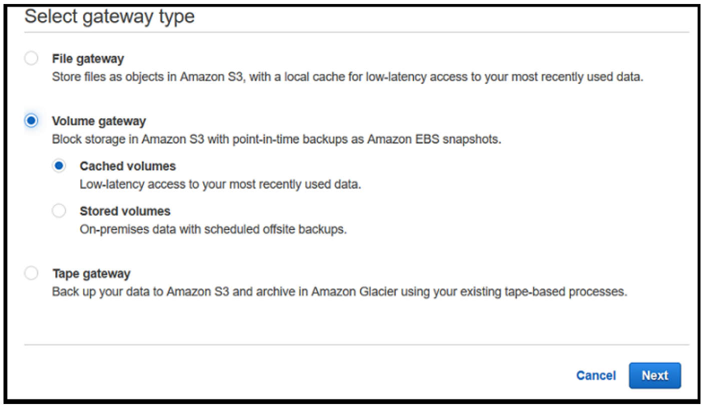

Select gateway type

○ **File gateway**
Store files as objects in Amazon S3, with a local cache for low-latency access to your most recently used data.

● **Volume gateway**
Block storage in Amazon S3 with point-in-time backups as Amazon EBS snapshots.

● **Cached volumes**
Low-latency access to your most recently used data.

○ **Stored volumes**
On-premises data with scheduled offsite backups.

○ **Tape gateway**
Back up your data to Amazon S3 and archive in Amazon Glacier using your existing tape-based processes.

Choosing a Host Platform

If you are creating the gateway on-premises, you download and deploy the gateway VM and then activate the gateway. If you are creating the gateway on an Amazon EC2 instance, you launch an Amazon Machine Image (AMI) that contains the gateway VM image and then activate the gateway.

Note
Stored volumes cannot be run on Amazon EC2 instance.

To choose a host platform

1. On the **Select host platform** page, choose the virtualization platform you want to run your gateway on.

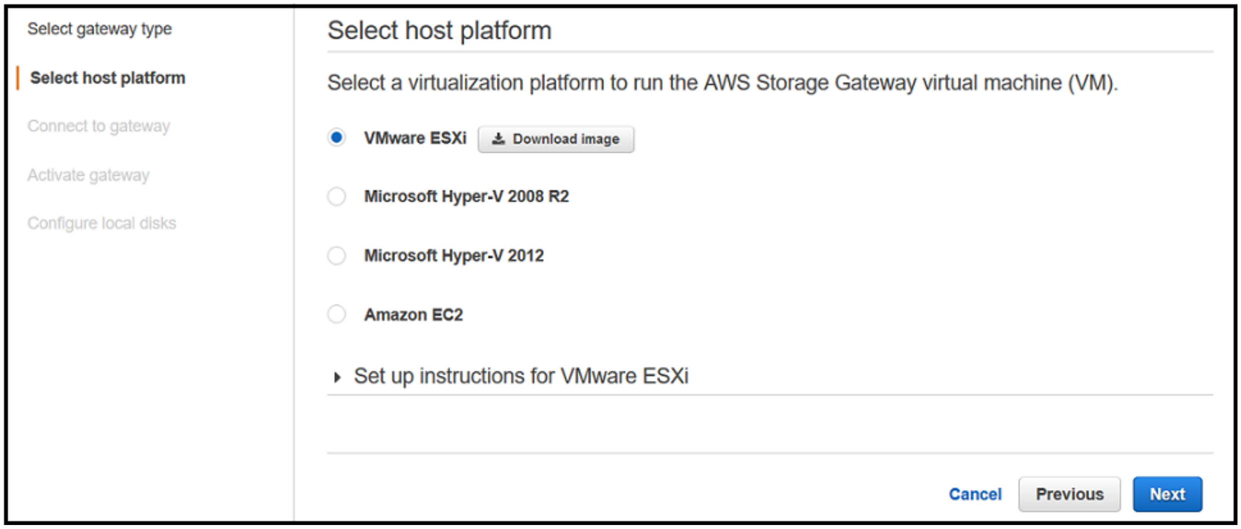

Select host platform

Select a virtualization platform to run the AWS Storage Gateway virtual machine (VM).

● VMware ESXi
○ Microsoft Hyper-V 2008 R2
○ Microsoft Hyper-V 2012
○ Amazon EC2

2. Choose **Download image** next to your virtualization platform to download a .zip file that contains the .ova file for your virtualization platform.

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

When your download is complete, you provision a host and deploy your gateway VM you downloaded.

Depending on your host platform, choose one of the following links.

Topics

- Provisioning a VMware Host (p. 26)
- Provisioning a Hyper-V Host (p. 27)
- Provisioning an Amazon EC2 Host (p. 29)

Provisioning a VMware Host

Following, you can find how to provision an on-premises VMware host and deploy your gateway. You perform this task in the VMware vSphere client and return to the AWS Storage Gateway console to continue your gateway setup.

To provision a VMware host

1. Review the minimum host requirements in Requirements (p. 11).
2. Provision a host in your data center with the VMware ESXi hypervisor. For a minimum set of instructions to install the hypervisor, see Configuring a VMware ESXi Host for AWS Storage Gateway (p. 208).

   Note
   If you plan to deploy AWS Storage Gateway using VMware High Availability (HA) for failover protection, see Using AWS Storage Gateway with VMware High Availability (p. 225). In this tutorial exercise, you deploy your AWS Storage Gateway VM on a single host with no clustering or failover.

To deploy your gateway on a VMware host

1. Connect to your gateway host's hypervisor by using your VMware vSphere client.
2. Deploy the OVF template package that you downloaded.
3. Choose **Next** through the following three screens. You might be prompted to select a data store on which to store the `.ova` package.
4. Allocate disks with **Thick provisioned format**.
5. Synchronize the time of your gateway VM to match your gateway host's time.

   You must also ensure that the clock on your gateway host is synchronized with a Network Time Protocol (NTP) server.

This procedure provides minimal instructions to help you deploy your gateway quickly. If you are not familiar with using the VMware vSphere client software, see Deploying the AWS Storage Gateway VM to Your VMWare Host (p. 212) for detailed instructions.

Next, you allocate local disks for your deployed gateway VM. After you activate your gateway, you configure the disks as local storage for the gateway's use.

To allocate a local disk from direct-attached storage

> **Important**
> All gateways require a disk of a minimum size of 150 GiB for use as an upload buffer. Cached and tape gateway setups require additional disks for use as cache storage. The minimum size of a cache disk is 1.1 times your upload buffer size. For more information, see AWS Storage Gateway Limits (p. 296).

1. Decide the number and size of disks to allocate for your gateway.
2. Start your VMware vSphere client, and then connect to your host.
3. In the client, follow the instructions provided by your host's client and add two virtual disks.
4. Configure the disks according to the sizes you decided for your gateway.

   For your disk configuration, we recommend thick provisioning. We also strongly recommend that you choose the **Specify a datastore or datastore cluster** option and select a data store for these disks different from the one that is used for the gateway VM. In addition, we strongly recommend that you dedicate one data store for the disks that you will use for the upload buffer and a different data store for other disks.
5. Configure your VM to use paravirtualized controllers.

This procedure provides minimal instructions to help you to quickly allocate disks for your gateway. If you are not familiar with using the VMware vSphere client software, see Provisioning Local Disk Storage for the Gateway VM (VMWare) (p. 218) for detailed instructions.

Provisioning a Hyper-V Host

Following, you can find how to provision an on-premises Microsoft Hyper-V host and deploy your gateway. You perform this task in the Hyper-V Manager and return to the AWS Storage Gateway console to continue your gateway setup.

To provision a Hyper-V host

1. Review the minimum host requirements in Requirements (p. 11).
2. Set up a host in your data center with the Microsoft Hyper-V host. For a minimum set of instructions to install the hypervisor, see Configuring a Hyper-V Host for AWS Storage Gateway (p. 227).

To deploy your gateway on a Hyper-V host

1. Connect to the Microsoft Hyper-V Manager on your Windows client.
2. Create locations on the hypervisor host for the gateway virtual hard disks and VM.
   a. Navigate to the hypervisor drive.
   b. Create a folder with two subfolders, `unzippedSourceVM` and `gateway`.
3. Configure the Hyper-V Manager to point to the `gateway` folder you created. The running VM stores its configuration in this folder.
4. Copy the unzipped source VM files to the folder you created on the host computer. Import the AWS Storage Gateway VM to the host. You must have 75 GiB of disk space for installation of the VM image and system data.
5. Rename the VM to avoid confusion with other VMs that you might import to the host.
6. Confirm that **Time synchronization** is selected for the VM.

   You must also ensure that the clock on your gateway host is synchronized with a Network Time Protocol (NTP) server.

This procedure provides minimal instructions to help you deploy your gateway quickly. If you are not familiar with using the Microsoft Hyper-V manager software, see Deploying a AWS Storage Gateway VM on a Microsoft Hyper-V Host (p. 236) for detailed instructions.

Next, you allocate local disks for your deployed gateway VM. After you activate your gateway, you configure the disk as local storage for the gateway's use.

To allocate a local disk from direct-attached storage

> **Important**
> All gateways require a disk of a minimum size of 150 GiB for use as an upload buffer. Gateway-cached and tape gateway setups require additional disks for use as cache storage. The minimum size of a cache disk is 1.1 times your upload buffer size. For more information, see AWS Storage Gateway Limits (p. 296)

1. Decide the number and size of disks to allocate for your gateway.
2. Start the Microsoft Hyper-V Manager, and then connect to the hypervisor.
3. In the Hyper-V Manager, follow the instructions provided by your host's client and add two virtual disks.
4. Configure the disks according to the sizes that you decided for your gateway.

   For your disk configuration, we recommend thick provisioning. We also strongly recommend that you choose the **Specify a datastore or datastore cluster** option and select a data store for these disks different from the one that is used for the gateway VM. In addition, we strongly recommend that you dedicate one data store for the disks that you will use for the upload buffer and a different data store for other disks.
5. > **Note**
   > AWS Storage Gateway supports the .vhdx file type. This file type enables you to create larger virtual disks than other file types. If you create a .vhdx type virtual disk, make sure that the size of the virtual disks you create does not exceed the recommended disk size for your gateway.

   If you are using the Microsoft Hyper-V 2012 Hypervisor, you will be prompted to choose a disk format (**VHD** or **VHDX**).
6. Configure the disks as **Fixed size** for **Disk Type**.

   When you use fixed-size provisioning, the disk storage is allocated immediately, resulting in better performance. If you don’t use fixed-size provisioning, the storage is allocated on demand. On-demand allocation can have a negative impact on the functioning of AWS Storage Gateway. For AWS Storage Gateway to function properly, the VM disks must be stored in fixed-size provisioned format.

This procedure provides minimal instructions to help you quickly allocate disks. If you are not familiar with using the Microsoft Hyper-V manager software, see Provision Local Storage for the AWS Storage Gateway VM (Hyper-V) (p. 245) for detailed instructions.

Provisioning an Amazon EC2 Host

You can deploy and activate a gateway on an Amazon EC2 instance. Gateways hosted on Amazon EC2 instances can be used for disaster recovery and data mirroring. AWS Storage Gateway is available as an Amazon Machine Image (AMI) that contains the gateway VM image in the AWS Marketplace for AWS Storage Gateway.

Provisioning an Amazon EC2 Host by Using an AMI

You can host and activate a gateway on an Amazon EC2 instance. Gateways hosted on Amazon EC2 instances can be used for disaster recovery and data mirroring. An Amazon Machine Image (AMI) is available in the AWS Marketplace for AWS Storage Gateway.

To deploy a gateway on an Amazon EC2 instance

1. On the **Select host platform** page, choose **Amazon EC2**.
2. Choose **Launch with AWS Marketplace**. You are redirected to AWS Marketplace where you launch the EC2 AMI.

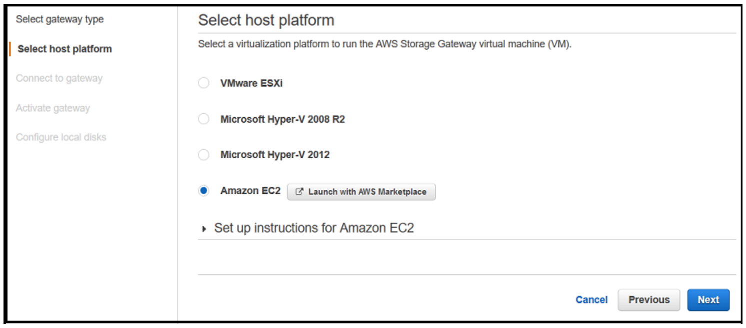

Select host platform

Select a virtualization platform to run the AWS Storage Gateway virtual machine (VM).

○ VMware ESXi

○ Microsoft Hyper-V 2008 R2

○ Microsoft Hyper-V 2012

● Amazon EC2 Launch with AWS Marketplace

▸ Set up instructions for Amazon EC2

3. On AWS Marketplace, choose **Continue**.
4. Choose **1-Click Launch**. Doing this launches the AMI with default settings.
5. If this is your first time using an AWS Storage Gateway AMI, choose **Accept Terms** to accept the terms of service.
6. Review the default settings. You can accept and use these default settings or modify them to meet your needs.

   The 1-Click Launch feature comes with an autogenerated security group that is named AWS Storage Gateway-1-0-AutogenByAWSMP. For information about security group settings, see Configuring Security Groups for Your Amazon EC2 Gateway Instance (p. 15).
7. After reviewing all your settings, choose **Launch with 1-Click**.
8. Choose **Return to Product Page** and locate your instance on the Amazon EC2 console.

   Important
   EC2 instances launched with the 1-Click Launch functionality come with one root Amazon EBS volume. You need to add additional EBS volumes to your instance as a separate step after the instance is launched. For information about how to add EBS volumes, see Attaching an Amazon EBS Volume to an Instance.

9. In the Amazon EC2 console, choose your Amazon EC2 instance, choose the **Description** tab at the bottom, and then note the IP address. You will use this IP address to connect to the gateway.

Connecting to Your Gateway

For gateways deployed and activated on an on-premises host, you can get the IP address from your gateway VM local console or your hypervisor client. For gateways deployed and activated on an Amazon EC2 instance, you can get the IP address from the Amazon EC2 console. The activation process associates your gateway with your AWS account. Your gateway VM must be running for activation to succeed.

To get the IP address for you gateway VM from the local console

1. Log on to your gateway VM local console. For detailed instructions, see Configuring Your Gateway Network (p. 162).
2. Get the IP address from the top of the menu page.
3. Take note of the gateway IP address.

To get the IP address from an EC2 instance

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.
3. Choose the **Description** tab at the bottom, and then note the IP address. You will use this IP address to activate the gateway.

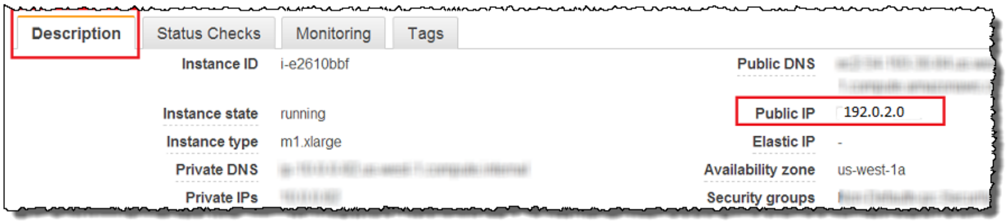

Instance ID i-e2610bbf
Instance state running
Instance type m1.xlarge
Private DNS [redacted]
Private IPs [redacted]
Public DNS [redacted]
Public IP 192.0.2.0
Elastic IP -
Availability zone us-west-1a
Security groups [redacted]

For activation, you can use the public or private IP address assigned to the gateway. You must be able to reach the IP address that you use from the browser from which you perform the activation. In this walkthrough, we use the public IP address to activate the gateway.

To associate your gateway with your AWS account

1. If the **Connect to gateway** page isn't already open, open it and navigate to the **Connect to gateway** page.
2. Type the IP address of your gateway for **IP address**, and then choose **Connect gateway**.

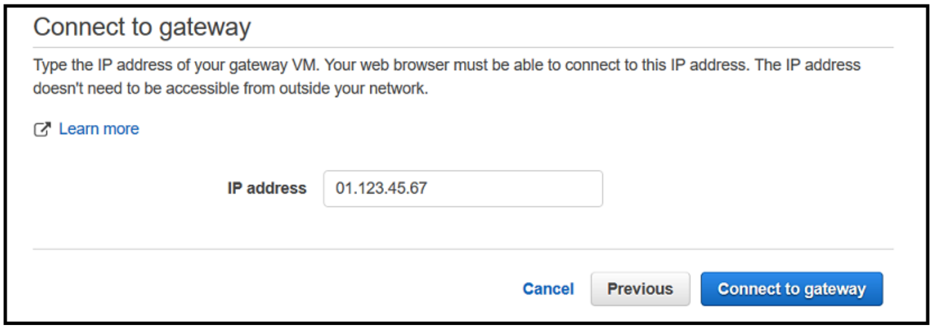

Connect to gateway

Type the IP address of your gateway VM. Your web browser must be able to connect to this IP address. The IP address doesn't need to be accessible from outside your network.

IP address 01.123.45.67

Activating Your Gateway

When your gateway VM is deployed and running, you configure your gateway settings and activate your gateway.

To activate your gateway

1. To complete the activation process, provide the information on the activation page to configure your gateway setting. The activation page you see depends on your gateway configuration.

   - **Gateway Time Zone** specifies the time zone to use for your gateway.
   - **Gateway Name** identifies your gateway. You use this name to manage your gateway in the console; you can change it after the gateway is activated. This name must be unique to your account.

     The following screenshot shows the activation page for volume gateways.

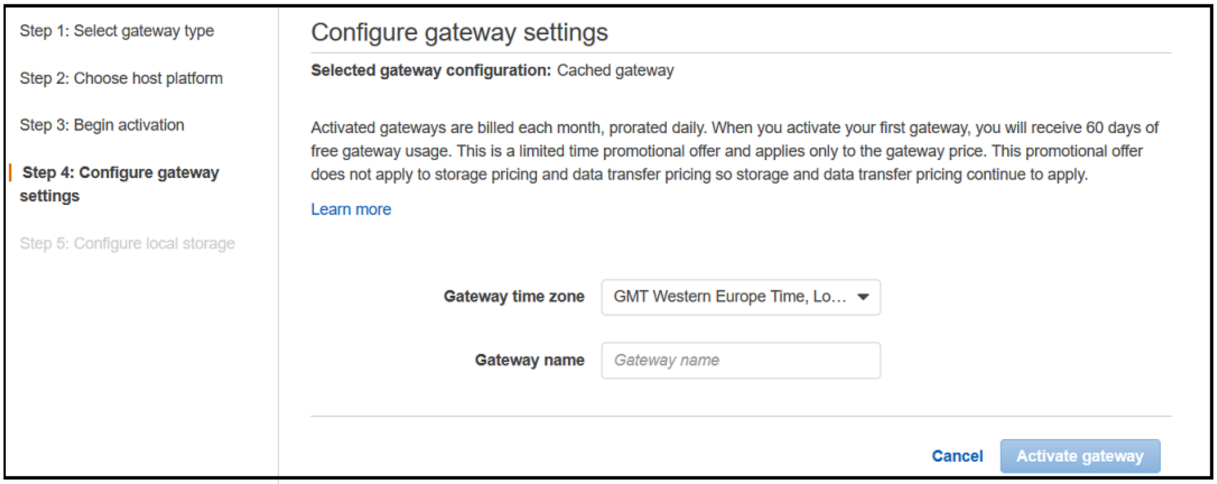

Configure gateway settings

Selected gateway configuration: Cached gateway

Activated gateways are billed each month, prorated daily. When you activate your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies only to the gateway price. This promotional offer does not apply to storage pricing and data transfer pricing so storage and data transfer pricing continue to apply.

Gateway time zone: GMT Western Europe Time, Lo...

Gateway name: Gateway name

Step 1: Select gateway type
Step 2: Choose host platform
Step 3: Begin activation
Step 4: Configure gateway settings
Step 5: Configure local storage
Learn more
Cancel
Activate gateway

2. Choose **Activate Gateway**.
3. When the gateway is successfully activated, the AWS Storage Gateway console displays the **Configure local disks** page. If activation fails, check that the IP address you entered is correct. If the IP address is correct, confirm that your network is configured to let your browser access the gateway VM. For more information, see Troubleshooting Your Gateway (p. 191) for possible solutions.

Configuring Local Disks

When you deployed the VM, you allocated local disks for your gateway. Now, you will configure your gateway to use these disks:

- For cached volumes (p. 4), you configure at least one disk for an upload buffer and the other for cache storage.
- For stored volumes (p. 5), you configure at least one disk for an upload buffer and allocate the rest of the storage for your application data.

To configure local disks

1. On the **Configure local disks** page, identify the disks you allocated and decide which ones you want to use for upload buffer and cached storage. For information about disk size limits, see Configuration and Performance Limits (p. 297).

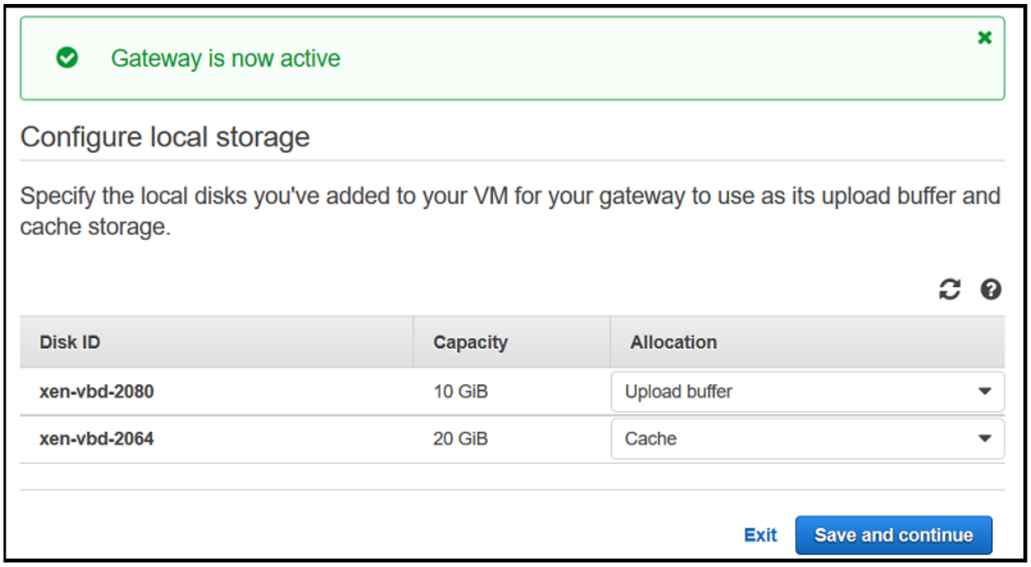

Gateway is now active

Configure local storage

Specify the local disks you've added to your VM for your gateway to use as its upload buffer and cache storage.

Disk ID	Capacity	Allocation
xen-vbd-2080	10 GiB	Upload buffer
xen-vbd-2064	20 GiB	Cache

2. From the list next to your upload buffer disk, choose **Upload Buffer**.
3. For cached volumes and tapes, choose **Cache** for the disk you want to configure as a cache storage.

   If you don't see your disks, choose **Refresh**.
4. Choose **Save and continue** to save your configuration settings.

For stored volumes, you configure one of the two disks for use by your application's data and the other disk as an upload buffer, and also optionally create alarms to monitor your disks.

Creating Volumes

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

To create a volume

1. On AWS Storage Gateway console, choose **Create volume**.
2. In the **Create volume** dialog box, choose a gateway from the **Gateway** list.
3. For the cached volumes, type the capacity in the**Capacity**.

   For stored volumes, select the **Disk ID** from the list.
4. For **Volume content**, specify whether you want to create a new empty volume or create a volume based on a snapshot and type the snapshot ID.

   For stored volumes, you can also preserve existing data on the disk. For more information, see Restoring a Snapshot to a Storage Volume (p. 262).

   Caution
   Make sure you don't have any existing data on the virtual disk. Any existing data on the disk will be lost.
5. Type a name in the **iSCSI target name** box.

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

6. Verify that the **Network interface** setting is the IP address of your gateway, and then choose **Create volume**. The **Configure CHAP Authentication** dialog box appears.

Configure CHAP Authentication for Your Volumes

You can configure Challenge-Handshake Authentication Protocol (CHAP) for your volume at this point, or you can choose **Cancel** and configure CHAP later.

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

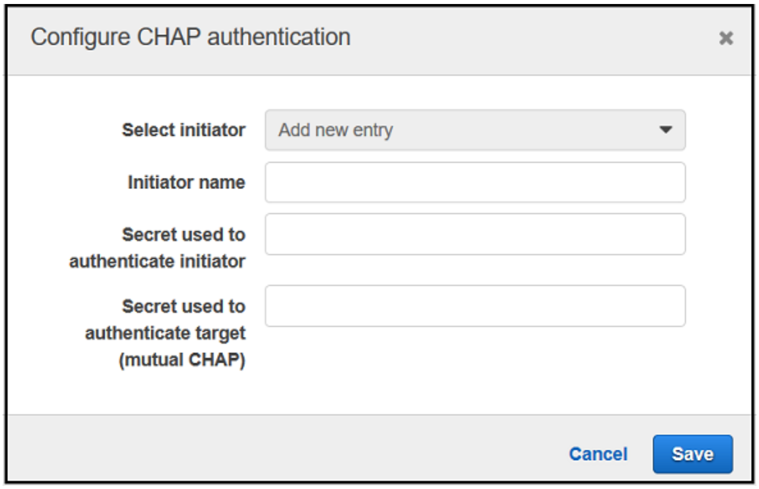
Configure CHAP authentication
×
Initiator name
Secret used to authenticate initiator
Secret used to authenticate target (mutual CHAP)
Cancel
Save

Select initiator: Add new entry

To configure CHAP

1. Select the volume for which you want to configure CHAP.
2. On the **Action** menu, choose **Configure CHAP Authentication**.
3. For **Initiator Name**, type the name of your initiator.
4. For **Secret Used to Authenticate Initiator**, type the secret phrase you used to authenticate your iSCSI initiator.
5. For **Secret Used to Authenticate Target (Mutual CHAP)**, type the secret phrase used to authenticate your target for mutual CHAP.
6. Choose **Save** to save your entries.

   For more information about setting up CHAP authentication, see Configuring CHAP Authentication for Your iSCSI Targets (p. 281).

Using Your Volumes

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

Topics

Microsoft Windows Client Volumes

Following, you can find instructions on how to connect and use your volumes on a Windows client.

Topics

Connecting Your Volumes to Your Microsoft Windows Client

You use the Microsoft Windows iSCSI initiator to connect to your volumes. At the end of the following procedure, the volumes become available as local devices on your Windows client. For instructions on accessing the iSCSI storage volume from Linux, see Recommended Red Hat Linux iSCSI Settings (p. 44).

> **Important**
> AWS Storage Gateway enables applications that are clustered using Windows Server Failover Clustering (WSFC) to use the iSCSI initiator to access your gateway's volumes. However, connecting multiple hosts to the same iSCSI target is not supported.

To connect your Windows client to a storage volume

1. On the **Start** menu of your Windows client computer, type **`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.

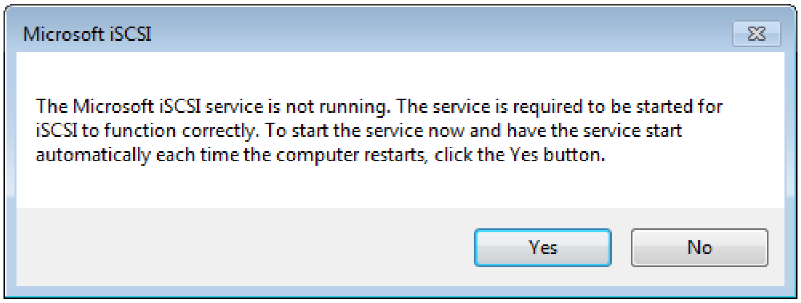

Microsoft iSCSI

The Microsoft iSCSI service is not running. The service is required to be started for iSCSI to function correctly. To start the service now and have the service start automatically each time the computer restarts, click the Yes button.

3. In the **iSCSI Initiator Properties** dialog box, choose the **Discovery** tab, and then choose the **Discovery Portal** button.

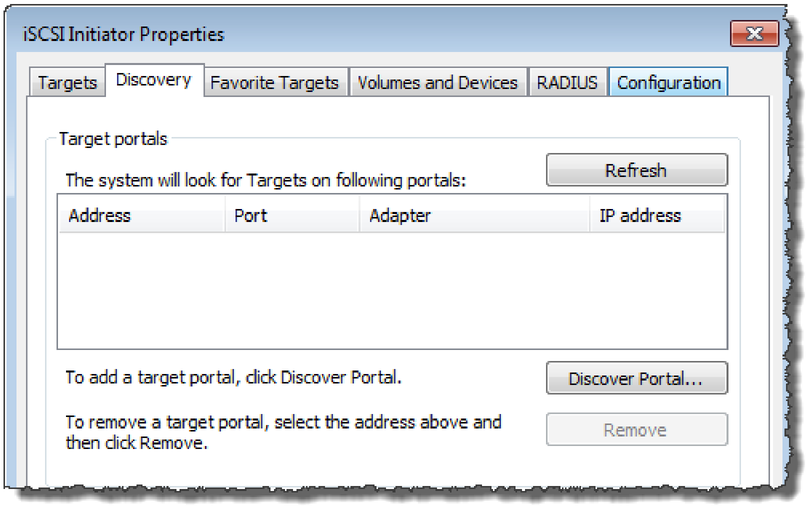

The system will look for Targets on following portals:

Address	Port	Adapter	IP address

To add a target portal, click Discover Portal.

To remove a target portal, select the address above and then click Remove.

4. In the **Discover Target Portal** dialog box, type the IP address of your iSCSI target for **IP address or DNS name**, and then choose **OK**. To get the IP address of your gateway, check the **Gateway** tab on the AWS 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.

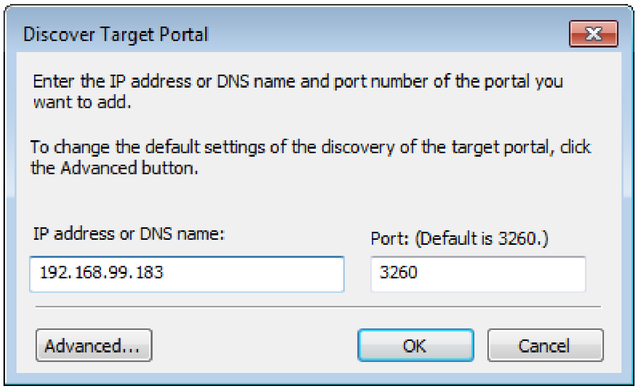

Enter the IP address or DNS name and port number of the portal you want to add.

To change the default settings of the discovery of the target portal, click the Advanced button.

IP address or DNS name: 192.168.99.183

Port: (Default is 3260.) 3260

The IP address now appears in the **Target portals** list on the **Discovery** tab.

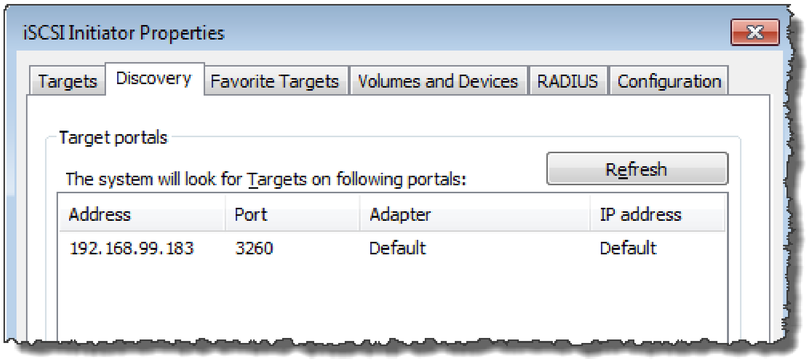

The system will look for Targets on following portals:

Address	Port	Adapter	IP address
192.168.99.183	3260	Default	Default

5. Connect the new target portal to the storage volume target on the gateway:

   a. Choose the **Targets** tab.

The new target portal is shown with an inactive status. Note that the target name shown should be the same as the name you specified for your storage volume in step 1.

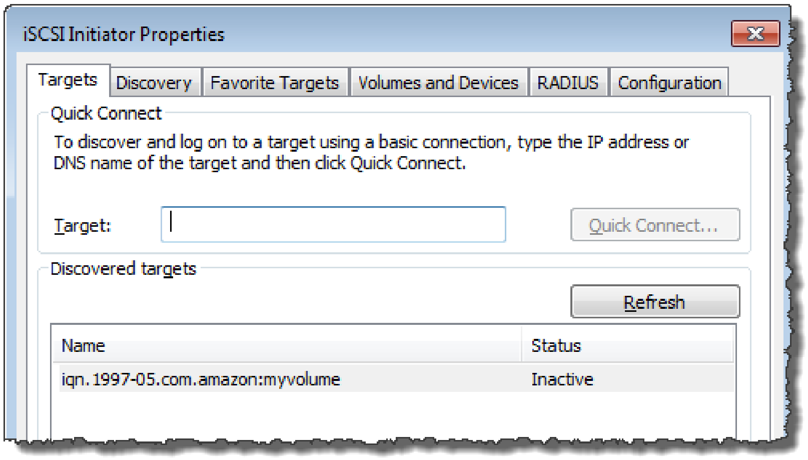

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Name	Status
iqn.1997-05.com.amazon:myvolume	Inactive

b. Select the target, and then choose **Connect**.

If the target name is not populated already, type the name of the target as shown in step 1 in the **Connect to Target** dialog box, select the check box next to **Add this connection to the list of Favorite Targets**, and then choose **OK**.

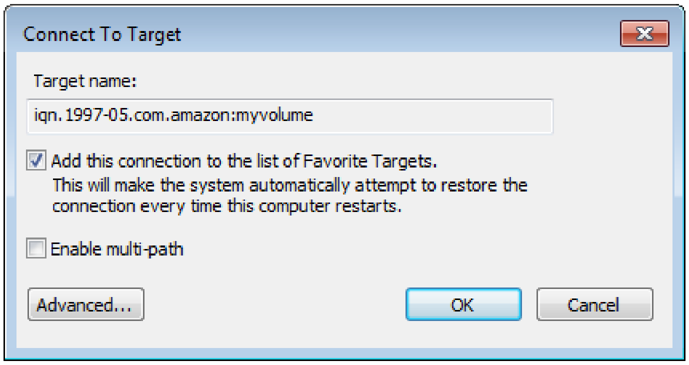

Target name:
iqn.1997-05.com.amazon:myvolume

☑ Add this connection to the list of Favorite Targets.
This will make the system automatically attempt to restore the connection every time this computer restarts.

□ Enable multi-path

c. In the **Targets** tab, ensure that the target **Status** has the value **Connected** indicating the target is connected, and then choose **OK**.

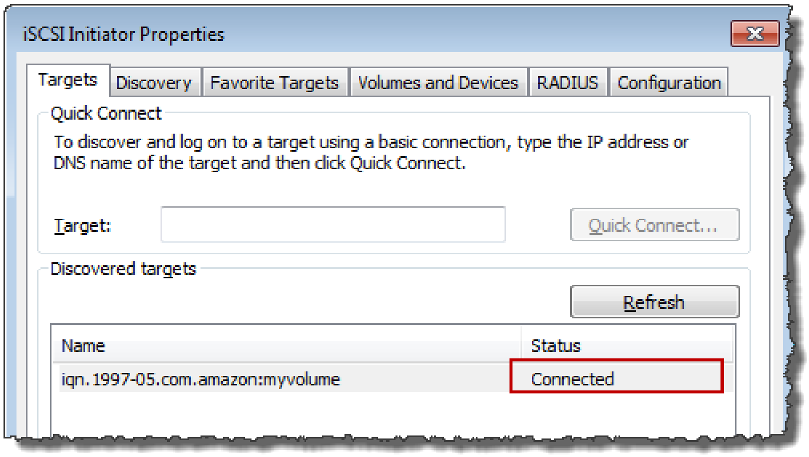

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Name	Status
iqn.1997-05.com.amazon:myvolume	Connected

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

> **Note**
> Although it is not required for this exercise, we highly recommend that you customize your iSCSI settings for a real-world application as discussed in Customizing Your Windows iSCSI Settings (p. 280).

To initialize and format the storage volume you just mapped

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

Partition Style	Use in the Following Conditions
MBR (Master Boot Record)	• If your gateway is a stored volume and the storage volume is limited to 1 TiB in size. • If your gateway is a cached volume and the storage volume is less than 2 TiB in size.
GPT (GUID Partition Table)	If your gateway's storage volume is 2 TiB or greater in size.

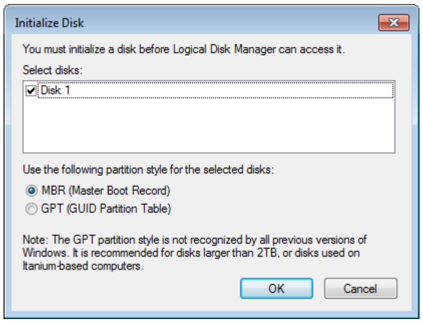

You must initialize a disk before Logical Disk Manager can access it.

Select disks:

☑ Disk 1

Use the following partition style for the selected disks:

● MBR (Master Boot Record)
○ GPT (GUID Partition Table)

Note: The GPT partition style is not recognized by all previous versions of Windows. It is recommended for disks larger than 2TB, or disks used on Itanium-based computers.

3. Create a simple volume:

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

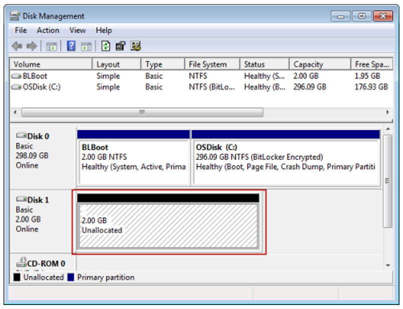

Volume	Layout	Type	File System	Status	Capacity	Free Spa...
BLBoot	Simple	Basic	NTFS	Healthy (S...	2.00 GB	1.95 GB
OSDisk (C:)	Simple	Basic	NTFS (BitLo...	Healthy (B...	296.09 GB	176.93 GB

Disk 0
Basic
298.09 GB
Online

BLBoot
2.00 GB NTFS
Healthy (System, Active, Prima

OSDisk (C:)
296.09 GB NTFS (BitLocker Encrypted)
Healthy (Boot, Page File, Crash Dump, Primary Partiti

Disk 1
Basic
2.00 GB
Online

2.00 GB
Unallocated

CD-ROM 0

   b. Open the context (right-click) menu for the disk, and then choose **New Simple Volume**.

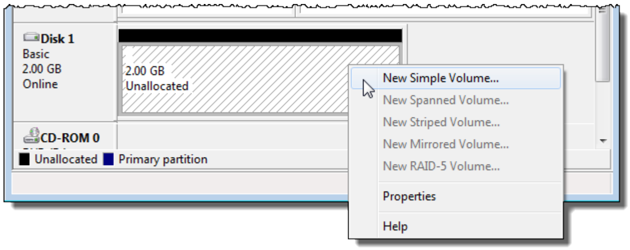

Important
Be careful not to format the wrong disk. Check to ensure that the disk you are formatting matches the size of the local disk you allocated to the gateway VM and that it has a status of **Unallocated**.

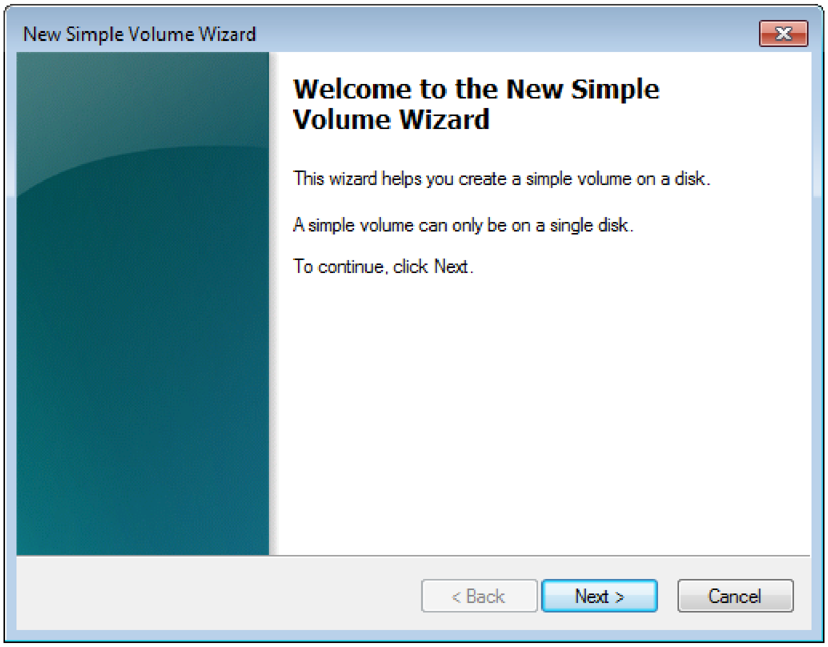

Welcome to the New Simple Volume Wizard

This wizard helps you create a simple volume on a disk.

A simple volume can only be on a single disk.

To continue, click Next.

c. In the **New Simple Volume Wizard**, choose **Next**.

d. In the **Specify Volume Size** dialog box, keep the default values, and then choose **Next**.

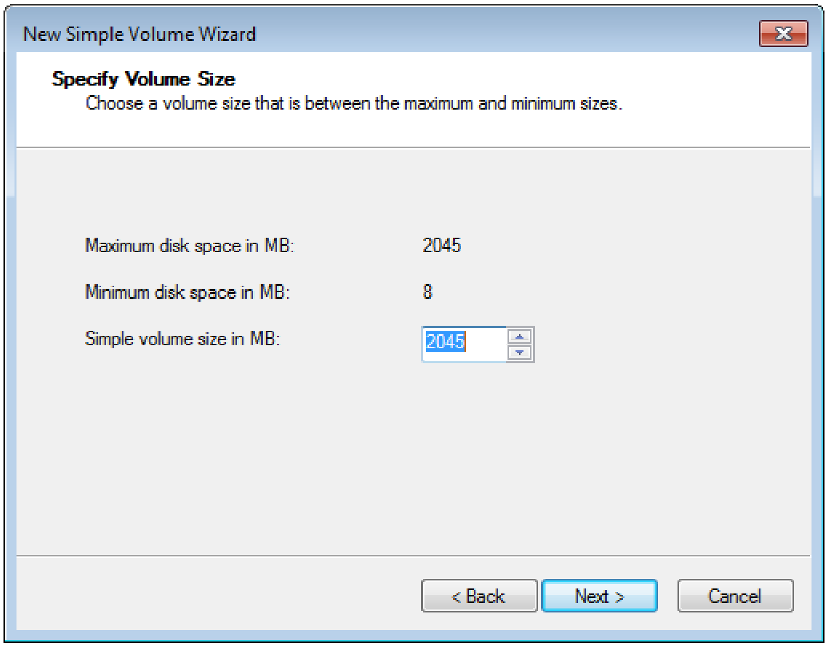

New Simple Volume Wizard

Specify Volume Size

Choose a volume size that is between the maximum and minimum sizes.

Maximum disk space in MB:	2045
Minimum disk space in MB:	8
Simple volume size in MB:	2045

e. In the **Assign Drive Letter or Path** dialog box, keep the default values, and then choose **Next**.

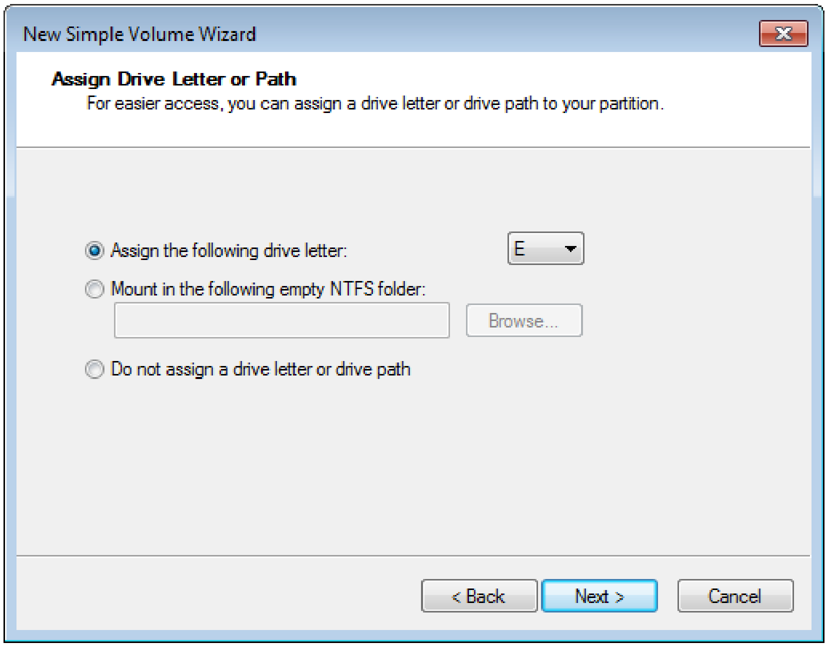

New Simple Volume Wizard

Assign Drive Letter or Path

For easier access, you can assign a drive letter or drive path to your partition.

● Assign the following drive letter: E

○ Mount in the following empty NTFS folder:

○ Do not assign a drive letter or drive path

f. In the **Format Partition** dialog box, type a label for **Volume label**, and ensure that **Perform a quick format** is selected. Choose **Next**.

> **Caution**
>
> Selecting **Perform a quick format** is highly recommended for cached volumes because it results in less initialization I/O, smaller initial snapshot size, and the fastest time to a usable volume. It also avoids cached volume usage that comes from the full format process and not from application data activity.

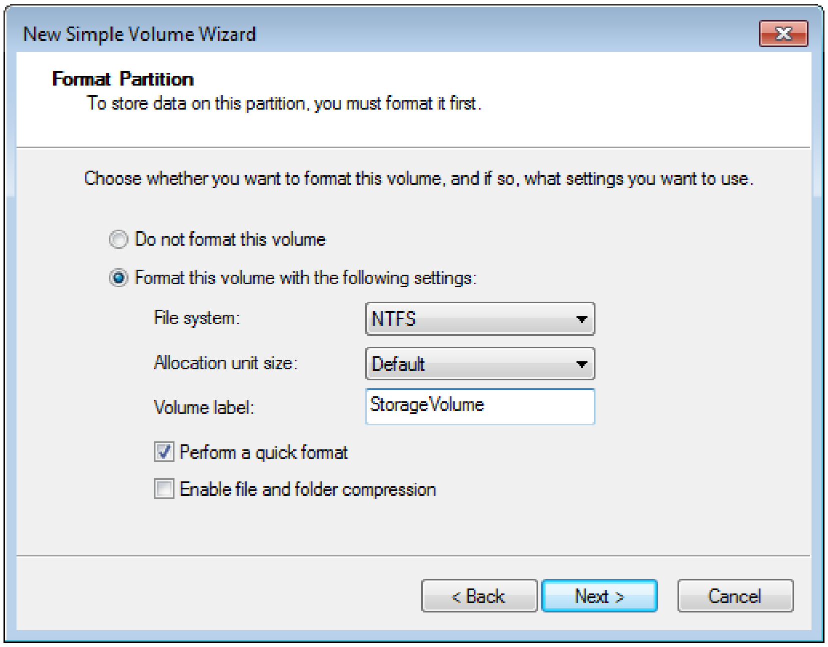
New Simple Volume Wizard

Format Partition
To store data on this partition, you must format it first.

Choose whether you want to format this volume, and if so, what settings you want to use.

○ Do not format this volume
● Format this volume with the following settings:

File system: NTFS
Allocation unit size: Default
Volume label: StorageVolume

☑ Perform a quick format
□ Enable file and folder compression

g. Choose **Finish** to close the wizard.

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

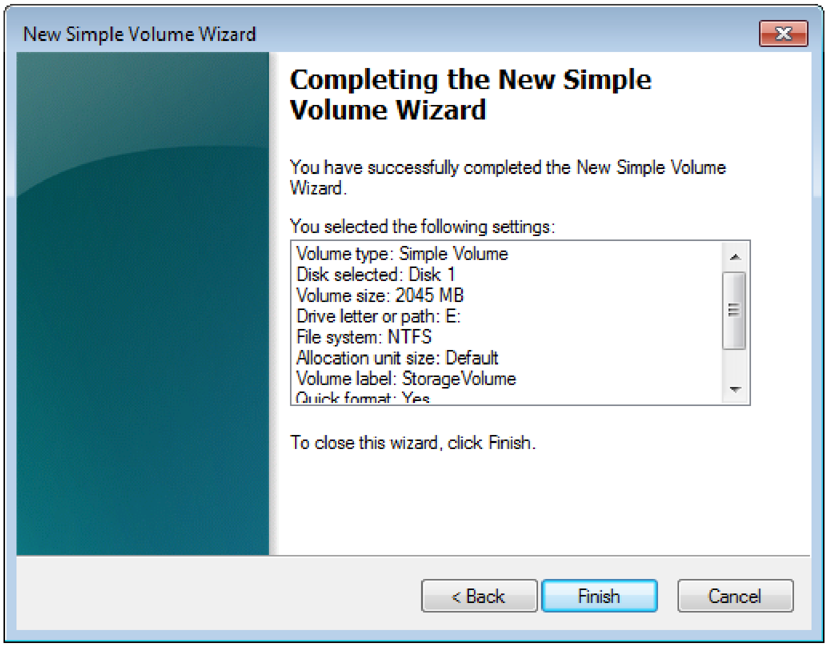
New Simple Volume Wizard

Completing the New Simple Volume Wizard

You have successfully completed the New Simple Volume Wizard.

You selected the following settings:

Volume type: Simple Volume
Disk selected: Disk 1
Volume size: 2045 MB
Drive letter or path: E:
File system: NTFS
Allocation unit size: Default
Volume label: StorageVolume
Quick format: Yes

To close this wizard, click Finish.

Testing Your Gateway on Microsoft Windows

You test your volume gateway setup by performing the following tasks:

1. Write data to the volume.

2. Take a snapshot.
3. Restore the snapshot to another volume.

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

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

To create a snapshot of a storage volume on Microsoft Windows

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

   The amount of data copied doesn't matter for this demonstration. A small file is enough to demonstrate the restore process.
2. In the navigation pane of the AWS Storage Gateway console, choose **Volumes**.
3. Select the storage volume that you created for the gateway.

   This gateway should have only one storage volume. Selecting the volume displays its properties.
4. For **Action**, choose **Create Snapshot** to create a snapshot of the volume.

   Depending on the amount of data on the disk and the upload bandwidth, it might take a few seconds to complete the snapshot. Note the volume ID for the volume from which you create a snapshot. You will use the ID to find the snapshot.
5. In the **Create Snapshot** dialog box, provide a description for your snapshot, and then choose **Create Snapshot**.

   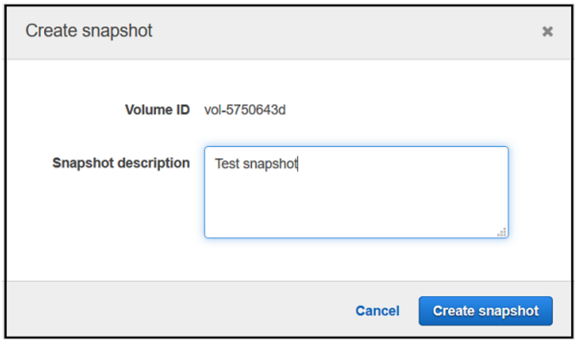

   Create snapshot

   Volume ID vol-5750643d

   Snapshot description: Test snapshot

   Your snapshot is stored as an Amazon EBS snapshot. Take note of your snapshot ID.

   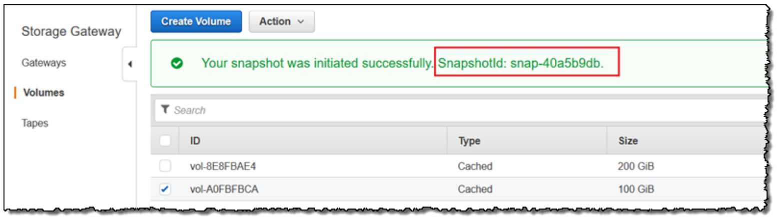

   Your snapshot was initiated successfully. SnapshotId: snap-40a5b9db.

	ID	Type	Size
□	vol-8E8FBAE4	Cached	200 GiB
☑	vol-A0FBFBCA	Cached	100 GiB

   The number of snapshots created for your volume is displayed in the snapshot column.
6. For **Snapshot**, choose the link for the volume you created the snapshot for to see your EBS snapshot on the Amazon EC2 console.

You can restore a snapshot to a new storage location. For more information, see Restoring a Snapshot to a Storage Volume (p. 262).

Recommended iSCSI Settings

After setting up your initiator, we highly recommend that you customize your iSCSI settings to prevent the initiator from disconnecting from targets.

Topics

Recommended Microsoft Windows iSCSI Settings

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*.

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 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.

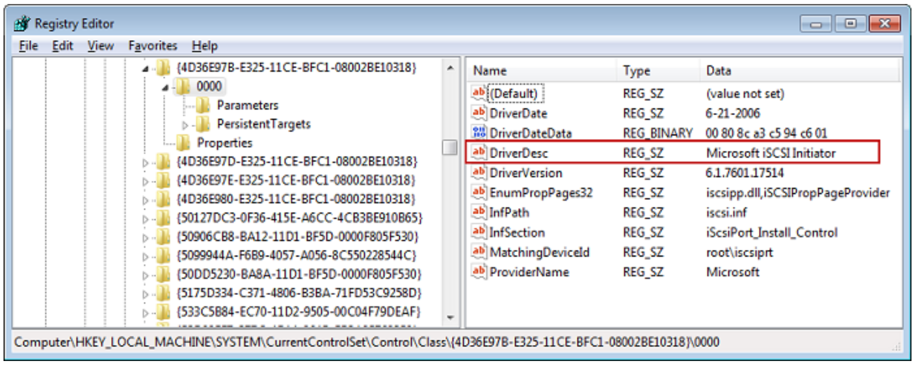

```
{4D36E97B-E325-11CE-BFC1-08002BE10318}
    0000
        Parameters
        PersistentTargets
    Properties
{4D36E97D-E325-11CE-BFC1-08002BE10318}
{4D36E97E-E325-11CE-BFC1-08002BE10318}
{4D36E980-E325-11CE-BFC1-08002BE10318}
{50127DC3-0F36-415E-A6CC-4CB3BE910B65}
{50906CB8-BA12-11D1-BF5D-0000F805F530}
{5099944A-F6B9-4057-A056-8C550228544C}
{50DD5230-BA8A-11D1-BF5D-0000F805F530}
{5175D334-C371-4806-B3BA-71FD53C9258D}
{533C5B84-EC70-11D2-9505-00C04F79DEAF}
```

Name	Type	Data
(Default)	REG_SZ	(value not set)
DriverDate	REG_SZ	6-21-2006
DriverDateData	REG_BINARY	00 80 8c a3 c5 94 c6 01
DriverDesc	REG_SZ	Microsoft iSCSI Initiator
DriverVersion	REG_SZ	6.1.7601.17514
EnumPropPages32	REG_SZ	iscsipp.dll,iSCSIPropPageProvider
InfPath	REG_SZ	iscsi.inf
InfSection	REG_SZ	iScsiPort_Install_Control
MatchingDeviceId	REG_SZ	root\iscsiprt
ProviderName	REG_SZ	Microsoft

Computer\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\0000

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.

   This value represents a hold time of 600 seconds. The example following shows the **MaxRequestHoldTime** DWORD value with a value of 600.

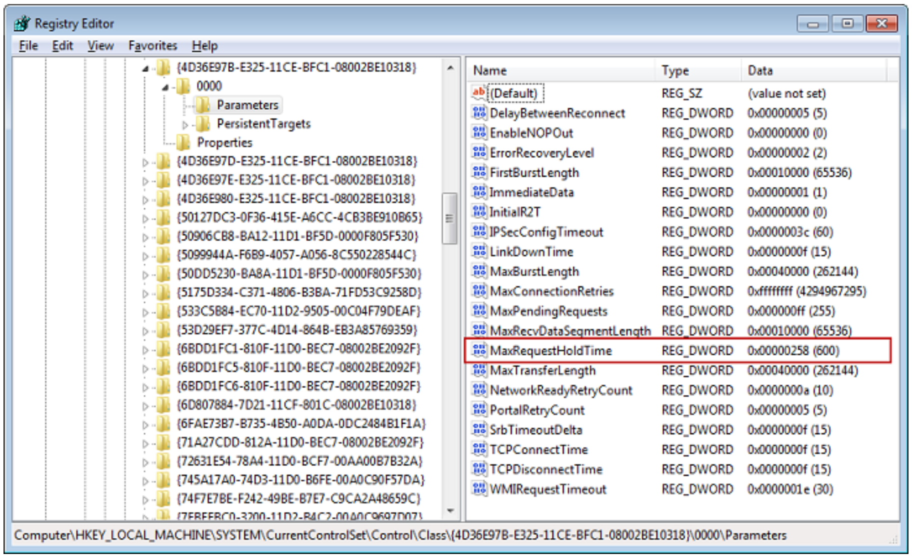

```
{4D36E97B-E325-11CE-BFC1-08002BE10318}
    0000
        Parameters
        PersistentTargets
    Properties
{4D36E97D-E325-11CE-BFC1-08002BE10318}
{4D36E97E-E325-11CE-BFC1-08002BE10318}
{4D36E980-E325-11CE-BFC1-08002BE10318}
{50127DC3-0F36-415E-A6CC-4CB3BE910B65}
{50906CB8-BA12-11D1-BF5D-0000F805F530}
{5099944A-F6B9-4057-A056-8C550228544C}
{50DD5230-BA8A-11D1-BF5D-0000F805F530}
{5175D334-C371-4806-B3BA-71FD53C9258D}
{533C5B84-EC70-11D2-9505-00C04F79DEAF}
{53D29EF7-377C-4D14-864B-EB3A85769359}
{6BDD1FC1-810F-11D0-BEC7-08002BE2092F}
{6BDD1FC5-810F-11D0-BEC7-08002BE2092F}
{6BDD1FC6-810F-11D0-BEC7-08002BE2092F}
{6D807884-7D21-11CF-801C-08002BE10318}
{6FAE73B7-B735-4B50-A0DA-0DC2484B1F1A}
{71A27CDD-812A-11D0-BEC7-08002BE2092F}
{72631E54-78A4-11D0-BCF7-00AA00B7B32A}
{745A17A0-74D3-11D0-B6FE-00A0C90F57DA}
{74F7E7BE-F242-49BE-B7E7-C9CA2A48659C}
```

Name	Type	Data
(Default)	REG_SZ	(value not set)
DelayBetweenReconnect	REG_DWORD	0x00000005 (5)
EnableNOPOut	REG_DWORD	0x00000000 (0)
ErrorRecoveryLevel	REG_DWORD	0x00000002 (2)
FirstBurstLength	REG_DWORD	0x00010000 (65536)
ImmediateData	REG_DWORD	0x00000001 (1)
InitialR2T	REG_DWORD	0x00000000 (0)
IPSecConfigTimeout	REG_DWORD	0x0000003c (60)
LinkDownTime	REG_DWORD	0x0000000f (15)
MaxBurstLength	REG_DWORD	0x00040000 (262144)
MaxConnectionRetries	REG_DWORD	0xffffffff (4294967295)
MaxPendingRequests	REG_DWORD	0x000000ff (255)
MaxRecvDataSegmentLength	REG_DWORD	0x00010000 (65536)
MaxRequestHoldTime	REG_DWORD	0x00000258 (600)
MaxTransferLength	REG_DWORD	0x00040000 (262144)
NetworkReadyRetryCount	REG_DWORD	0x0000000a (10)
PortalRetryCount	REG_DWORD	0x00000005 (5)
SrbTimeoutDelta	REG_DWORD	0x0000000f (15)
TCPConnectTime	REG_DWORD	0x0000000f (15)
TCPDisconnectTime	REG_DWORD	0x0000000f (15)
WMIRequestTimeout	REG_DWORD	0x0000001e (30)

Computer\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\0000\Parameters

2. 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.

   This value represents a timeout of 600 seconds. The example following shows the **TimeOutValue** DWORD value with a value of 600.

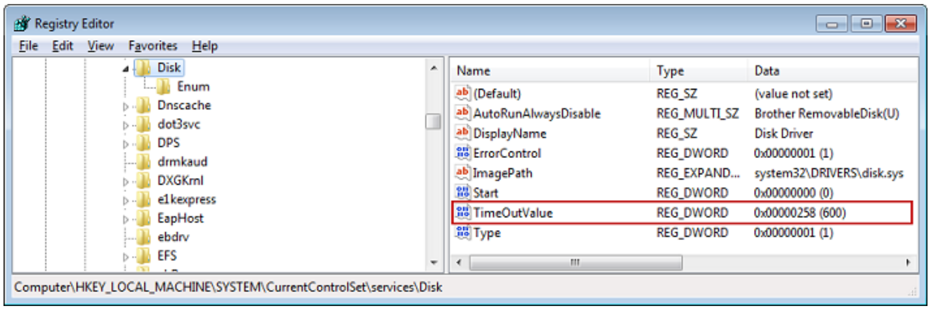

Name	Type	Data
(Default)	REG_SZ	(value not set)
AutoRunAlwaysDisable	REG_MULTI_SZ	Brother RemovableDisk(U)
DisplayName	REG_SZ	Disk Driver
ErrorControl	REG_DWORD	0x00000001 (1)
ImagePath	REG_EXPAND...	system32\DRIVERS\disk.sys
Start	REG_DWORD	0x00000000 (0)
TimeOutValue	REG_DWORD	0x00000258 (600)
Type	REG_DWORD	0x00000001 (1)

Computer\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\Disk

3. 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.

Recommended Red Hat Linux iSCSI Settings

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*, then 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 /etc/init.d/iscsi 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 AWS 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 AWS 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:myvolume --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
   ```

   After setting up your initiator, we highly recommend that you customize your iSCSI settings as discussed in Customizing Your Linux iSCSI Settings (p. 45).

Customizing Your Linux iSCSI Settings

After setting up your initiator, we highly recommend that you customize your iSCSI settings to prevent the initiator from disconnecting from targets. By increasing the iSCSI timeout values as shown following, you make your application better at dealing with write operations that take a long time and other transient issues such as network interruptions.

Note
Commands might be slightly different for other types of Linux. The following examples are based on Red Hat Linux.

To customize your Linux iSCSI settings

1. Increase the maximum time for which requests are queued.

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

      ```
      node.session.timeo.replacement_timeout = [replacement_timeout_value]
      node.conn[0].timeo.noop_out_interval = [noop_out_interval_value]
      node.conn[0].timeo.noop_out_timeout = [noop_out_timeout_value]
      ```

   b. Set the *`[replacement_timeout_value]`* value to 600.

      Set the *`[noop_out_interval_value]`* value to 60.

      Set the *`[noop_out_timeout_value]`* value to 600.

      All three values are in seconds.

      > **Note**
      > The `iscsid.conf` settings must be made before discovering the gateway. If you have already discovered your gateway or logged in to the target, or both, you can delete the entry from the discovery database using the following command. Then you can rediscover or log in again to pick up the new configuration.

      ```
      iscsiadm -m discoverydb -t sendtargets -p gateway_ip:3260 -o
       delete
      ```

2. Increase the disk timeout value in the rules file.

   a. 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'"
      ```

      > **Note**
      > This rules file does not exist in RHEL 6 or 7 initiators, so you must create it.

      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.

3. Restart your system to ensure that the new configuration values take effect.

   Before restarting, you must make sure that the results of all write operations to your volumes are flushed. To do this, unmount storage volumes before restarting.

4. 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.

Where Do I Go from Here?

Topics

The AWS Storage Gateway service provides an easy way for you to back your application storage with the storage infrastructure of the AWS Cloud. In Using Your Volumes (p. 33), you created and provisioned a gateway, and then connected your Windows host to the gateway's storage volume. You added data to the gateway's iSCSI volume, took a snapshot of the volume and restored it to a new volume, and connected to the new volume and verified that the data shows up on it.

After you finish the exercise, consider the following:

- If you plan on continuing to use your gateway, you should read about sizing the upload buffer more appropriately for real-world workloads. For more information, see Sizing Your Gateway's Storage for Real-World Workloads (p. 47).
- If you don't plan on continuing to use your gateway, consider deleting the gateway to avoid incurring any charges. For more information, see Clean Up Resources You Don't Need (p. 48).

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

- Learn more about storage volumes and how to manage them (see Managing Your Gateway (p. 99)).
- Troubleshoot gateway problems (see Troubleshooting Your Gateway (p. 191)).
- Optimize your gateway (see Optimizing Gateway Performance (p. 149)).
- Understand Storage Gateway metrics and how you can monitor how your gateway performs (see Monitoring Your Gateway (p. 120)).
- Connect to the gateway's iSCSI targets to store data (see Connecting to Volumes on Your Volume Gateway (p. 279)).

Sizing Your Gateway's Storage for Real-World Workloads

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

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

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

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

- Use the formula shown in Adding and Removing Upload Buffer (p. 144) for sizing the upload buffer. We strongly recommend that you allocate at least 150 GiB for the upload buffer. If the upload buffer formula yields a value less than 150 GiB, use 150 GiB as your allocated upload buffer.

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

To size your upload buffer for a stored setup

- Use the formula discussed in Adding and Removing Upload Buffer (p. 144). We strongly recommend that you allocate at least 150 GiB for your upload buffer. If the upload buffer formula yields a value less than 150 GiB, use 150 GiB as your allocated upload buffer.

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

To monitor your upload buffer

1. View your gateway's current upload buffer.
   - In the **Gateway** tab in the AWS Storage Gateway console, Choose the **Details** tab and find the **Upload Buffer Used** field.
2. Set one or more alarms to notify you about upload buffer use.

   We highly recommend that you create one or more upload buffer alarms in the CloudWatch console. For example, you can set an alarm for a level of use you want to be warned about and an alarm for a level of use that, if exceeded, is cause for action. The action might be adding more upload buffer space. For more information, see To set an upper threshold alarm for a gateway's upload buffer (p. 130).

Clean Up Resources You Don't Need

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

If you plan to continue using your gateway, see additional information in Where Do I Go from Here? (p. 47)

To clean up resources you don't need

1. Delete any snapshots. For instructions, see Deleting a Snapshot (p. 105).
2. Unless you plan to continue using the gateway, delete it. For more information, see Deleting Your Gateway by Using the AWS Storage Gateway Console and Removing Associated Resources (p. 186).
3. Delete the AWS Storage Gateway VM from your on-premises host. If you created your gateway on an Amazon EC2 instance, terminate the instance.

Creating a Tape Gateway

In this section, you can find instructions about how to create and use a tape gateway.

Topics

Creating a Gateway

In this section, you can find instructions about how to download, deploy, and activate a tape gateway.

Topics

Selecting a Gateway Type

For a Tape Gateway (p. 6), you store and archive your data on virtual tapes in AWS. A tape gateway eliminates some of the challenges associated with owning and operating an on-premises physical tape infrastructure.

To create a tape gateway

1. Open the AWS Console at http://console.aws.amazon.com/storagegateway/home, and choose the AWS Region you want to create your gateway in.

   If you have a gateway in this region, the console shows your gateway in the console.
2. In the navigation pane, choose **Gateways**, and then choose **Create gateway** to open the **Select gateway type** page.
3. On the **Select gateway type page**, choose **Tape gateway**, and then choose **Next**.

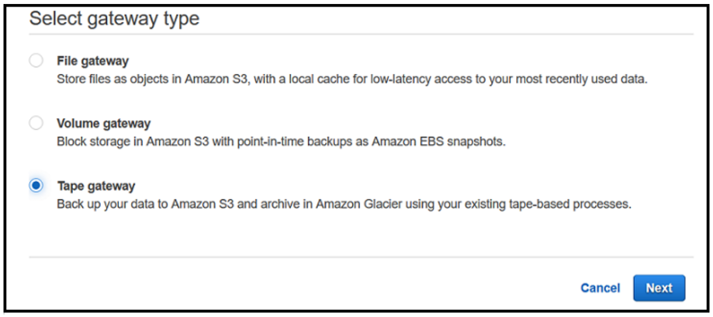

Select gateway type

○ **File gateway**
Store files as objects in Amazon S3, with a local cache for low-latency access to your most recently used data.

○ **Volume gateway**
Block storage in Amazon S3 with point-in-time backups as Amazon EBS snapshots.

● **Tape gateway**
Back up your data to Amazon S3 and archive in Amazon Glacier using your existing tape-based processes.

Choosing a Host Platform

If you are creating the gateway on-premises, you download and deploy the gateway VM and then activate the gateway. If you are creating the gateway on an Amazon EC2 instance, you launch an Amazon Machine Image (AMI) that contains the gateway VM image and then activate the gateway.

To choose a host platform

1. On the **Select host platform** page, choose the virtualization platform you want to run your gateway on.

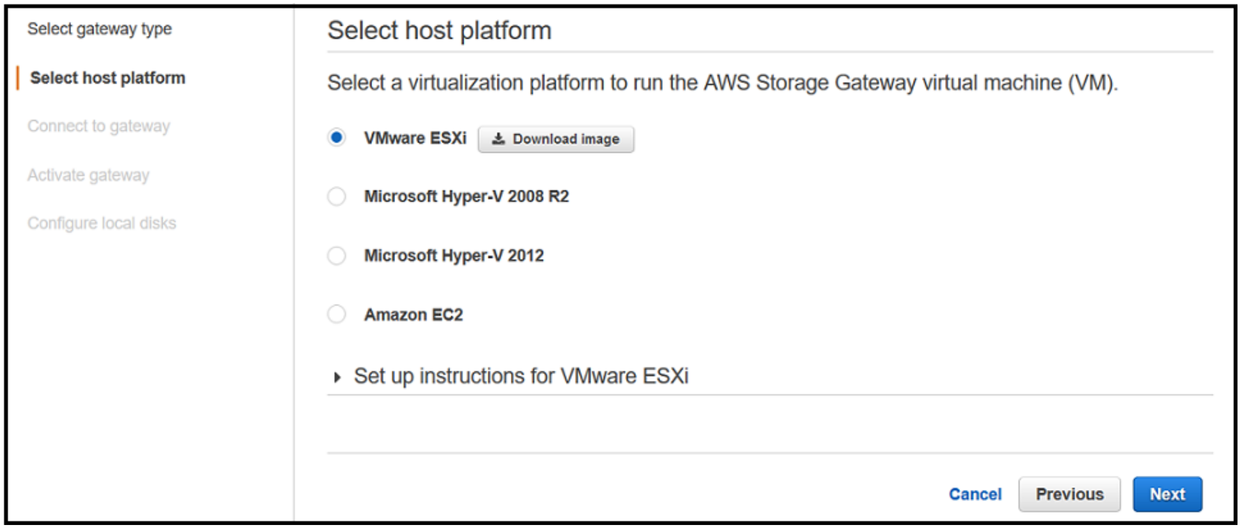

Select gateway type

Select host platform

Connect to gateway

Activate gateway

Configure local disks

Select host platform

Select a virtualization platform to run the AWS Storage Gateway virtual machine (VM).

● **VMware ESXi** Download image

○ **Microsoft Hyper-V 2008 R2**

○ **Microsoft Hyper-V 2012**

○ **Amazon EC2**

▸ Set up instructions for VMware ESXi

2. Choose the **Download image** button next to your virtualization platform to download a .zip file that contains the .ova file for your virtualization platform.

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

When your download is complete, you provision a host and deploy your gateway VM you downloaded. If you haven't provisioned a host yet, provision it now. You perform this task outside the AWS Storage Gateway console and return when you are done.

- Provisioning a VMware Host (p. 50)
- Provisioning a Hyper-V Host (p. 51)
- Provisioning an Amazon EC2 Host (p. 53)

Provisioning a VMware Host

Following, you can find instructions how to provision an on-premises VMware host and deploy your gateway. You perform this task in the VMware vSphere client and return to the AWS Storage Gateway console to continue your gateway setup.

To provision a VMware host

1. Review the minimum host requirements in Requirements (p. 11).
2. Provision a host in your data center with the VMware ESXi hypervisor. For a minimum set of instructions to install the hypervisor, see Configuring a VMware ESXi Host for AWS Storage Gateway (p. 208).

   Note
   If you plan to deploy AWS Storage Gateway using VMware High Availability (HA) for failover protection, see Using AWS Storage Gateway with VMware High Availability (p. 225).

To deploy your gateway on a VMware host

1. Connect to your gateway host's hypervisor by using your VMware vSphere client.
2. Deploy the OVF template package that you downloaded.
3. Choose **Next** through the following three screens. You might be prompted to select a data store on which to store the `.ova` package.
4. Allocate disks with **Thick provisioned format**.
5. Synchronize the time of your gateway VM to match your gateway host's time.

   You must also ensure that the clock on your gateway host is synchronized with a Network Time Protocol (NTP) server.

This procedure provides minimal instructions to help you deploy your gateway quickly. If you are not familiar with using the VMware vSphere client software, see Deploying the AWS Storage Gateway VM to Your VMWare Host (p. 212) for detailed instructions.

Next, you allocate local disks for your deployed gateway VM. After you activate your gateway, you configure the disks as local storage for the gateway's use.

To allocate a local disk from direct-attached storage

> **Important**
> All gateways require a disk of a minimum size of 150 GiB for use as an upload buffer. Gateway-cached and tape gateway setups require additional disks for use as cache storage. The minimum size of a cache disk is 1.1 times your upload buffer size. For more information, see AWS Storage Gateway Limits (p. 296).

1. Decide the number and size of disks to allocate for your gateway.
2. Start your VMware vSphere client, and then connect to your host.
3. In the client, follow the instructions provided by your host's client and add two virtual disks.
4. Configure the disks according to the sizes you decided for your gateway.

   For your disk configuration, we recommend thick provisioning. We also strongly recommend that you choose the **Specify a datastore or datastore cluster** option and select a data store for these disks different from the one that is used for the gateway VM. In addition, we strongly recommend that you dedicate one data store for the disks that you will use for the upload buffer and a different data store for other disks.
5. Configure your VM to use paravirtualized controllers.

This procedure provides minimal instructions to help you to quickly allocate disks for your gateway. If you are not familiar with using the VMware vSphere client software, see Provisioning Local Disk Storage for the Gateway VM (VMWare) (p. 218) for detailed instructions.

Provisioning a Hyper-V Host

Following, you can find instructions how to provision an on-premises Microsoft Hyper-V host and deploy your gateway. You perform this task in the Hyper-V Manager and return to the AWS Storage Gateway console to continue your gateway setup.

To provision a Hyper-V host

1. Review the minimum host requirements in Requirements (p. 11).
2. Set up a host in your data center with the Microsoft Hyper-V host. For a minimum set of instructions to install the hypervisor, see Configuring a Hyper-V Host for AWS Storage Gateway (p. 227).

To deploy your gateway on a Hyper-V host

1. Connect to the Microsoft Hyper-V Manager on your Microsoft Windows client.
2. Create locations on the hypervisor host for the gateway virtual hard disks and VM.

   a. Navigate to the hypervisor drive.

   b. Create a folder with two subfolders, `unzippedSourceVM` and `gateway`.
3. Configure the Hyper-V Manager to point to the `gateway` folder you created. The running VM stores its configuration in this folder.
4. Copy the unzipped source VM files to the folder you created on the host computer. Import the AWS Storage Gateway VM to the host. You must have 75 GiB of disk space for installation of the VM image and system data.
5. Rename the VM to avoid confusion with other VMs that you might import to the host.
6. Confirm that **Time synchronization** is selected for the VM.

   You must also ensure that the clock on your gateway host is synchronized with a Network Time Protocol (NTP) server.

This procedure provides minimal instructions to help you deploy your gateway quickly. If you are not familiar with using the Microsoft Hyper-V manager software, see Deploying a AWS Storage Gateway VM on a Microsoft Hyper-V Host (p. 236) for detailed instructions.

Next, you allocate local disks for your deployed gateway VM. After you activate your gateway, you configure the disk as local storage for the gateway's use.

To allocate a local disk from direct-attached storage

> **Important**
> All gateways require a disk of a minimum size of 150 GiB for use as an upload buffer. Volume gateways and tape gateway setups require additional disks for use as cache storage. The minimum size of a cache disk is 1.1 times your upload buffer size. For more information, see AWS Storage Gateway Limits (p. 296)

1. Decide the number and size of disks to allocate for your gateway.
2. Start the Microsoft Hyper-V Manager, and then connect to the hypervisor.
3. In the Hyper-V Manager, follow the instructions provided by your host's client and add two virtual disks.
4. Configure the disks according to the sizes that you decided for your gateway.

   For your disk configuration, we recommend thick provisioning. We also strongly recommend that you choose the **Specify a datastore or datastore cluster** option and select a data store for these disks different from the one that is used for the gateway VM. In addition, we strongly recommend that you dedicate one data store for the disks that you will use for the upload buffer and a different data store for other disks.
5. > **Note**
   > AWS Storage Gateway supports the .vhdx file type. This file type enables you to create larger virtual disks than other file types. If you create a .vhdx type virtual disk, make sure that the size of the virtual disks you create does not exceed the recommended disk size for your gateway.

   If you are using the Microsoft Hyper-V 2012 Hypervisor, you will be prompted to choose a disk format (**VHD** or **VHDX**).
6. Configure the disks as **Fixed size** for **Disk Type**.

When you use fixed-size provisioning, the disk storage is allocated immediately, resulting in better performance. If you don't use fixed-size provisioning, the storage is allocated on demand. On-demand allocation can have a negative impact on the functioning of AWS Storage Gateway. For AWS Storage Gateway to function properly, the VM disks must be stored in fixed-size provisioned format.

This procedure provides minimal instructions to help you quickly allocate disks. If you are not familiar with using the Microsoft Hyper-V manager software, see Provision Local Storage for the AWS Storage Gateway VM (Hyper-V) (p. 245) for detailed instructions.

Provisioning an Amazon EC2 Host

You can deploy and activate a gateway on an Amazon EC2 instance. Gateways hosted on Amazon EC2 instances can be used for disaster recovery and data mirroring. An Amazon Machine Image (AMI) is available in the AWS Marketplace for AWS Storage Gateway.

To deploy a gateway on an Amazon EC2 instance

1. On the **Step 2: Choose host platform** page, choose **Amazon EC2**.

   AWS Storage Gateway is available as an Amazon Machine Image (AMI) that contains the gateway VM image.
2. Choose **Setup a Storage Gateway on EC2**. You are redirected to Amazon Marketplace where you can launch the EC2 AMI.

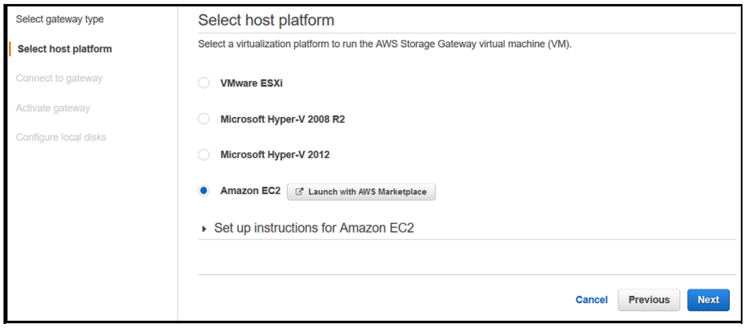

Select host platform

Select a virtualization platform to run the AWS Storage Gateway virtual machine (VM).

○ VMware ESXi

○ Microsoft Hyper-V 2008 R2

○ Microsoft Hyper-V 2012

● Amazon EC2

Provisioning an Amazon EC2 Host by Using an AMI

You can host and activate a gateway on an Amazon EC2 instance. Gateways hosted on Amazon EC2 instances can be used for disaster recovery and data mirroring. An Amazon Machine Image (AMI) is available in the AWS Marketplace for AWS Storage Gateway.

To deploy a gateway on an Amazon EC2 instance

1. On the **Select host platform** page, choose **Amazon EC2**.
2. Choose **Launch with AWS Marketplace**. You are redirected to AWS Marketplace where you launch the EC2 AMI.

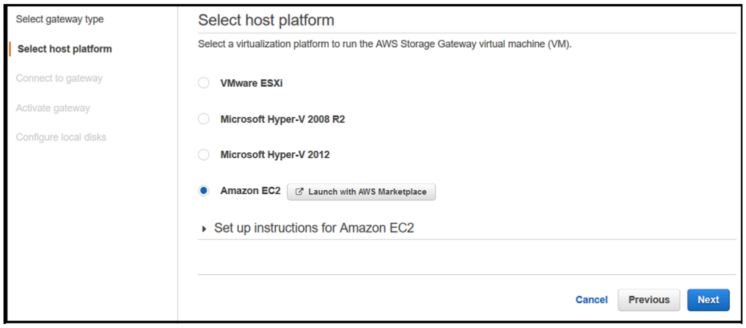

Select host platform

Select a virtualization platform to run the AWS Storage Gateway virtual machine (VM).

○ VMware ESXi

○ Microsoft Hyper-V 2008 R2

○ Microsoft Hyper-V 2012

● Amazon EC2

▸ Set up instructions for Amazon EC2

3. On AWS Marketplace, choose **Continue**.
4. Choose **1-Click Launch**. Doing this launches the AMI with default settings.
5. If this is your first time using an AWS Storage Gateway AMI, choose **Accept Terms** to accept the terms of service.
6. Review the default settings. You can accept and use these default settings or modify them to meet your needs.

   The 1-Click Launch feature comes with an autogenerated security group that is named AWS Storage Gateway-1-0-AutogenByAWSMP. For information about security group settings, see Configuring Security Groups for Your Amazon EC2 Gateway Instance (p. 15).
7. After reviewing all your settings, choose **Launch with 1-Click**.
8. Choose **Return to Product Page** and locate your instance on the Amazon EC2 console.

   > **Important**
   > EC2 instances launched with the 1-Click Launch functionality come with one root Amazon EBS volume. You need to add additional EBS volumes to your instance as a separate step after the instance is launched. For information about how to add EBS volumes, see Attaching an Amazon EBS Volume to an Instance.
9. In the Amazon EC2 console, choose your Amazon EC2 instance, choose the **Description** tab at the bottom, and then note the IP address. You will use this IP address to connect to the gateway.

Connecting to Your Gateway

To connect to your gateway, the first step is to get to the IP address of your gateway VM. You will use this IP address to activate your gateway. For gateways deployed and activated on an on-premises host, you can get the IP address from your gateway VM local console or your hypervisor client. For gateways deployed and activated on an Amazon EC2 instance, you can get the IP address from the Amazon EC2 console. The activation process associates your gateway with your AWS account. You gateway VM must be running for activation to succeed.

To get the IP address for you gateway VM from the local console

1. Log on to your gateway VM local console. For detailed instructions, see Configuring Your Gateway Network (p. 162).
2. Get the IP address from the top of the menu page.
3. Take note of the Gateway IP address.

To get the IP address from an EC2 instance

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.

3. Choose the **Description** tab at the bottom, and then note the IP address. You will use this IP address to activate the gateway.

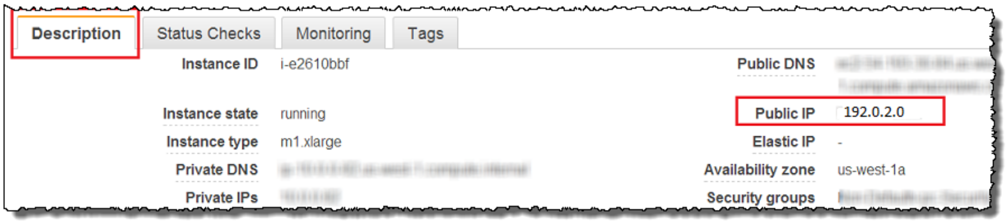

Instance ID	i-e2610bbf	Public DNS	[redacted]
Instance state	running	Public IP	192.0.2.0
Instance type	m1.xlarge	Elastic IP	-
Private DNS	[redacted]	Availability zone	us-west-1a
Private IPs	[redacted]	Security groups	[redacted]

For activation, you can use the public or private IP address assigned to the gateway. You must be able to reach the IP address that you use from the browser from which you perform the activation. In this walkthrough, we use the public IP address to activate the gateway.

To associate your gateway with your AWS account

1. If the **Connect to gateway** page isn't already open, open the console and navigate to the **Connect to gateway** page.
2. Type the IP address of your gateway for **IP address**, and then choose the **Connect gateway** button.

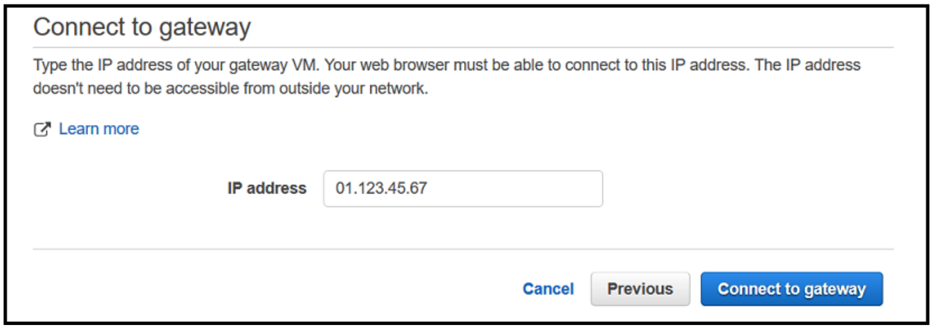

Connect to gateway

Type the IP address of your gateway VM. Your web browser must be able to connect to this IP address. The IP address doesn't need to be accessible from outside your network.

Learn more

IP address: 01.123.45.67

Activating Your Gateway

When your gateway VM is deployed and running, you configure your gateway settings and activate your gateway. If activation fails, check that the IP address you entered is correct. If the IP address is correct, confirm that your network is configured to let your browser access the gateway VM. For more information, see Troubleshooting On-Premises Gateway Issues (p. 191) or Troubleshooting Amazon EC2 Gateway Issues (p. 194) for troubleshooting guidelines.

To configure your gateway settings

1. To complete the activation process, type the information listed on the activation page to configure your gateway setting.

   The following screenshot shows the activation page for tape gateways.

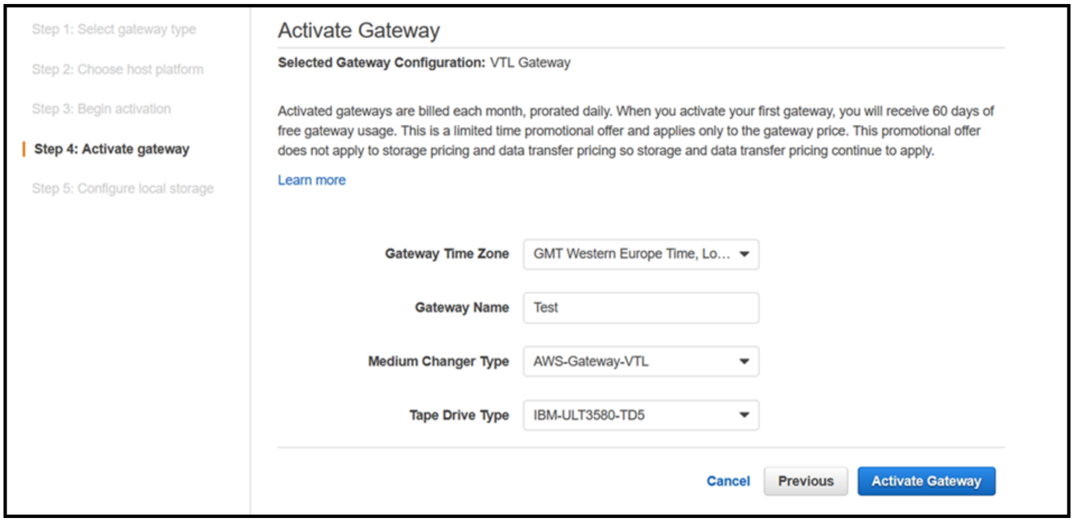

Step 1: Select gateway type

Step 2: Choose host platform

Step 3: Begin activation

Step 4: Activate gateway

Step 5: Configure local storage

Activate Gateway

Selected Gateway Configuration: VTL Gateway

Activated gateways are billed each month, prorated daily. When you activate your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies only to the gateway price. This promotional offer does not apply to storage pricing and data transfer pricing so storage and data transfer pricing continue to apply.

Learn more

Gateway Time Zone GMT Western Europe Time, Lo...

Gateway Name Test

Medium Changer Type AWS-Gateway-VTL

Tape Drive Type IBM-ULT3580-TD5

- **Gateway Time Zone** specifies the time zone to use for your gateway.
- **Gateway Name** identifies your gateway. You use this name to manage your gateway in the console; you can change it after the gateway is activated. This name must be unique to your account.
- **Medium Changer Type** specifies the type of medium changer to use for your backup software (available for tape gateways only).

Important

The type of medium changer you select depends on the backup software you plan to use. The following table shows which medium changer to select for your backup software. This list includes third-party backup software that has been tested and found to be compatible with tape gateway.

Backup Software	Medium Changer Type
Backup Exec 2012	`STK-L700`
Backup Exec 2014	`AWS-Gateway-VTL`
Backup Exec 15	`AWS-Gateway-VTL`
Backup Exec 16	`AWS-Gateway-VTL`
Dell NetVault Backup 10.0	`STK-L700`
EMC NetWorker 8.x	`STK-L700`
HPE Data Protector 9.x	`AWS-Gateway-VTL`
Microsoft System Center 2012 R2 Data Protection Manager	`STK-L700`
Symantec NetBackup Version 7.x	`AWS-Gateway-VTL`
Veeam Backup & Replication V7	`STK-L700`
Veeam Backup & Replication V8	`STK-L700`
Veeam Backup & Replication V9 Update 2 or later	`AWS-Gateway-VTL`

> **Note**
> You must select the medium changer that is recommended for your backup software. Other medium changers might not function properly. You can select a different medium changer after the gateway is activated. For more information, see Selecting a Medium Changer After Gateway Activation (p. 273).

- **Tape Drive Type** specifies the type of tape drive used by this gateway (available for tape gateways only).

2. Choose **Activate Gateway**.
3. When the gateway is successfully activated, the AWS Storage Gateway console displays the **Configure local storage** page.

   If activation is not successful, see Troubleshooting Your Gateway (p. 191) for possible solutions.

Configuring Local Disks

When you deployed the VM, you allocated local disks for your gateway. Now, you will configure your gateway to use these disks:

- For cached volumes (p. 4), you configure at least one disk for an upload buffer and the other for cache storage.
- For gateway-stored volumes (p. 5), you configure at least one disk for an upload buffer and allocate the rest of the storage for your application data.

To configure local disks

1. On the **Configure local disks** page, identify the disks you allocated and decide which ones you want to use for upload buffer and cached storage. For information about disk size limits, see Configuration and Performance Limits (p. 297).

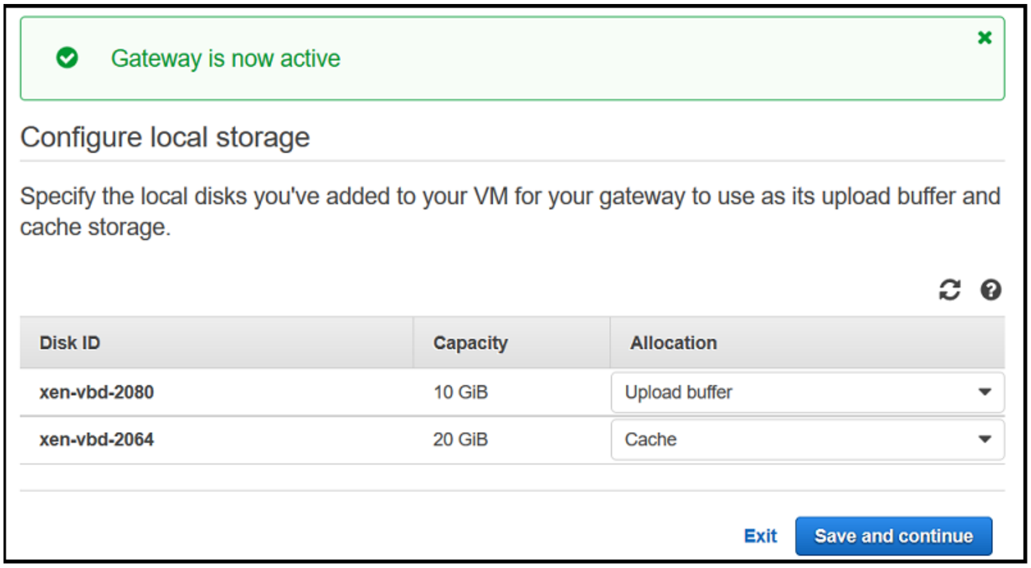

Gateway is now active

Configure local storage

Specify the local disks you've added to your VM for your gateway to use as its upload buffer and cache storage.

Disk ID	Capacity	Allocation
xen-vbd-2080	10 GiB	Upload buffer
xen-vbd-2064	20 GiB	Cache

2. From the list next to your upload buffer disk, choose **Upload Buffer**.
3. For cached volumes and tape, choose **Cache** for the disk you want to configure as a cache storage.

If you don't see your disks, choose **Refresh**.

4. Choose **Save and continue** to save your configuration settings.

For stored volumes, you configure one of the two disks for use by your application's data and the other disk as an upload buffer, and also optionally create alarms to monitor your disks.

Next Step

The remaining steps to create storage are specific to the type of gateway you created:

- For cached volumes, see Creating Volumes (p. 32).
- For stored volumes, see Creating Volumes (p. 32).

Creating Tapes

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

To create virtual tapes

1. In the navigation pane, choose the **Gateways** tab.
2. Choose **Create tapes** to open the **Create tape** dialog box.
3. For **Gateway**, choose a gateway. The tape is created for this gateway.
4. For **Number of tapes**, choose the number of tapes you want to create. For more information about tape limits, see AWS Storage Gateway Limits (p. 296).
5. For **Capacity**, type the size of the virtual tape you want to create. Tapes must be larger than 100 GiB. For information about capacity limits, see AWS Storage Gateway Limits (p. 296).
6. For **Barcode prefix**, type the prefix 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 for your virtual tapes. The prefix must be uppercase letters (A–Z) and must be one to four characters long.
7. Choose **Create tapes**.
8. In the navigation pane, choose the **Tapes** tab to see your tapes.

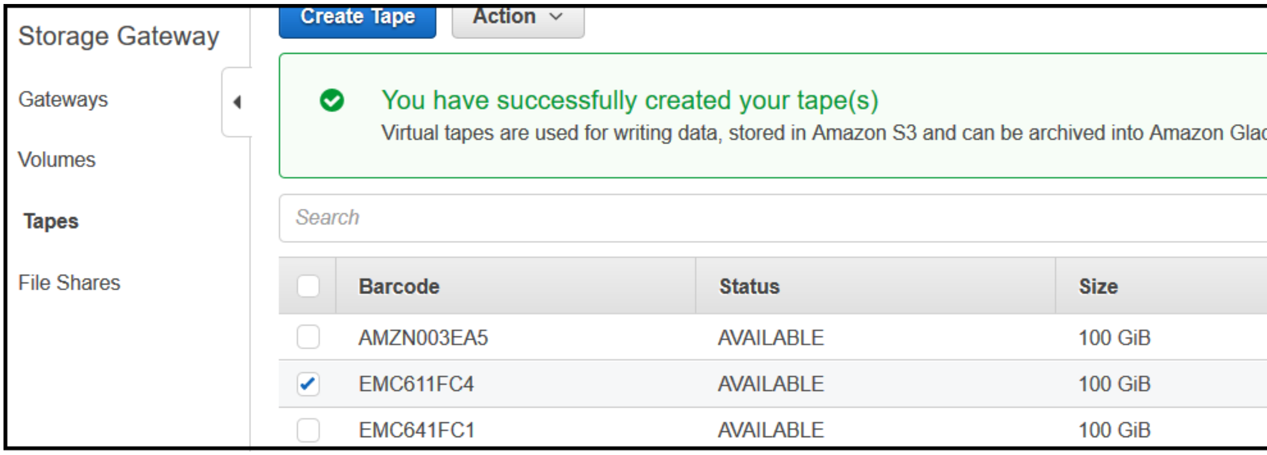

You have successfully created your tape(s)
Virtual tapes are used for writing data, stored in Amazon S3 and can be archived into Amazon Glac

□	Barcode	Status	Size
□	AMZN003EA5	AVAILABLE	100 GiB
☑	EMC611FC4	AVAILABLE	100 GiB
□	EMC641FC1	AVAILABLE	100 GiB

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 Working With Tapes (p. 275).

Using Your Tape Gateway

Following, you can find instructions about how to use your tape gateway.

Topics

Connect Your Tape Gateway Devices to Your Windows Client

Following, you can find instructions about how to connect you VTL devices to your Windows client. For instructions on accessing the iSCSI VTL devices from Linux, see Recommended Red Hat Linux iSCSI Settings (p. 44).

To connect your Windows client to the VTL devices

1. On the **Start** menu of your Windows client computer, type `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.

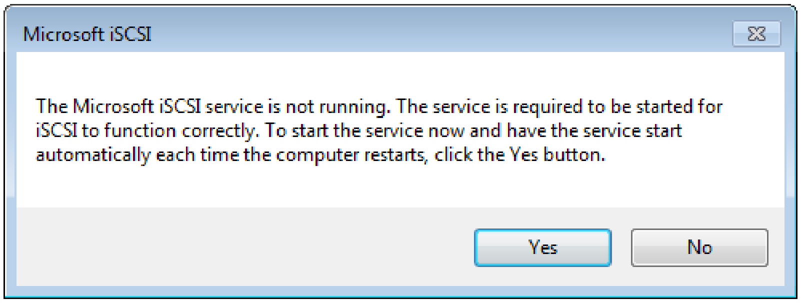

Microsoft iSCSI

The Microsoft iSCSI service is not running. The service is required to be started for iSCSI to function correctly. To start the service now and have the service start automatically each time the computer restarts, click the Yes button.

3. In the **iSCSI Initiator Properties** dialog box, choose the **Discovery** tab, and then choose the **Discover Portal** button.

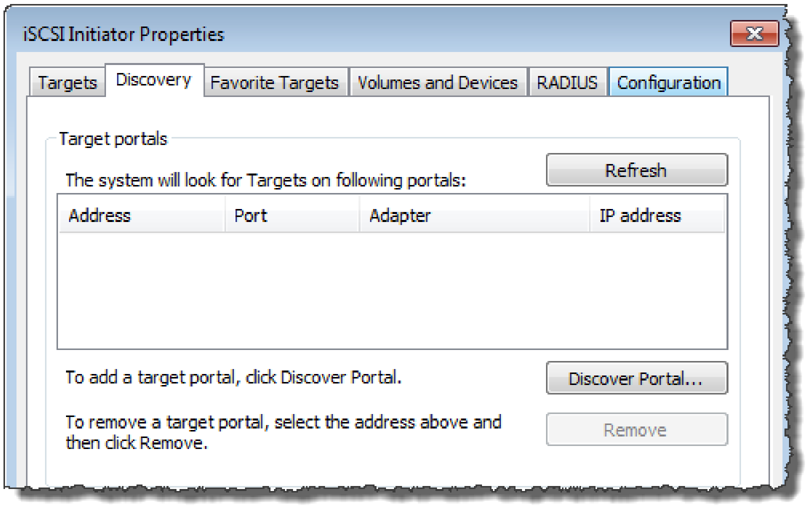

The system will look for Targets on following portals:

To add a target portal, click Discover Portal.

To remove a target portal, select the address above and then click Remove.

4. In the **Discover Target Portal** dialog box, type 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 AWS 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.

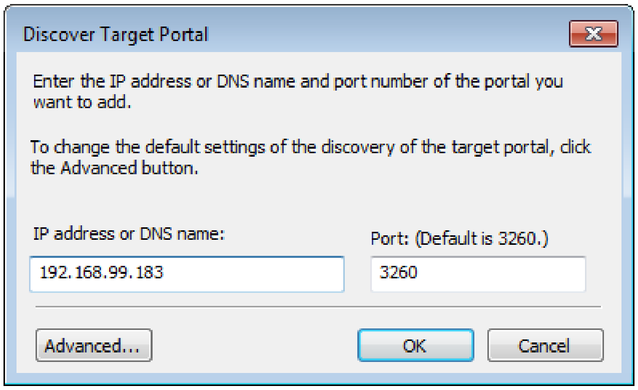

Enter the IP address or DNS name and port number of the portal you want to add.

To change the default settings of the discovery of the target portal, click the Advanced button.

IP address or DNS name: 192.168.99.183

Port: (Default is 3260.) 3260

5. Choose the **Targets** tab, and then choose **Refresh**. All ten tape drives and the medium changer appear in the **Discovered targets** box. The status for the targets is **Inactive**.

   The following screenshot shows the discovered targets.

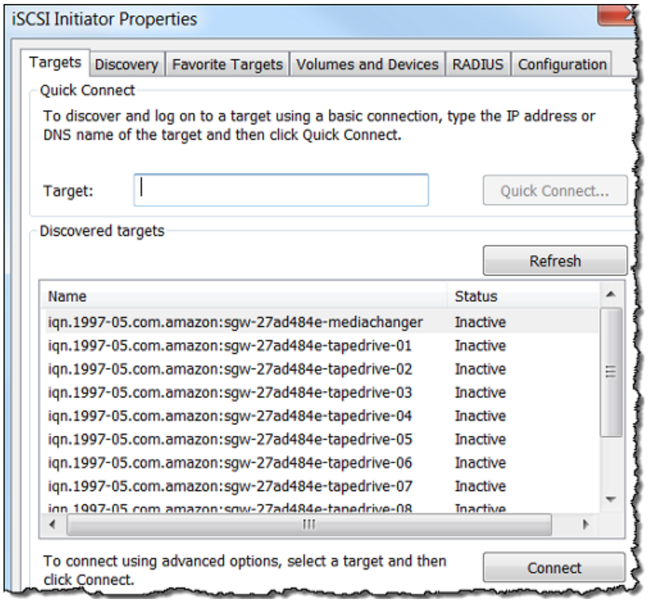

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Name	Status
iqn.1997-05.com.amazon:sgw-27ad484e-mediachanger	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-01	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-02	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-03	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-04	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-05	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-06	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-07	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-08	Inactive

To connect using advanced options, select a target and then click Connect.

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.
9. 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.

   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**.

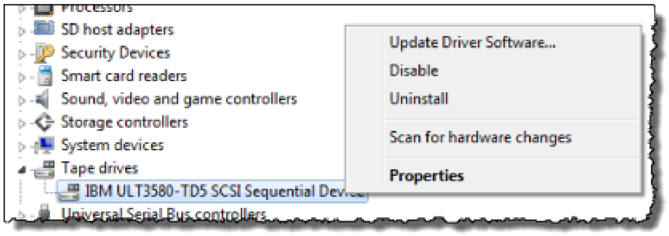
Processors
SD host adapters
Security Devices
Smart card readers
Sound, video and game controllers
Storage controllers
System devices
Tape drives
IBM ULT3580-TD5 SCSI Sequential Device
Universal Serial Bus controllers
Update Driver Software...
Disable
Uninstall
Scan for hardware changes
Properties

   3. In the **Driver** tab of the **Device Properties** dialog box, verify **Driver Provider** is Microsoft.

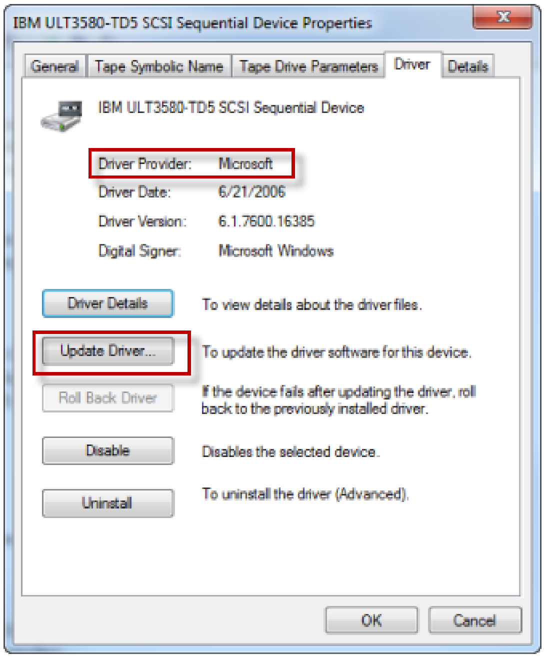

IBM ULT3580-TD5 SCSI Sequential Device Properties

IBM ULT3580-TD5 SCSI Sequential Device

Driver Provider: Microsoft
Driver Date: 6/21/2006
Driver Version: 6.1.7600.16385
Digital Signer: Microsoft Windows

To view details about the driver files.

To update the driver software for this device.

If the device fails after updating the driver, roll back to the previously installed driver.

Disables the selected device.

To uninstall the driver (Advanced).

4. If **Driver Provider** is not Microsoft, set the value as follows:

   1. Choose **Update Driver**.
   2. In the **Update Driver Software** dialog box, choose **Browse my computer for driver software**.

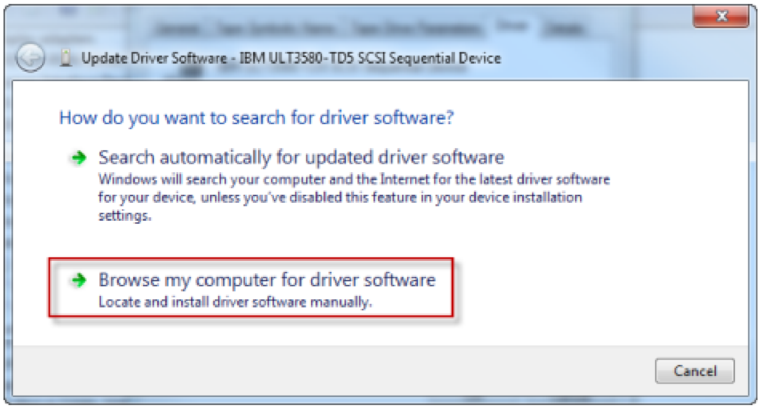

Update Driver Software - IBM ULT3580-TD5 SCSI Sequential Device

How do you want to search for driver software?

Search automatically for updated driver software
Windows will search your computer and the Internet for the latest driver software for your device, unless you've disabled this feature in your device installation settings.

Browse my computer for driver software
Locate and install driver software manually.

   3. In the **Update Driver Software** dialog box, choose **Let me pick from a list of device drivers on my computer**.

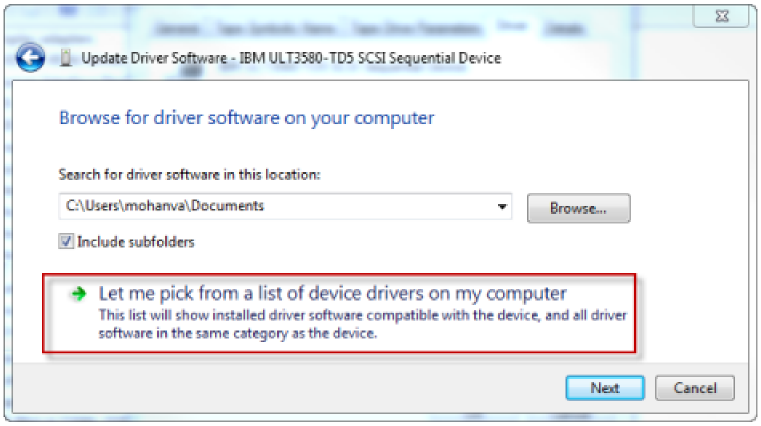
Update Driver Software - IBM ULT3580-TD5 SCSI Sequential Device

Browse for driver software on your computer

Search for driver software in this location:

C:\Users\mohanva\Documents

☑ Include subfolders

Let me pick from a list of device drivers on my computer

This list will show installed driver software compatible with the device, and all driver software in the same category as the device.

4. Select **LTO Tape drive** and choose **Next**.

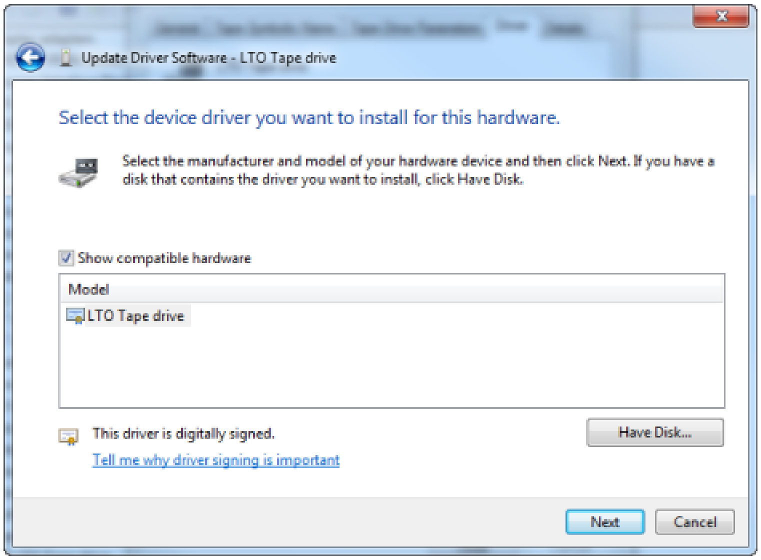
Update Driver Software - LTO Tape drive

Select the device driver you want to install for this hardware.

Select the manufacturer and model of your hardware device and then click Next. If you have a disk that contains the driver you want to install, click Have Disk.

☑ Show compatible hardware

Model
LTO Tape drive

This driver is digitally signed.

Tell me why driver signing is important

5. Choose **Close** to close the **Update Driver Software** window, and verify that the **Driver Provider** value is now set to Microsoft.
6. Repeat steps 9.2 through 9.5 to update all the tape drives.

Testing Your Gateway Setup

You test your tape gateway setup by performing the following tasks using your backup software:

1. Configure the backup software to detect your storage devices.
2. Back up data to a tape.
3. Archive the tape.
4. Retrieve the tape from the archive.
5. Restore data from the tape.

You can test your setup with one of the following types of compatible backup software:

- Testing Your Setup by Using Backup Exec (p. 64).
- Testing Your Setup by Using Dell NetVault Backup (p. 67).

- Testing Your Setup by Using EMC NetWorker (p. 70).
- Testing Your Setup by Using HPE Data Protector (p. 73).
- Testing Your Setup by Using Microsoft System Center 2012 R2 Data Protection Manager (p. 78).
- Testing Your Setup by Using Symantec NetBackup Version 7.x (p. 81).
- Testing Your Setup by Using Veeam Backup & Replication (p. 94).

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Testing Your Setup by Using Backup Exec

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

- Backup Exec 2014
- Backup Exec 15
- Backup Exec 16

The procedure for using these versions of Backup Exec with AWS Storage Gateway–VTL is the same. For detailed information about how to use Backup Exec, see the How to Create Secure Backups with Backup Exec video on the Backup Exec website. For Backup Exec support information on hardware compatibility, see the Software Compatibility Lists (SCL), Hardware Compatibility Lists (HCL), and Administrator Guides for Backup Exec (all versions) on the Backup Exec website. For information about best practices, see Best Practices for using Symantec Backup products (NetBackup, Backup Exec) with the Amazon Web Services (Tape Gateway) on the Symantec website.

Using Symantec Backup Exec, you can configure storage, write data to a tape, archive a tape, and restore the data.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

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 Connect Your Tape Gateway Devices to Your Windows Client (p. 59).

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.

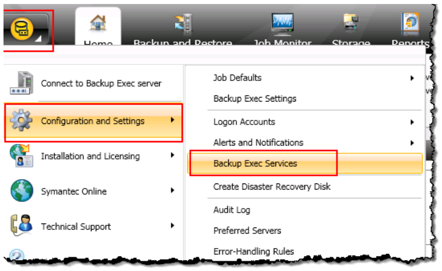

3. Choose **Restart All Services**. Backup Exec then recognizes the VTL devices (that is, the medium changer and tape drives). The restart process might take a few minutes.

   Note
   Tape Gateway provides 10 tape drives. However, your Backup Exec license agreement might require your backup software to work with fewer than 10 tape drives. In that case, you must disable tape drives in the Backup Exec robotic library to leave only the number of tape drives allowed by your license agreement enabled. For instructions, see Disabling a Tape Drive in Backup Exec (p. 67).

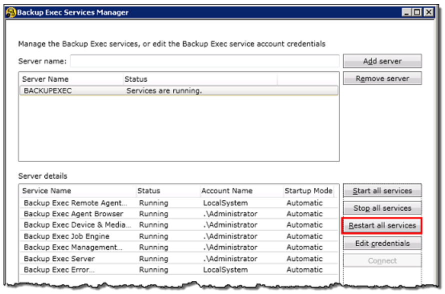

Manage the Backup Exec services, or edit the Backup Exec service account credentials

Server name:

Server Name	Status
BACKUPEXEC	Services are running.

Server details

Service Name	Status	Account Name	Startup Mode
Backup Exec Remote Agent...	Running	LocalSystem	Automatic
Backup Exec Agent Browser	Running	.\Administrator	Automatic
Backup Exec Device & Media...	Running	.\Administrator	Automatic
Backup Exec Job Engine	Running	.\Administrator	Automatic
Backup Exec Management...	Running	.\Administrator	Automatic
Backup Exec Server	Running	.\Administrator	Automatic
Backup Exec Error...	Running	LocalSystem	Automatic

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
   Symantec 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 software. For information about how to select a different medium changer type, see Selecting a Medium Changer After Gateway Activation (p. 273).

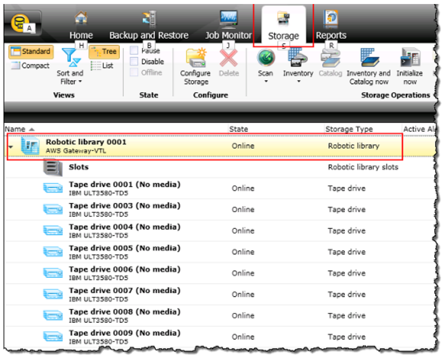

Name	State	Storage Type	Active Al
Robotic library 0001 AWS Gateway-VTL	Online	Robotic library	
Slots		Robotic library slots	
Tape drive 0001 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0003 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0004 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0005 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0006 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0007 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0008 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0009 (No media) IBM ULT3580-TD5	Online	Tape drive	

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 will be 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 (p. 116).

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

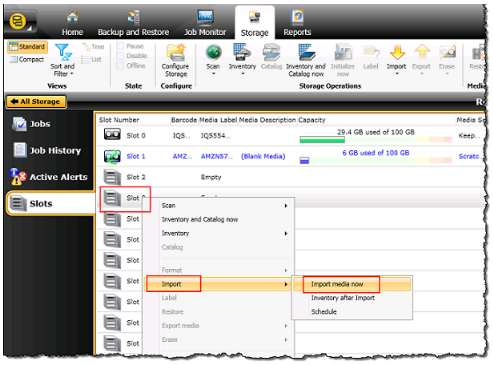

Slot Number	Barcode	Media Label	Media Description	Capacity	Media S
Slot 0	IQS...	IQS554...		29.4 GB used of 100 GB	Keep...
Slot 1	AMZ...	AMZN57...	(Blank Media)	6 GB used of 100 GB	Scratc...
Slot 2		Empty			

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

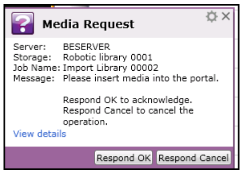

Media Request

Server: BESERVER
Storage: Robotic library 0001
Job Name: Import Library 00002
Message: Please insert media into the portal.

Respond OK to acknowledge.
Respond Cancel to cancel the operation.

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

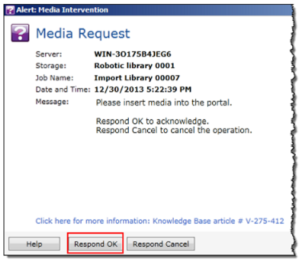

Alert: Media Intervention

Media Request

Server: WIN-3O175B4JEG6
Storage: Robotic library 0001
Job Name: Import Library 00007
Date and Time: 12/30/2013 5:22:39 PM
Message: Please insert media into the portal.

Respond OK to acknowledge.
Respond Cancel to cancel the operation.

Click here for more information: Knowledge Base article # V-275-412

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.

Archiving a Tape Using Backup Exec

When you archive a tape, AWS Storage 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.

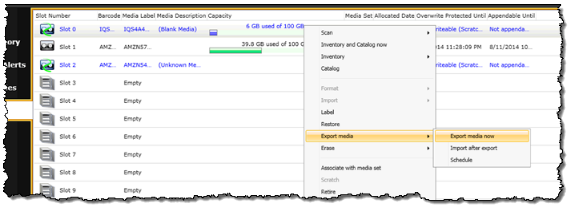

Slot Number	Barcode	Media Label	Media Description	Capacity	Media Set	Allocated Date	Overwrite Protected Until	Appendable Until
Slot 0	IQS...	IQS4A4...	(Blank Media)	6 GB used of 100 GB			riteable (Scratc...	Not appenda...
Slot 1	AMZ...	AMZN57...		39.8 GB used of 100 G			14 11:28:09 PM	8/11/2014 10...
Slot 2	AMZ...	AMZN54...	(Unknown Me...				riteable (Scratc...	Not appenda...
Slot 3		Empty						
Slot 4		Empty						
Slot 5		Empty						
Slot 6		Empty						
Slot 7		Empty						
Slot 8		Empty						
Slot 9		Empty						

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

   In the AWS 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 will be listed in the tape gateway's 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 will no longer be listed in the VTL.
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 AWS Storage Gateway console, choose **Tapes**. Verify that your tapes 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 the Archived Tape from Archive Back to Your Tape Gateway (p. 92).
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.

Disabling a Tape Drive in Backup Exec

Tape Gateway–provides 10 tape drives, but you might decide to use fewer tape drives. In that case, you disable the tape drives you will not 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 disable, and then choose **Disable**.

Testing Your Setup by Using Dell NetVault Backup

You can back up your data to virtual tapes, archive the tapes and manage your virtual tape library (VTL) devices by using Dell NetVault Backup version 10.0. In this topic, you can find basic documentation on how to configure the NetVault Backup software for a tape gateway and perform a backup. For detailed information about how to use the NetVault Backup software, see the NetVault Backup 10.0.1 – Administration Guide. In this topic, you can find out how to configure storage devices, write data to a tape, archive a tape and restore the data. For additional setup information, see Backing up to Amazon AWS with Dell NetVault Backup on the Dell website.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

- Configuring the NetVault Backup Software to Work with VTL Devices (p. 68)
- Backing Up Data to a Tape in the NetVault Backup Software (p. 69)
- Archiving a Tape by Using the NetVault Backup Software (p. 69)
- Restoring Data from a Tape Archived in the NetVault Backup Software (p. 70)

Configuring the NetVault Backup Software to Work with VTL Devices

After you have connected the virtual tape library (VTL) devices to the Windows client, you configure NetVault Backup to recognize your devices. For information about how to connect VTL devices to the Windows client, see Connect Your Tape Gateway Devices to Your Windows Client (p. 59).

By default, the NetVault Backup software does not automatically recognize tape gateway devices. You must manually add the devices to expose them to the NetVault Backup software and then discover the VTL devices.

Adding VTL Devices

To add the VTL devices

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

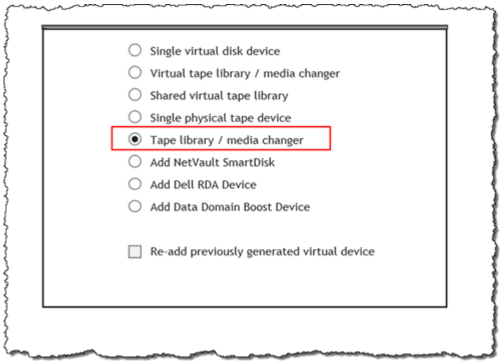

○ Single virtual disk device
○ Virtual tape library / media changer
○ Shared virtual tape library
○ Single physical tape device
● Tape library / media changer
○ Add NetVault SmartDisk
○ Add Dell RDA Device
○ Add Data Domain Boost Device

□ Re-add previously generated virtual device

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 will be displayed. In this case, your medium changer is displayed in the device box.
6. Select your medium changer and choose **Next**. Detailed information about the device is displayed in the wizard.
7. On the Add Tapes to Bays page, select **Scan For Devices**, choose your client machine, and then choose **Next**.

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

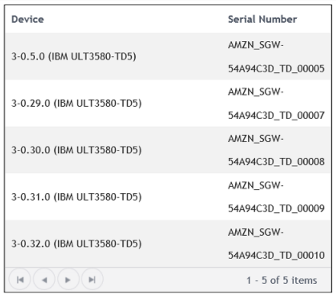

Device	Serial Number
3-0.5.0 (IBM ULT3580-TD5)	AMZN_SGW-54A94C3D_TD_00005
3-0.29.0 (IBM ULT3580-TD5)	AMZN_SGW-54A94C3D_TD_00007
3-0.30.0 (IBM ULT3580-TD5)	AMZN_SGW-54A94C3D_TD_00008
3-0.31.0 (IBM ULT3580-TD5)	AMZN_SGW-54A94C3D_TD_00009
3-0.32.0 (IBM ULT3580-TD5)	AMZN_SGW-54A94C3D_TD_00010

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 you added.

Backing Up Data to a Tape in the NetVault Backup Software

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 Dell NetVault Backup documentation.

Archiving a Tape by Using the NetVault Backup Software

When you archive a tape, AWS Storage Gateway moves the tape from the NetVault Backup 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 software—that is, the NetVault Backup software.

To archive a tape in NetVault Backup

1. In the 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.

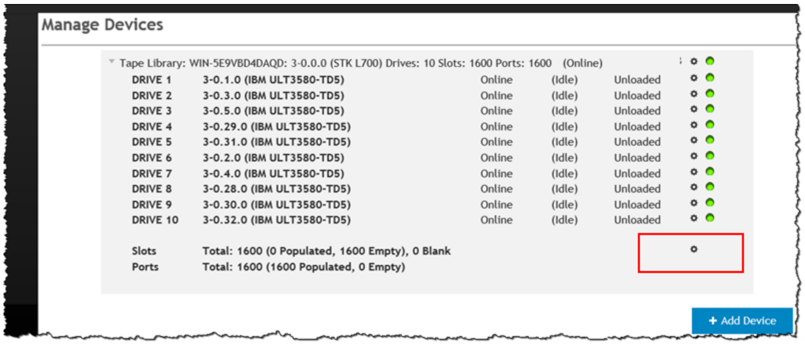

Manage Devices

Tape Library: WIN-5E9VBD4DAQD: 3-0.0.0 (STK L700) Drives: 10 Slots: 1600 Ports: 1600 (Online)

DRIVE 1	3-0.1.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 2	3-0.3.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 3	3-0.5.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 4	3-0.29.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 5	3-0.31.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 6	3-0.2.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 7	3-0.4.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 8	3-0.28.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 9	3-0.30.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 10	3-0.32.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
Slots	Total: 1600 (0 Populated, 1600 Empty), 0 Blank			
Ports	Total: 1600 (1600 Populated, 0 Empty)			

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

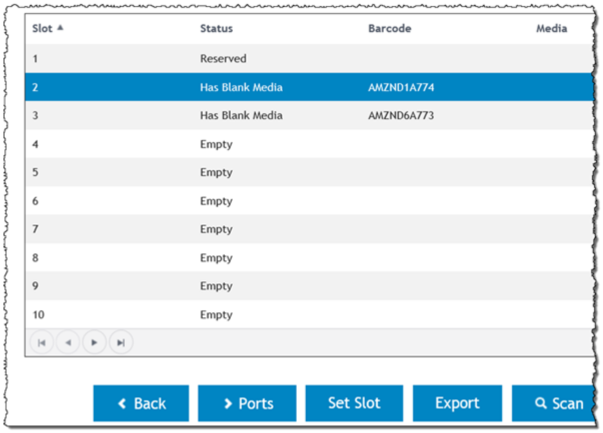

Slot ▲	Status	Barcode	Media
1	Reserved		
2	Has Blank Media	AMZND1A774	
3	Has Blank Media	AMZND6A773	
4	Empty		
5	Empty		
6	Empty		
7	Empty		
8	Empty		
9	Empty		
10	Empty		

Restoring Data from a Tape Archived in the NetVault Backup Software

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 the Archived Tape from Archive Back to Your Tape Gateway (p. 92).
2. Use the NetVault Backup 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 NetVault Backup 10.0.1 – Administration Guide (Creating a restore job) in the NetVault Backup documentation.

Testing Your Setup by Using EMC NetWorker

You can back up your data to virtual tapes, archive the tapes and manage your virtual tape library (VTL) devices by using EMC NetWorker version 8.1 or 8.2. In this topic, you can find basic documentation on how to configure the 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 EMC NetWorker software, see the *EMC NetWorker Administration Guide*.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

Configuring the EMC NetWorker Software to Work with VTL Devices

After you have connected your virtual tape library (VTL) devices to your Microsoft Windows client, you configure EMC NetWorker to recognize your devices. For information about how to connect VTL devices to the Windows client, see Connect Your Tape Gateway Devices to Your Windows Client (p. 59).

EMC NetWorker 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 EMC NetWorker software and that you are familiar with the EMC NetWorker Management Console. For more information about the EMC NetWorker Management Console, see the NetWorker Management Console interface section of the *EMC NetWorker Administration Guide*.

The following screenshot shows the EMC NetWorker Management Console.

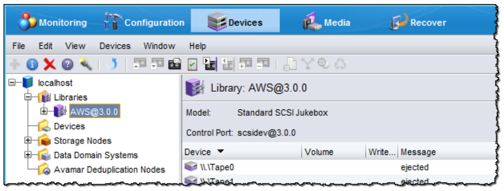

Library: AWS@3.0.0

Model: Standard SCSI Jukebox

Control Port: scsidev@3.0.0

Device	Volume	Write...	Message
\\.\Tape0			ejected
\\.\Tape1			ejected

To configure the EMC NetWorker software for VTL devices

1. Start the 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. This process might take a few seconds to load the devices into the library. In the example shown preceding, we have one library (`AWS@.3.0.0`).
6. Open the context (right-click) menu for your library, and then choose **Configure All Libraries**.
7. In the **Provide General Configuration Information** box, choose the configuration settings you want, and then choose **Next**.
8. In the **Select Target Storage Nodes** box, verify that a storage node is selected, and then choose **Start Configuration**.
9. In the Start Configuration wizard, choose **Finish**.
10. Choose your library to see your tapes in the left pane and the corresponding empty volume slots list in the right pane.

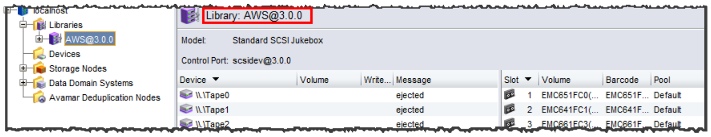

Library: AWS@3.0.0

Model: Standard SCSI Jukebox

Control Port: scsidev@3.0.0

Device	Volume	Write...	Message
\\.\Tape0			ejected
\\.\Tape1			ejected
\\.\Tape2			ejected

Slot	Volume	Barcode	Pool
1	EMC651FC0(...	EMC651F...	Default
2	EMC641FC1(...	EMC641F...	Default
3	EMC661FC3(...	EMC661F...	Default

11. In the volume list, select the volumes you want to enable (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.
12. In the dialog box that appears, choose **Yes**, and then in the **Load the Cartridges into** dialog box, choose **Yes**.
13. If you don't have any more tapes to deposit, choose **No** or **Ignore**. Otherwise, choose **Yes** to deposit additional tapes.

Enabling Import of WORM Tapes into EMC NetWorker

You are now ready to import tapes from your tape gateway into the EMC NetWorker library.

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

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

> **Note**
> You use the 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.

Archiving a Tape in EMC NetWorker

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

To archive a tape by using 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 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 an Archived Tape in EMC NetWorker

Restoring your archived data is a two-step process:

1. Retrieve the archived tape a tape gateway. For instructions, see Retrieving the Archived Tape from Archive Back to Your Tape Gateway (p. 92).
2. Use the 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 *EMC NetWorker Administration Guide.*

Testing Your Setup by Using 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 HPE Data Protector v9.x. In this topic, you can find basic documentation on how to configure the HPE Data Protector software for a tape gateway and perform a backup. This documentation includes how to configure storage devices, write data to a tape, archive a tape and restore the data. For detailed information about how to use the HPE Data Protector software, see the Hewlett Packard documentation.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

Configuring the HPE Data Protector Software to Work with VTL Devices

After you have connected the virtual tape library (VTL) devices to the client, you configure HPE Data Protector to recognize your devices. For information about how to connect VTL devices to the client, see Connect Your Tape Gateway Devices to Your Windows Client (p. 59).

By default, the 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 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**.

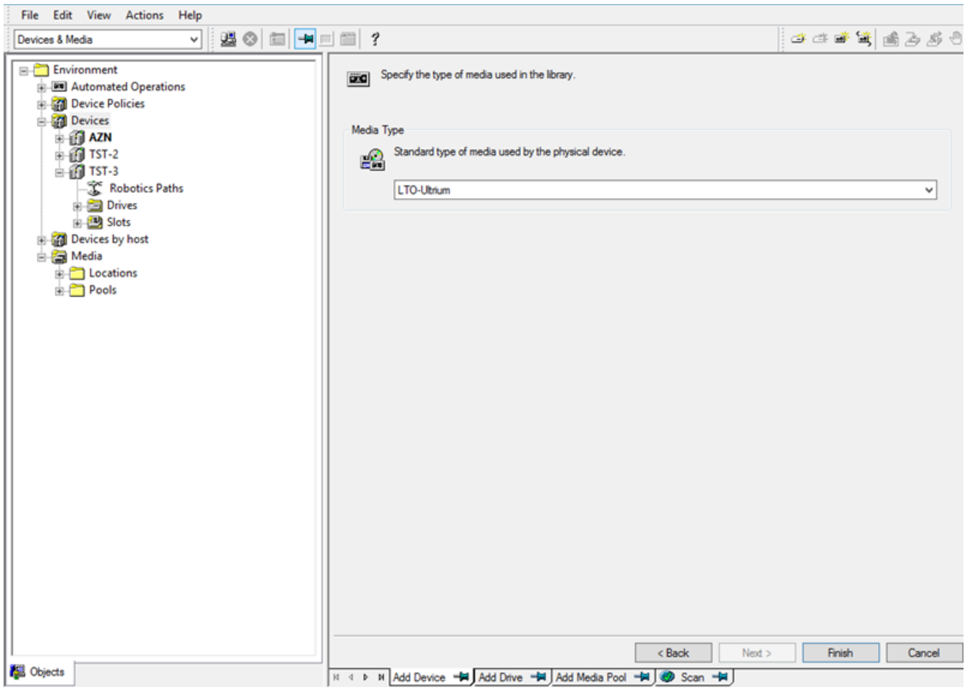

Specify the type of media used in the library.

Media Type

Standard type of media used by the physical device.

LTO-Ultrium

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 enable 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** popup 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 disabled (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 a tape is already loaded into the slot.
   - A grey slot means the slot is empty.
   - A cyan question mark means 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 organized 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 (grey) to a status of tape inserted (green). A series of messages appear to confirm 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 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.

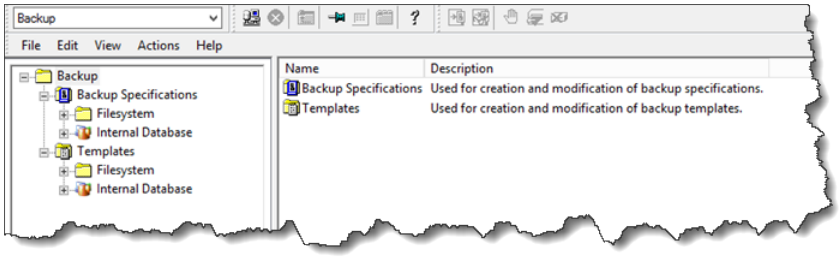

Name	Description
Backup Specifications	Used for creation and modification of backup specifications.
Templates	Used for creation and modification of backup templates.

2. Open the context (right-click) menu for **Filesystem**, and choose **Add Backup**.
3. On the **Create New Backup** screen, under **Filesystem**, choose **Blank File System Backup**, and then choose **OK**.

4. 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.
5. 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**.
6. Choose your media pool, choose **OK**, and then choose **Next**.
7. For the next three screens, accept the default settings and choose **Next**.
8. On the **Perform finishing steps in your backup/template design** screen, choose **Save as** to save this session. In the popup window, give the backup a name and assign it to the group where you want to save your new backup specification.
9. 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.

Archiving a Tape

When you archive a tape, AWS Storage 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 AWS 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 Amazon Glacier. 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.

Retrieving an Archived Tape

You can retrieve your archived tapes from Amazon Glacier with the following procedure.

To retrieve an archived tape

1. Open the AWS Storage Gateway console, and in the navigation pane choose **Tapes**.
2. Choose **Actions**, and then choose **Retrieve tape**.

   Note
   It takes about 24 hours to retrieve a tape from archive to a gateway.

When you have retrieved the tape, you can use an empty slot to reload the tape.

For more information, see Retrieving the Archived Tape from Archive Back to Your Tape Gateway (p. 92).

Restoring Data from a Tape

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.

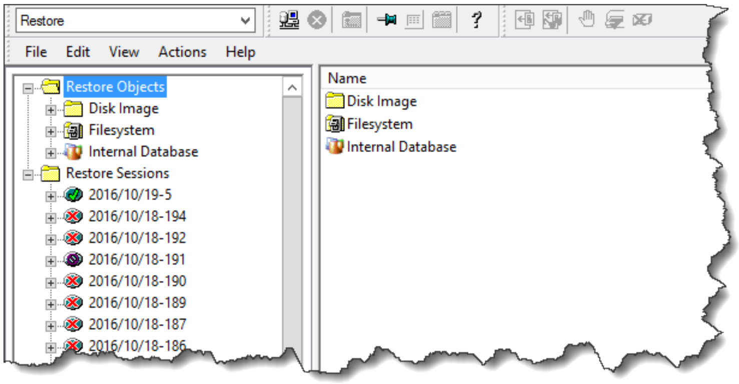

2. Choose the file system or database system you want to restore. For the backup that you want to restore, make sure the box is selected. Choose **Restore**.
3. In the **Start Restore Session** window, choose **Needed Media**. Choose **All media**, and you should see the tape originally used for the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

Testing Your Setup by Using Microsoft System Center 2012 R2 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 Data Protection Manager (DPM). In this topic, you can find basic documentation on how to configure the DPM backup software for a tape gateway and perform a backup. For detailed information about how to use DPM, see the DPM documentation on the Microsoft System Center website. In this topic, you can find out how to configure storage, write data to a tape, archive a tape, and restore the data.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

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 Connect Your Tape Gateway Devices to Your Windows Client (p. 59).

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 (p. 274).

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 (p. 58).
2. Import the tape into the DPM library. For information on how to do this, see Importing a Tape into DPM (p. 79).
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 2012 R2\DPM\DPM\Config\DPMLA.xml`.

   Note
   If this file exists, the DPMDriveMappingTool will overwrite 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 2012 R2\DPM\DPM\Bin`, and run the following command.

```
C:\Microsoft System Center 2012 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 software library.

To import tapes into the DPM backup software 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.

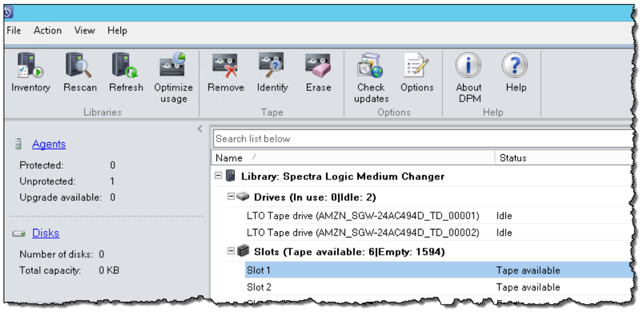

Agents

Protected: 0
Unprotected: 1
Upgrade available: 0

Disks

Number of disks: 0
Total capacity: 0 KB

Name	Status
Library: Spectra Logic Medium Changer	
Drives (In use: 0|Idle: 2)	
LTO Tape drive (AMZN_SGW-24AC494D_TD_00001)	Idle
LTO Tape drive (AMZN_SGW-24AC494D_TD_00002)	Idle
Slots (Tape available: 6|Empty: 1594)	
Slot 1	Tape available
Slot 2	Tape available

2. 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.

3. 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.

   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.

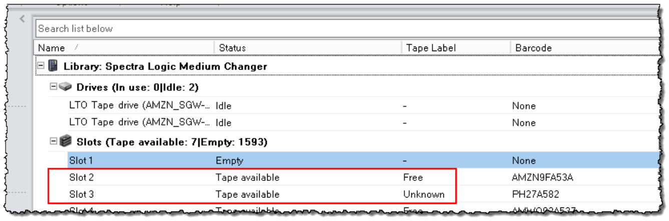

Name	Status	Tape Label	Barcode
Library: Spectra Logic Medium Changer			
Drives (In use: 0|Idle: 2)			
LTO Tape drive (AMZN_SGW-...	Idle	-	None
LTO Tape drive (AMZN_SGW-...	Idle	-	None
Slots (Tape available: 7|Empty: 1593)			
Slot 1	Empty	-	None
Slot 2	Tape available	Free	AMZN9FA53A
Slot 3	Tape available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

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.

Archiving a Tape by Using DPM

When you archive a tape, AWS Storage 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 software—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)**.

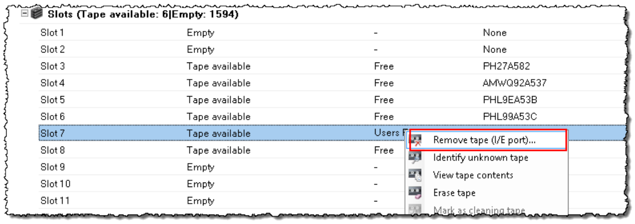

Slots (Tape available: 6|Empty: 1594)

Slot 1	Empty	-	None
Slot 2	Empty	-	None
Slot 3	Tape available	Free	PH27A582
Slot 4	Tape available	Free	AMWQ92A537
Slot 5	Tape available	Free	PHL9EA53B
Slot 6	Tape available	Free	PHL99A53C
Slot 7	Tape available	Users F	
Slot 8	Tape available	Free	
Slot 9	Empty	-	
Slot 10	Empty	-	
Slot 11	Empty	-	

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 AWS 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 the Archived Tape from Archive Back to Your Tape Gateway (p. 92).
2. Use the DPM backup software 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.

Testing Your Setup by Using Symantec NetBackup Version 7.x

You can back up your data to virtual tapes, archive the tapes and manage your virtual tape library (VTL) devices by using Symantec NetBackup version 7.x. In this topic, you can find basic documentation on how to configure the NetBackup software for a tape gateway and perform a backup. For detailed information about how to use NetBackup, see the NetBackup documentation on the Symantec website. For Symantec support information on hardware compatibility, see the Symantec NetBackup Enterprise Server and Server 7.x Hardware Compatibility List on the Symantec website.

In this topic, you can find out how to configure storage, write data to a tape, archive a tape and restore the data.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

Configuring NetBackup Storage Devices

After you have connected the virtual tape library (VTL) devices to the Windows client, you configure Symantec NetBackup version 7.x storage to recognize your devices. For information about how to connect VTL devices to the Windows client, see Connect Your Tape Gateway Devices to Your Windows Client (p. 59).

To configure NetBackup to use storage devices on your tape gateway

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

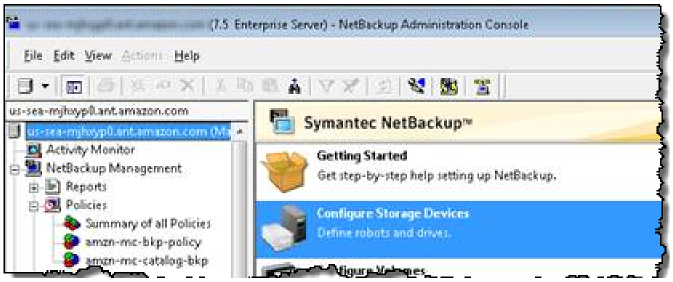
[redacted] (7.5 Enterprise Server) - NetBackup Administration Console
File Edit View Actions Help
us-sea-mjhxyp0.ant.amazon.com
us-sea-mjhxyp0.ant.amazon.com (Ma
Activity Monitor
NetBackup Management
Reports
Policies
Summary of all Policies
amzn-mc-bkp-policy
amzn-mc-catalog-bkp
Symantec NetBackup™
Getting Started
Get step-by-step help setting up NetBackup.
Configure Storage Devices
Define robots and drives.
Configure Volumes

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

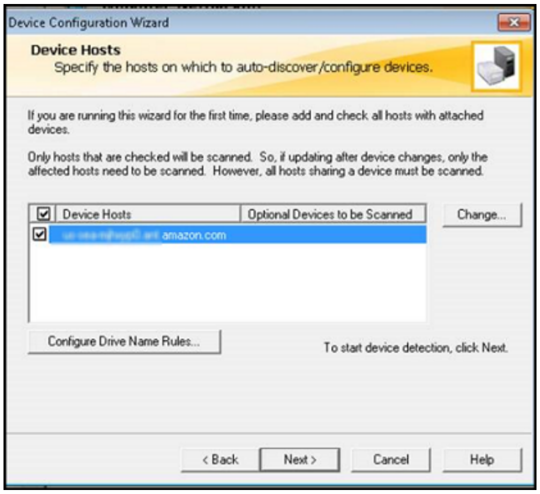

Specify the hosts on which to auto-discover/configure devices.

If you are running this wizard for the first time, please add and check all hosts with attached devices.

Only hosts that are checked will be scanned. So, if updating after device changes, only the affected hosts need to be scanned. However, all hosts sharing a device must be scanned.

☑ Device Hosts	Optional Devices to be Scanned
☑ [redacted] amazon.com	

To start device detection, click Next.

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

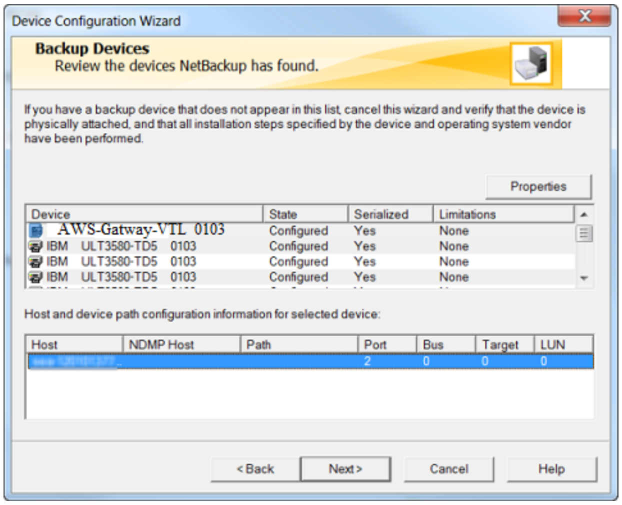

Backup Devices

Review the devices NetBackup has found.

If you have a backup device that does not appear in this list, cancel this wizard and verify that the device is physically attached, and that all installation steps specified by the device and operating system vendor have been performed.

Device	State	Serialized	Limitations
AWS-Gatway-VTL 0103	Configured	Yes	None
IBM ULT3580-TD5 0103	Configured	Yes	None
IBM ULT3580-TD5 0103	Configured	Yes	None
IBM ULT3580-TD5 0103	Configured	Yes	None

Host and device path configuration information for selected device:

Host	NDMP Host	Path	Port	Bus	Target	LUN
[redacted]			2	0	0	0

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

To verify your devices in the NetBackup software

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.

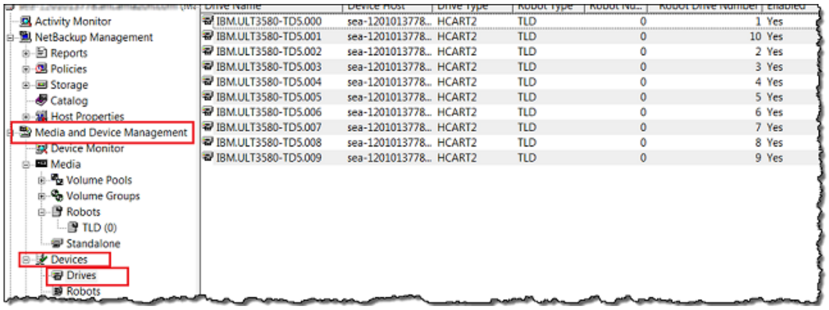

Drive Name	Device Host	Drive Type	Robot Type	Robot Nu...	Robot Drive Number	Enabled
IBM.ULT3580-TD5.000	sea-1201013778...	HCART2	TLD	0	1	Yes
IBM.ULT3580-TD5.001	sea-1201013778...	HCART2	TLD	0	10	Yes
IBM.ULT3580-TD5.002	sea-1201013778...	HCART2	TLD	0	2	Yes
IBM.ULT3580-TD5.003	sea-1201013778...	HCART2	TLD	0	3	Yes
IBM.ULT3580-TD5.004	sea-1201013778...	HCART2	TLD	0	4	Yes
IBM.ULT3580-TD5.005	sea-1201013778...	HCART2	TLD	0	5	Yes
IBM.ULT3580-TD5.006	sea-1201013778...	HCART2	TLD	0	6	Yes
IBM.ULT3580-TD5.007	sea-1201013778...	HCART2	TLD	0	7	Yes
IBM.ULT3580-TD5.008	sea-1201013778...	HCART2	TLD	0	8	Yes
IBM.ULT3580-TD5.009	sea-1201013778...	HCART2	TLD	0	9	Yes

2. In the **Devices** node, choose **Robots** to display all your medium changers. In the NetBackup software, 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**.

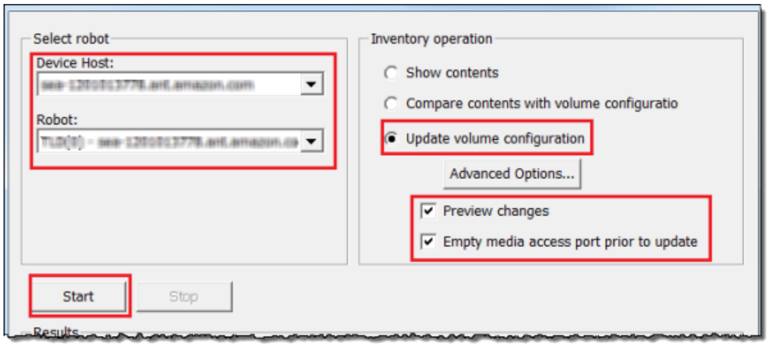

The process then inventories your medium changer and virtual tapes in the NetBackup Enterprise Media Management (EMM) database. NetBackup stores media information, device configuration, and tape status in the EMM.

7. 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.

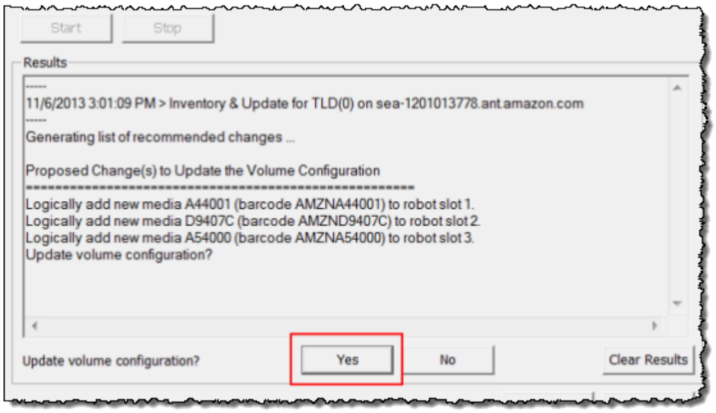

```
-----
11/6/2013 3:01:09 PM > Inventory & Update for TLD(0) on sea-1201013778.ant.amazon.com
-----
Generating list of recommended changes ...

Proposed Change(s) to Update the Volume Configuration
==========================================================
Logically add new media A44001 (barcode AMZNA44001) to robot slot 1.
Logically add new media D9407C (barcode AMZND9407C) to robot slot 2.
Logically add new media A54000 (barcode AMZNA54000) to robot slot 3.
Update volume configuration?
```

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

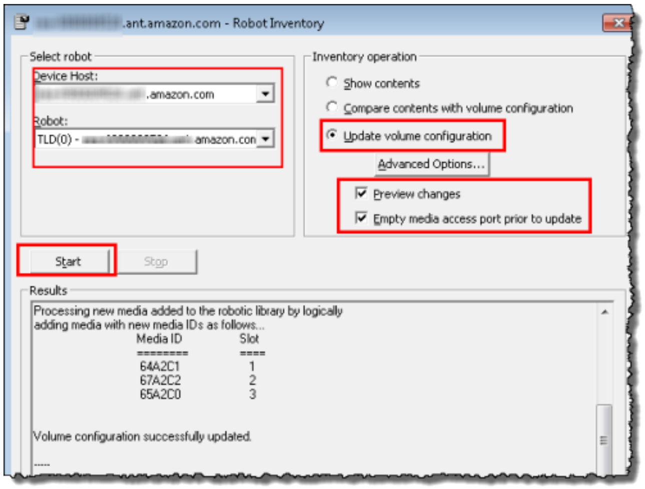

[redacted].ant.amazon.com - Robot Inventory
Select robot
Device Host:
[redacted].amazon.com
Robot:
TLD(0) - [redacted].amazon.con
Inventory operation
○ Show contents
○ Compare contents with volume configuration
● Update volume configuration
Advanced Options...
☑ Preview changes
☑ Empty media access port prior to update
Start
Stop
Results

```
Processing new media added to the robotic library by logically
adding media with new media IDs as follows...
                    Media ID          Slot
                    ========          ====
                    64A2C1             1
                    67A2C2             2
                    65A2C0             3

Volume configuration successfully updated.

-----
```

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 software, you might have multiple robots. Make sure you select the right robot.

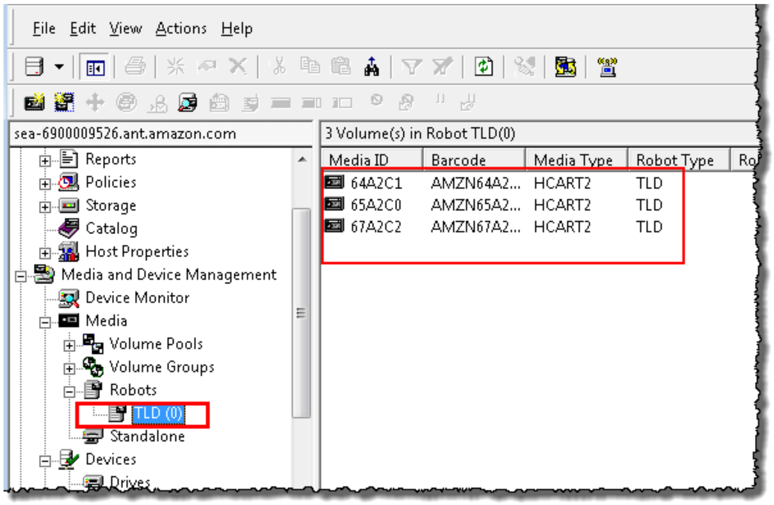

3 Volume(s) in Robot TLD(0)

Media ID	Barcode	Media Type	Robot Type	Ro
64A2C1	AMZN64A2...	HCART2	TLD	
65A2C0	AMZN65A2...	HCART2	TLD	
67A2C2	AMZN67A2...	HCART2	TLD	

Now that you have connected your devices and made them available to your backup software, 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 Sample Data to a Tape on Your Tape Gateway

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 AWS Storage Gateway detail page.

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.

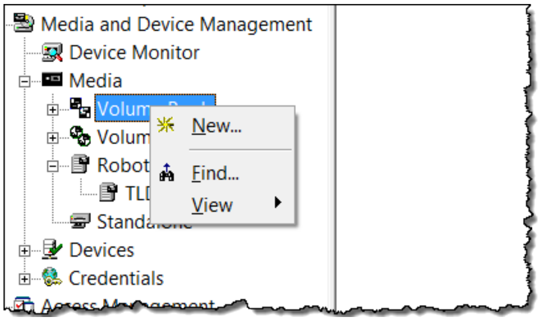

3. For **Name**, type a name for your volume pool.
4. For **Description**, type a description for the volume pool, and then choose **OK**. The volume pool you just created is added to the volume pool list.

   The following screenshot shows a list of volume pools.

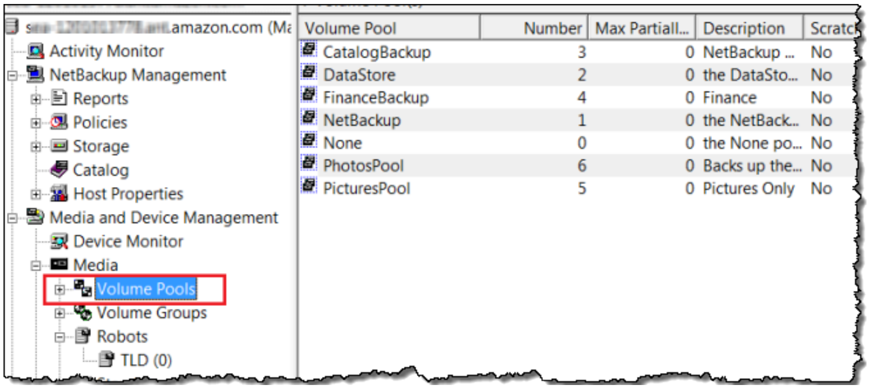

Volume Pool	Number	Max Partiall...	Description	Scratc
CatalogBackup	3	0	NetBackup ...	No
DataStore	2	0	the DataSto...	No
FinanceBackup	4	0	Finance	No
NetBackup	1	0	the NetBack...	No
None	0	0	the None po...	No
PhotosPool	6	0	Backs up the...	No
PicturesPool	5	0	Pictures Only	No

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.

   Note that 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.

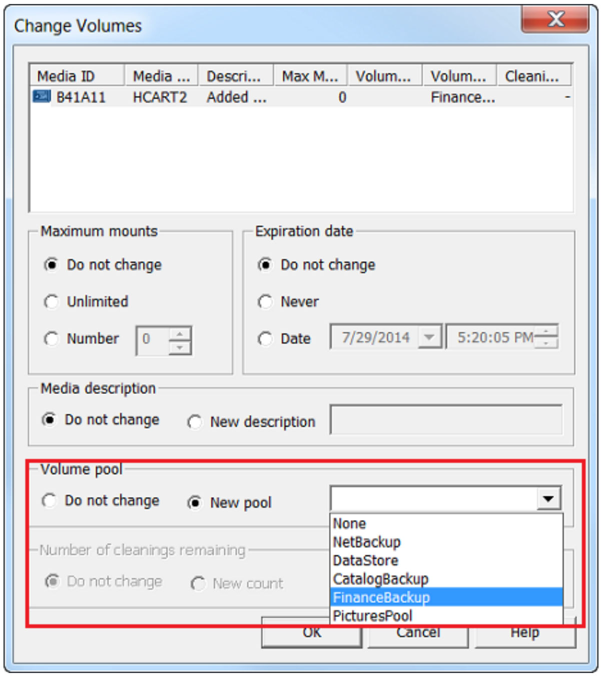

Change Volumes

Media ID	Media ...	Descri...	Max M...	Volum...	Volum...	Cleani...
B41A11	HCART2	Added ...	0		Finance...	-

Maximum mounts
● Do not change
○ Unlimited
○ Number 0

Expiration date
● Do not change
○ Never
○ Date 7/29/2014 5:20:05 PM

Media description
● Do not change ○ New description

Volume pool
○ Do not change ● New pool

Number of cleanings remaining
● Do not change ○ New count

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 Symantec NetBackup console.

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

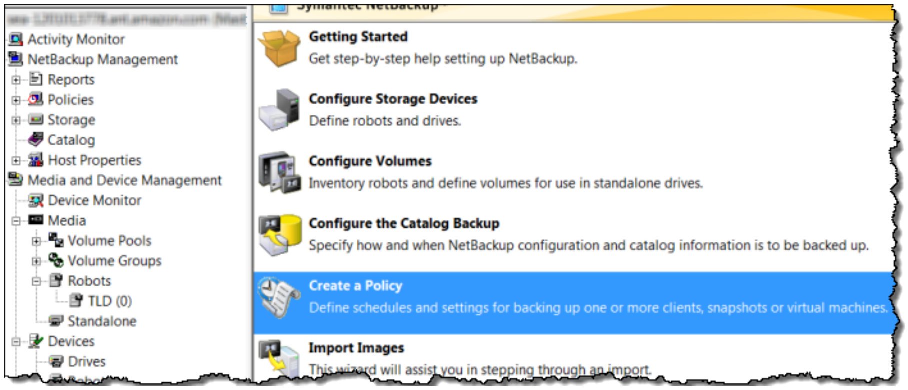

Getting Started
Get step-by-step help setting up NetBackup.

Configure Storage Devices
Define robots and drives.

Configure Volumes
Inventory robots and define volumes for use in standalone drives.

Configure the Catalog Backup
Specify how and when NetBackup configuration and catalog information is to be backed up.

Create a Policy
Define schedules and settings for backing up one or more clients, snapshots or virtual machines.

Import Images
This wizard will assist you in stepping through an import.

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).

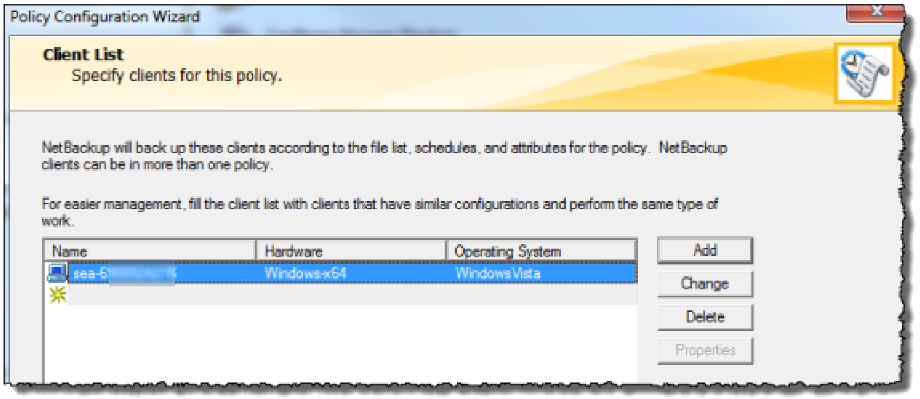

Client List
Specify clients for this policy.

NetBackup will back up these clients according to the file list, schedules, and attributes for the policy. NetBackup clients can be in more than one policy.

For easier management, fill the client list with clients that have similar configurations and perform the same type of work.

Name	Hardware	Operating System
sea-6[redacted]	Windows-x64	WindowsVista

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

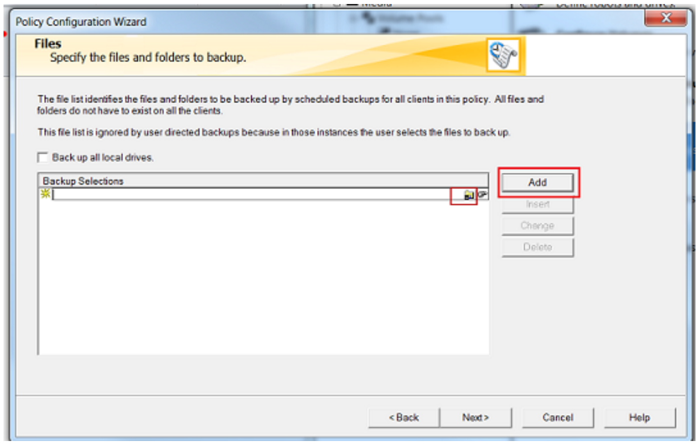

Policy Configuration Wizard

Files
Specify the files and folders to backup.

The file list identifies the files and folders to be backed up by scheduled backups for all clients in this policy. All files and folders do not have to exist on all the clients.

This file list is ignored by user directed backups because in those instances the user selects the files to back up.

□ Back up all local drives.

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**.

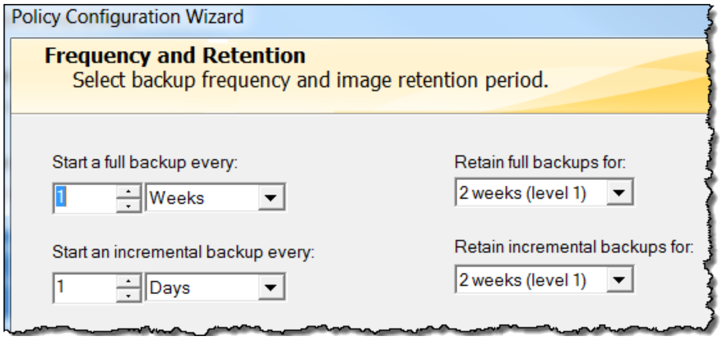

Policy Configuration Wizard

Frequency and Retention
Select backup frequency and image retention period.

Start a full backup every: 1 Weeks

Retain full backups for: 2 weeks (level 1)

Start an incremental backup every: 1 Days

Retain incremental backups for: 2 weeks (level 1)

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.

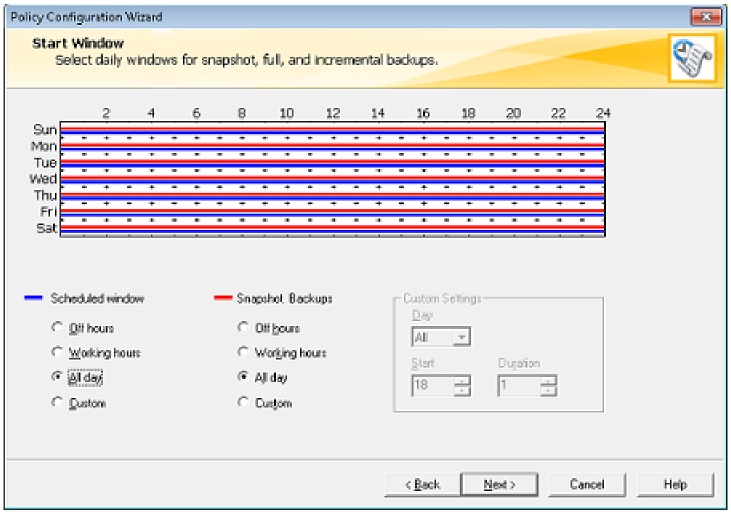

Policy Configuration Wizard

Start Window
Select daily windows for snapshot, full, and incremental backups.

Scheduled window: ○ Off hours ○ Working hours ● All day ○ Custom

Snapshot Backups: ○ Off hours ○ Working hours ● All day ○ Custom

Custom Settings: Day: All; Start: 18; Duration: 1

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 will 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**.

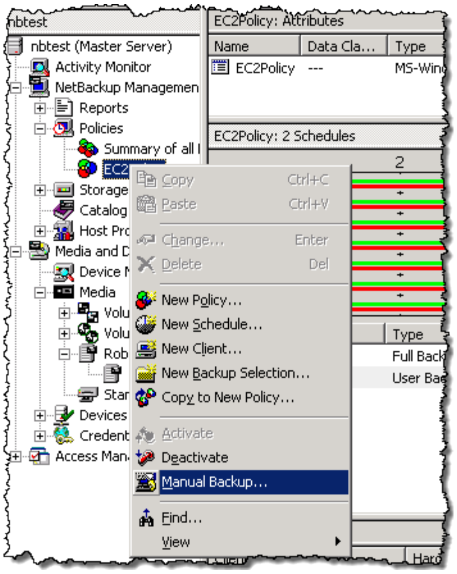

4. In the **Manual Backup** window, select a schedule, select a client, and then choose **OK**.

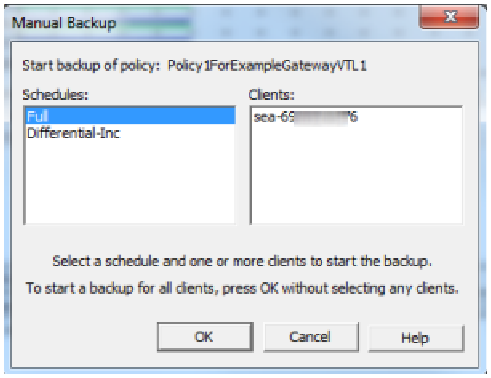

Manual Backup

Start backup of policy: Policy1ForExampleGatewayVTL1

Schedules:
- Full
- Differential-Inc

Clients:
- sea-69[redacted]6

Select a schedule and one or more clients to start the backup.

To start a backup for all clients, press OK without selecting any clients.

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.

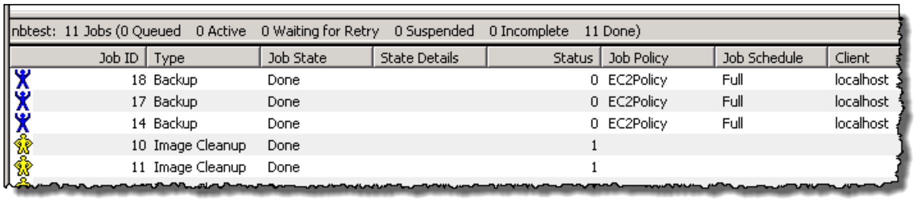

nbtest: 11 Jobs (0 Queued 0 Active 0 Waiting for Retry 0 Suspended 0 Incomplete 11 Done)

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 will 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.

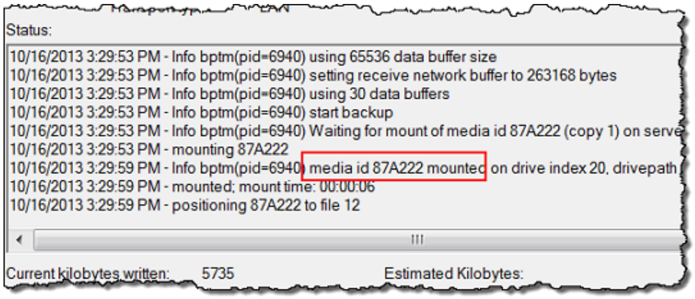

Status:

```
10/16/2013 3:29:53 PM - Info bptm(pid=6940) using 65536 data buffer size
10/16/2013 3:29:53 PM - Info bptm(pid=6940) setting receive network buffer to 263168 bytes
10/16/2013 3:29:53 PM - Info bptm(pid=6940) using 30 data buffers
10/16/2013 3:29:53 PM - Info bptm(pid=6940) start backup
10/16/2013 3:29:53 PM - Info bptm(pid=6940) Waiting for mount of media id 87A222 (copy 1) on serve
10/16/2013 3:29:53 PM - mounting 87A222
10/16/2013 3:29:59 PM - Info bptm(pid=6940) media id 87A222 mounted on drive index 20, drivepath
10/16/2013 3:29:59 PM - mounted; mount time: 00:00:06
10/16/2013 3:29:59 PM - positioning 87A222 to file 12
```

Current kilobytes written: 5735 Estimated Kilobytes:

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, AWS Storage 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**.

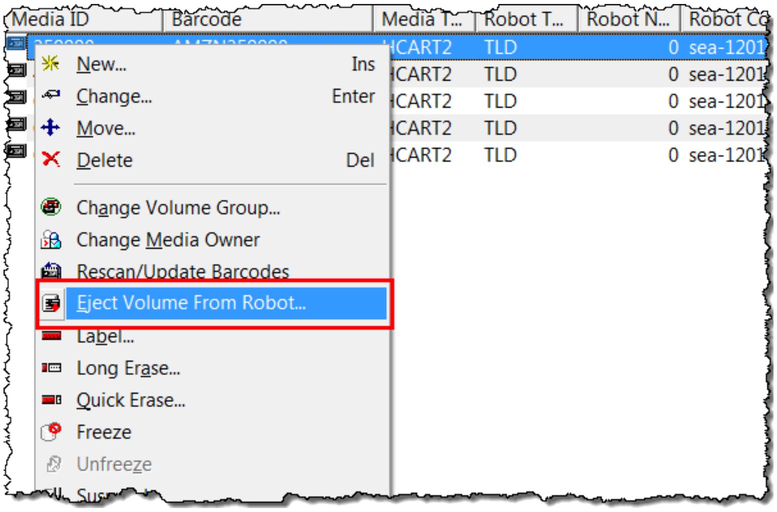

4. In the **Eject Volumes** window, make sure the **Media ID** matches the virtual tape you want to eject, and then choose **Eject**.

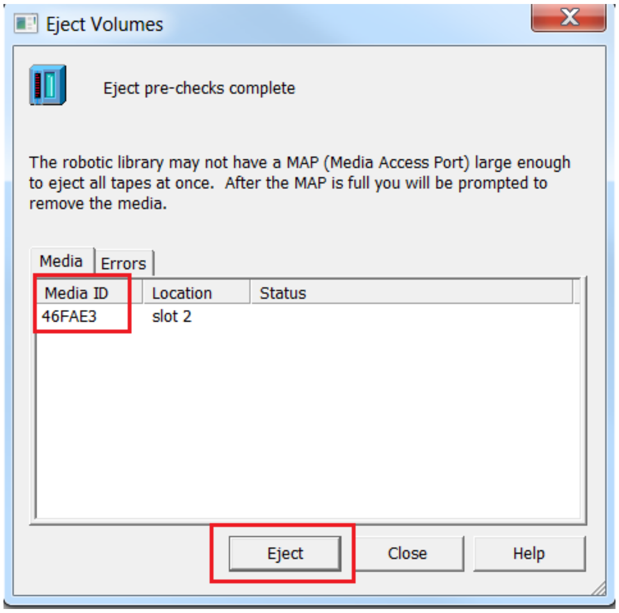

Eject pre-checks complete

The robotic library may not have a MAP (Media Access Port) large enough to eject all tapes at once. After the MAP is full you will be prompted to remove the media.

Media ID	Location	Status
46FAE3	slot 2	

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

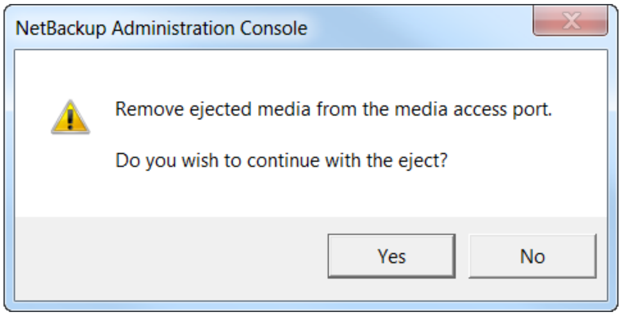

Remove ejected media from the media access port.

Do you wish to continue with the eject?

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

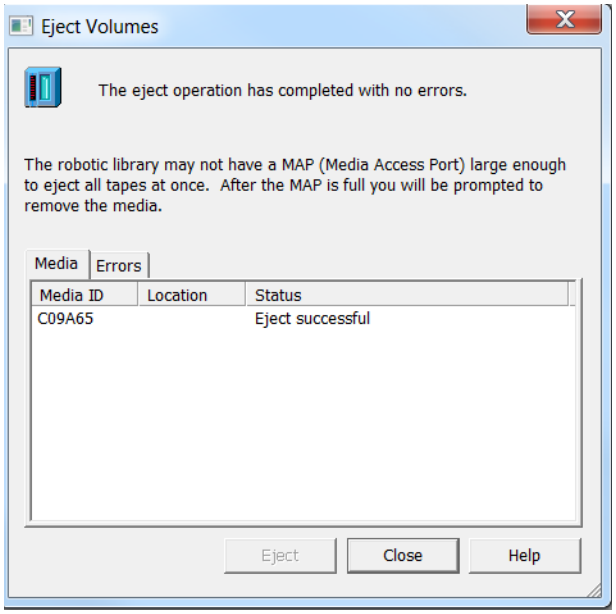

The eject operation has completed with no errors.

The robotic library may not have a MAP (Media Access Port) large enough to eject all tapes at once. After the MAP is full you will be prompted to remove the media.

Media ID	Location	Status
C09A65		Eject successful

6. Choose **Close** to close the **Eject Volumes** window.
7. In the AWS 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 will be listed in the gateway's VTL with the status IN TRANSIT TO VTS. When archiving starts, the status will be ARCHIVING. Once data upload has completed, the ejected tape will no longer be listed in the VTL.
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 AWS Storage Gateway console, choose **Tapes**. Verify that your archived tape's status is ARCHIVED.

Retrieving the Archived Tape from Archive Back to Your Tape Gateway

Archived tapes are stored in archive, which provides offline storage. If you want to access tape data, you must first retrieve the tape from the archive back to your gateway. In this step, you will retrieve the tape that you archived in the preceding step.

> **Note**
> It takes about 24 hours to retrieve a tape from the archive to a gateway.

To retrieve an archived tape

1. In the navigation pane of the AWS Storage Gateway console, choose **Tapes**. The console displays all virtual tapes and their status.
2. Select the virtual tape you want to retrieve, and choose **Retrieve Tape** The status of the tape must be ARCHIVED.
3. In the **Actions** drop-down list, choose **Retrieve tape** to open the **Retrieve tape** dialog box.
4. In the **Tape Barcode** field of the **Retrieve tape** dialog box, verify that the barcode identifies the virtual tape you want to retrieve.
5. For **Gateway**, choose the gateway you want to retrieve the archived tape to, and then choose **Proceed**.

The status of the tape changes from ARCHIVED to RETRIEVING. After all the data is moved, the status of the virtual tape in changes to RETRIEVED, and the tape appears in your gateway's VTL.

Note
Retrieved virtual tapes are read-only.

Restoring Data from the Tape

In this step, you restore data from the virtual tape to your client computer.

Note
This step does not use the NetBackup Administration Console that you used in the previous steps. Instead, you will use the Backup, Archive, and Restore software installed with Symantec NetBackup software.

To restore data

1. Start the Backup, Archive, and Restore software, and run it as an administrator.
2. Choose the **Select for Restore** tab.

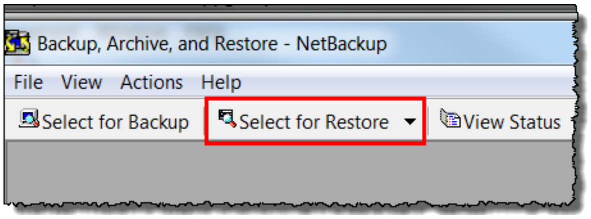

3. If you have previously backed up data, you will see backup icons in the **NetBackup History** pane. In this example, there is only one backup icon.

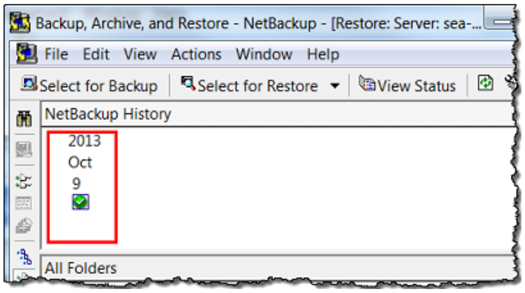
Backup, Archive, and Restore - NetBackup - [Restore: Server: sea-...
File Edit View Actions Window Help
Select for Backup
Select for Restore
View Status
NetBackup History
2013
Oct
9
All Folders

4. Select the backup icon that represents the backup you want to restore.
5. In the **All Folders** pane, select the folder you want to restore, and then choose the **Restore Marked Files** icon in the left pane. This icon is not labeled. The name appears when you rest your mouse on the icon.

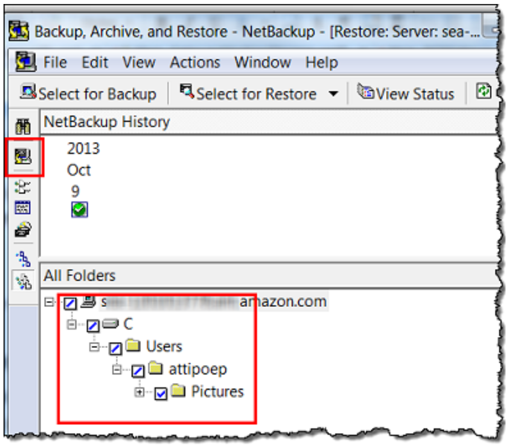

6. In the **Restore Marked Files** window, select the **Restore everything to a different location (maintaining existing structure)** button. This selection avoids overwriting your original data.

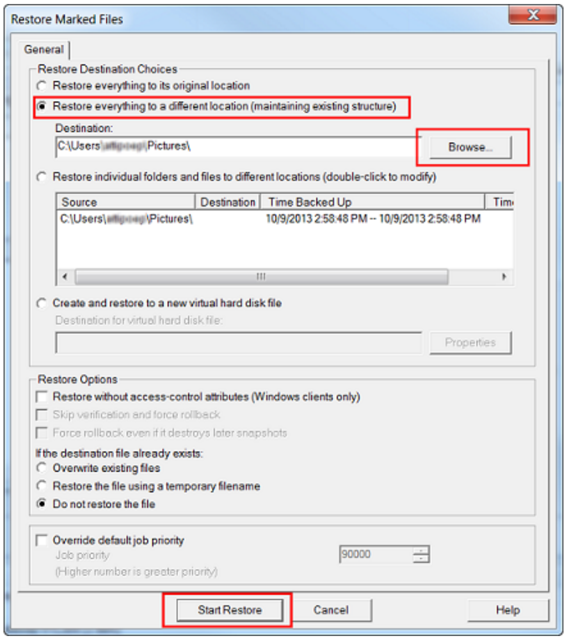

Restore Marked Files

General

Restore Destination Choices

○ Restore everything to its original location

● Restore everything to a different location (maintaining existing structure)

Destination:

C:\Users\[redacted]\Pictures\

○ Restore individual folders and files to different locations (double-click to modify)

Source	Destination	Time Backed Up	Tim
C:\Users\[redacted]\Pictures\		10/9/2013 2:58:48 PM -- 10/9/2013 2:58:48 PM	

○ Create and restore to a new virtual hard disk file

Destination for virtual hard disk file:

Restore Options

□ Restore without access-control attributes (Windows clients only)

□ Skip verification and force rollback

□ Force rollback even if it destroys later snapshots

If the destination file already exists:

○ Overwrite existing files

○ Restore the file using a temporary filename

● Do not restore the file

□ Override default job priority

Job priority 90000

(Higher number is greater priority)

7. For **Destination**, browse to the folder you want to restore the data to, and then choose **Start Restore**.

8. In the dialog box that appears, choose **Yes** to view the progress of the restore process.

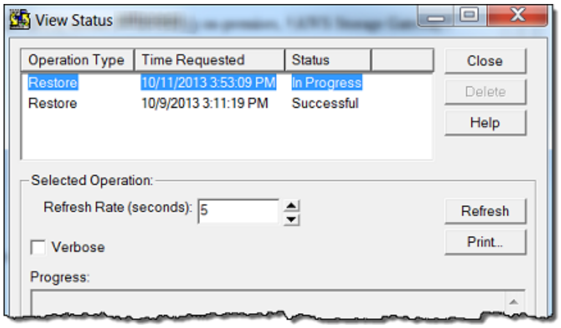

View Status [redacted]

Operation Type	Time Requested	Status
Restore	10/11/2013 3:53:09 PM	In Progress
Restore	10/9/2013 3:11:19 PM	Successful

Selected Operation:

Refresh Rate (seconds): 5

□ Verbose

Progress:

In the **View Status** window, you can see the status of the restore process. If the restore succeeds, the status changes to **Successful**.

Next Step

Clean Up Resources You Don't Need (p. 97)

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 V7, V8, or V9 Update 2 or later. 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. For detailed information about how to use the Veeam software, see the Veeam Backup & Replication documentation in the Veeam Help Center. In this topic, you can find out how to configure storage, write data to a tape, archive a tape, and restore the data.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

Configuring the Veeam Software 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 Connect Your Tape Gateway Devices to Your Windows Client (p. 59).

Updating VTL Device Drivers

By default, the Veeam V7 and V8 backup software does not recognize tape gateway devices. 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 (p. 274).

Discovering VTL Devices

For the Veeam 9 backup software, 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 Working with Tape Libraries.

To discover the VTL devices

1. In the Veeam software, choose **Backup Infrastructure**. When the tape gateway is connected, virtual tapes will be listed in the **Backup Infrastructure** tab.

   Note
   Depending on the version of the Veeam Backup & Replication backup software you are using, the user interface might differ somewhat from that shown in the screenshots in this documentation.

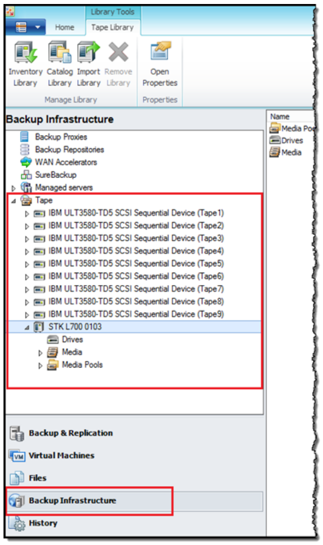
Library Tools
Home
Tape Library
Inventory Library
Catalog Library
Import Library
Remove Library
Open Properties
Manage Library
Properties
Backup Infrastructure
Backup Proxies
Backup Repositories
WAN Accelerators
SureBackup
Managed servers
Tape
IBM ULT3580-TD5 SCSI Sequential Device (Tape1)
IBM ULT3580-TD5 SCSI Sequential Device (Tape2)
IBM ULT3580-TD5 SCSI Sequential Device (Tape3)
IBM ULT3580-TD5 SCSI Sequential Device (Tape4)
IBM ULT3580-TD5 SCSI Sequential Device (Tape5)
IBM ULT3580-TD5 SCSI Sequential Device (Tape6)
IBM ULT3580-TD5 SCSI Sequential Device (Tape7)
IBM ULT3580-TD5 SCSI Sequential Device (Tape8)
IBM ULT3580-TD5 SCSI Sequential Device (Tape9)
STK L700 0103
Drives
Media
Media Pools
Name
Media Po
Drives
Media
Backup & Replication
Virtual Machines
Files
Backup Infrastructure
History

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 will 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 the Veeam Software

You are now ready to import tapes from your tape gateway into the Veeam backup software 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 the Veeam Software

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 Veeam documentation in the Veeam Help Center.

Archiving a Tape by Using the Veeam Software

When you archive a tape, AWS Storage 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 software—that is, the Veeam software.

To archive a tape in the Veeam library

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

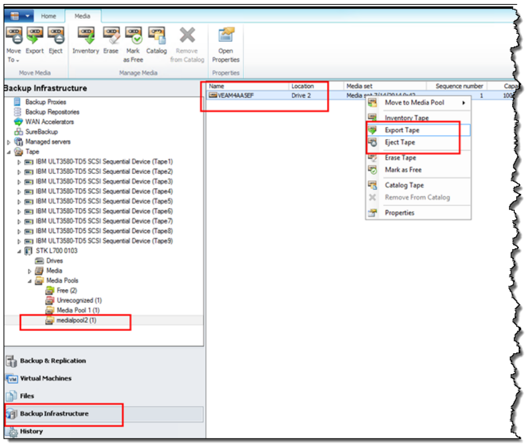

2. Open the context (right–click) menu for the tape that you want to archive, and then choose **Eject Tape**.
3. For **Ejecting tape** box, choose **Close**. The location of the tape changes from a tape drive to a slot.
4. Open the context (right–click) menu for the tape again, and then choose **Export**.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

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 the Archived Tape from Archive Back to Your Tape Gateway (p. 92).
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 Data from Tape in the Veeam Help Center.

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 (p. 99).

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 of the maintenance tasks on the gateway's local console, such as 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 preform specific maintenance tasks on the Amazon EC2 console, such as adding and removing Amazon EBS volumes. For more information, see

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 AWS Storage Gateway (p. 142). 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 Clean Up Resources You Don't Need (p. 97).

Clean Up Resources You Don't Need

If you created the gateway as 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? (p. 97).

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 (p. 186).

    a. Cancel retrieval for any tapes that have the **RETRIEVING** status in your gateway's VTL. For instructions, see Canceling Tape Retrieval (p. 277).
    b. Archive any tapes that have the **RETRIEVED** status in your gateway's VTL. For instructions, see Archiving Tapes (p. 276).
    c. Delete any remaining tapes from your gateway's VTL. For instructions, see Deleting Tapes (p. 117).
    d. Delete any tapes you have in the archive. For instructions, see Deleting Tapes (p. 117).
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 (p. 186).
3. Delete the AWS Storage Gateway VM from your on-premises host. If you created your gateway on an Amazon EC2 instance, terminate the instance.

Next Step

Where Do I Go from Here? (p. 97)

Managing Your Gateway

Managing your gateway includes tasks such as configuring cache storage and upload buffer space, working with volumes or virtual tapes, doing general maintenance, troubleshooting, and monitoring your gateway's performance. If you haven't created a gateway, see Getting Started (p. 10).

Topics

Managing Your File Gateway

Following, you can find information about how to manage your file gateway resources.

Topics

Adding a File Share

After your file gateway is activated and is running, you can add additional file shares. For information about now to add a file, see Creating a File Share (p. 22).

Deleting a File Share

If you no longer need a file share, you can delete it from the AWS Storage Gateway management console.

To delete a file share

1. On the AWS Storage Gateway management console, choose **File shares**, and select the file share you want to delete.
2. In the **Actions** menu, choose **Delete file share**.
3. In the confirmation dialog box, verify the file share you want to delete, then select the check box, and then choose **Delete**.

Updating a File Share

If you don't set metadata values for your files or directories, AWS Storage Gateway sets default metadata. You can change the default metadata on the AWS Storage Gateway management console by updating your file share.

To update a file share

1. On the AWS Storage Gateway management console, choose **File shares** and select the file share you want to update.
2. In the **Actions** menu, choose **Update file share**.
3. In the **Update file share dialog box**, provide the metadata information and choose **Save**.

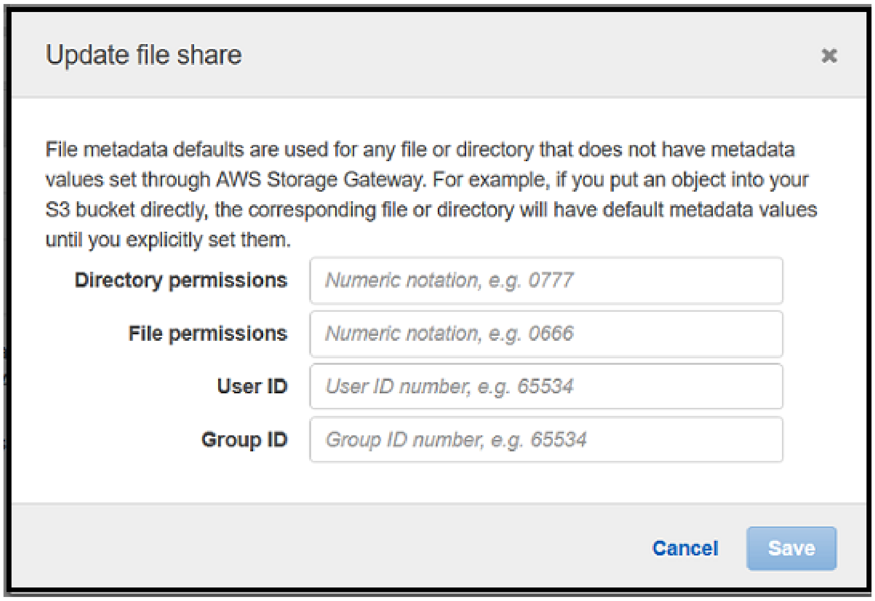

Update file share

File metadata defaults are used for any file or directory that does not have metadata values set through AWS Storage Gateway. For example, if you put an object into your S3 bucket directly, the corresponding file or directory will have default metadata values until you explicitly set them.

Directory permissions: Numeric notation, e.g. 0777

File permissions: Numeric notation, e.g. 0666

User ID: User ID number, e.g. 65534

Group ID: Group ID number, e.g. 65534

Understanding File Share Status

Each file share has an associated status that tells you at a glance what the health of the file share is. Most of the time, the status indicates that the file share is functioning normally and that no action is needed on your part. In some cases, the status indicates a problem that might or might not require action on your part.

You can see file share status on the AWS Storage Gateway console. File share status appears in the **Status** column for each file share in your gateway. A file share that is functioning normally has status as AVAILABLE.

The following table describes each file share status, and if and when you should take action based on the status. A file share should have AVAILABLE status all or most of the time it is in use.

Status	Meaning
AVAILABLE	The file share is configured properly and is available to use. The AVAILABLE status is the normal running status for a file share.
CREATING	The file share is being created and is not ready to be used. The CREATING status is transitional. No action is required. If file share is stuck in this status, it is probably because the gateway VM lost connection to AWS.
UPDATING	The file share configuration is being updated. If a file share is stuck in this status, it is probably because the gateway VM lost connection to AWS.
UPDATING	The file share configuration is being updated. If a file share is stuck in this status, it is probably because the gateway VM lost connection to AWS.
DELETING	The file share is being deleted. The File share is deleted until all data is uploaded to AWS. The DELETING status is transitional and no action is required.
UNAVAILABLE	The file share is in an unhealthy state. Issues such as role policy errors or if a Amazon S3 bucket the file share maps to does not exit, could cause the file share to go into an unhealthy state. Once the issue that caused the unhealthy state is resolved the file returns to AVAILABLE state.
ERROR	The Amazon S3 bucket you specified doesn’t exist or you don’t have permissions to access it.

Granting Access to an Amazon S3 Destination (File Gateway Only)

When you create a file share for file gateway, AWS Storage Gateway requires access to upload files into your Amazon S3 bucket. To grant this access, Storage Gateway creates an IAM access policy and role on your behalf. Your gateway assumes this role and performs the allowed actions on your behalf.

The following example trust policy allows Storage Gateway to assume the role.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "",
      "Effect": "Allow",
      "Principal": {
        "Service": "storagegateway.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "sts:ExternalId":"account-id"
        }
      }
    }
  ]
}
```

If you don’t want Storage Gateway to create a policy on your behalf, you create your own and attach it to your file share. For more information, see Creating a File Share (p. 22).

The following example policy allows Storage Gateway to perform all the Amazon S3 actions listed in the policy. The part of the statement allows all the actions listed to be performed on the S3 bucket named TestBucket. The second part allows the listed action on all objects in TestBucket.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "s3:GetAccelerateConfiguration",
                "s3:GetBucketLocation",
                "s3:GetBucketVersioning",
                "s3:ListBucket",
                "s3:ListBucketVersions",
                "s3:ListBucketMultipartUploads"
            ],
            "Resource": "arn:aws:s3:::TestBucket",
            "Effect": "Allow"
        },
        {
            "Action": [
                "s3:AbortMultipartUpload",
                "s3:DeleteObject",
                "s3:DeleteObjectVersion",
                "s3:GetObject",
                "s3:GetObjectVersion",
                "s3:ListMultipartUploadParts",
                "s3:PutObject"
            ],
            "Resource": "arn:aws:s3:::TestBucket/*",
            "Effect": "Allow"
        }
    ]
}
```

For more information, see Creating a Role to Delegate Permissions to an AWS Service in the *IAM User Guide*.

Managing Your Volume Gateway

Following, you can find information about how to manage you volume gateway resources.

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

Topics

- Creating a One-Time Snapshot (p. 104)
- Editing a Snapshot Schedule (p. 104)
- Deleting a Snapshot (p. 105)
- Understanding Volume Status (p. 111)

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

Resizing a volume is not supported. To change the size of a volume, create a snapshot of the volume, and then create a new cached volume from the snapshot. The new volume can be bigger than the volume from which the snapshot was created. For steps describing how to remove a volume, see To remove a volume (p. 103). For steps describing how to add a volume and preserve existing data, see Deleting a Volume (p. 103).

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

Adding a Volume

As your application needs grow, you might need to add more volumes to your gateway. As you add more volumes, you must consider the size of the cache storage and upload buffer you allocated to the gateway. The gateway must have sufficient buffer and cache space for new volumes. For more information, see Adding and Removing Upload Buffer (p. 144).

You can add volumes using the AWS Storage Gateway console or AWS Storage Gateway API. For information on using the AWS Storage Gateway API to add volumes, see CreateCachediSCSIVolume. The following procedure demonstrates using the console and assumes that you already have created a gateway.

For instructions on how to add a volume using the AWS Storage Gateway console, see Creating Volumes (p. 32).

Deleting a Volume

You might need to remove a volume as your application needs change—for example, if you migrate your application to use a larger storage volume. Before removing a volume, make sure that there are no applications currently writing to the volume. Also, make sure that there are no snapshots in progress for the volume. If a snapshot schedule is defined for the volume, you can check it on the **Snapshot Schedules** tab of the AWS Storage Gateway console. For more information, see Editing a Snapshot Schedule (p. 104).

You can remove volumes using the AWS Storage Gateway console or the AWS Storage Gateway API. For information on using the AWS Storage Gateway API to remove volumes, see DeleteVolume. The following procedure demonstrates using the console.

To remove a volume

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

2. On the **Volumes** tab, choose the volume and choose the confirmation box. Make sure the volume listed is the volume you intend to delete.

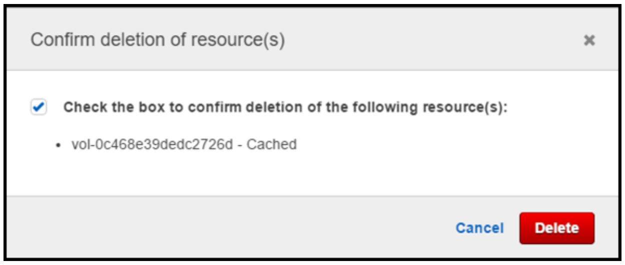

Confirm deletion of resource(s)

☑ Check the box to confirm deletion of the following resource(s):

- vol-0c468e39dedc2726d - Cached

3. Choose **Delete** to delete the volume.

Creating a One-Time Snapshot

In addition to scheduled snapshots, AWS Storage Gateway allows you to take one-time, ad hoc snapshots. By doing this, you can back up your storage volume immediately without waiting for the next scheduled snapshot.

To take a one-time snapshot of your storage volume

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Volumes**, and then choose the volume you want to create the snapshot from.
3. On the **Action** menu, choose **Create snapshot**.
4. In the **Create snapshot** dialog box, type the snapshot description, and then choose **Create snapshot**.

   You can verify that the snapshot was created using the console. For more information, see Finding a Snapshot (p. 255).
5. Your snapshot is listed in the **Snapshots** in the same row as the volume.

Editing a Snapshot Schedule

For stored volumes, AWS Storage Gateway creates a default snapshot schedule of once a day. This schedule helps ensure that your gateway can keep up with the rate of incoming write operations on your local storage volumes. You cannot remove the default snapshot schedule, but you can change it by specifying either the time the snapshot occurs each day or the frequency (every 1, 2, 4, 8, 12, or 24 hours), or both.

For cached volumes, AWS Storage Gateway does not create a default snapshot schedule. However, you can set up a snapshot schedule at any time if you need to. For cached volumes, because your data is stored in Amazon S3, you don't need snapshots or a snapshot schedule for disaster recovery purposes.

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

To edit the snapshot schedule for a volume

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

2. In the navigation pane, choose **Volumes**, and then choose the volume the snapshot was created from.
3. On the **Action** menu, choose **Edit snapshot schedule**.
4. In the **Edit snapshot schedule** dialog box, modify the schedule, and then choose **Save**.

Deleting a Snapshot

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

Topics

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

To delete multiple snapshots at a time, you can use one of the AWS SDKs that supports AWS Storage Gateway operations. For examples, see Deleting Snapshots by Using the AWS SDK for Java (p. 105), Deleting Snapshots by Using the AWS SDK for .NET (p. 107), and Deleting Snapshots by Using the AWS Tools for Windows PowerShell (p. 109).

Deleting Snapshots by Using the AWS SDK for Java

To delete many snapshots associated with a volume, you can use a programmatic approach. The example following demonstrates how to delete snapshots using the AWS SDK for Java. To use the example code, you should be familiar with running a Java console application. For more information, see Getting Started in the *AWS SDK for Java Developer Guide*. If you need to just delete a few snapshots, use the console as described in Deleting a Snapshot (p. 105).

Example : Deleting Snapshots by Using the AWS SDK for Java

```
            int daysBack2) {

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

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

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

            List<Snapshot> snapshots =
 describeSnapshotsResult.getSnapshots();
            System.out.println("volume-id = " + volumeId);
            for (Snapshot s : snapshots){
                StringBuilder sb = new StringBuilder();
                boolean meetsCriteria = !CompareDates(daysBack,
 s.getStartTime());
                sb.append(s.getSnapshotId() + ", " +
 s.getStartTime().toString());
                sb.append(", meets criteria for delete? " + meetsCriteria);
                sb.append(", deleted? ");
                if (!viewOnly & meetsCriteria) {
                    sb.append("yes");
                    DeleteSnapshotRequest deleteSnapshotRequest =
                        new
 DeleteSnapshotRequest().withSnapshotId(s.getSnapshotId());
                    ec2Client.deleteSnapshot(deleteSnapshotRequest);
                }
                else {
                    sb.append("no");
                }
                System.out.println(sb.toString());
            }
        }
    }

    private static String OutputVolumeInfo(VolumeInfo vi) {

        String volumeInfo = String.format(
                 "Volume Info:\n" +
                 "  ARN: %s\n" +
                 "  Type: %s\n",
                 vi.getVolumeARN(),
                 vi.getVolumeType());
        return volumeInfo;
     }

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

}
```

Deleting Snapshots by Using the AWS SDK for .NET

To delete many snapshots associated with a volume, you can use a programmatic approach. The following example demonstrates how to delete snapshots using the AWS SDK for .NET version 2 and 3. To use the example code, you should be familiar with running a .NET console application. For more information, see Getting Started in the *AWS SDK for .NET Developer Guide*. If you need to just delete a few snapshots, use the console as described in Deleting a Snapshot (p. 105).

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

```
                    }
                }
            }
            catch (AmazonEC2Exception ex)
            {
                Console.WriteLine(ex.Message);
            }
            return SelectedSnapshots;
        }

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

                    DeleteSnapshotRequest deleteSnapshotRequest = new
 DeleteSnapshotRequest(s.SnapshotId);
                    DeleteSnapshotResponse response =
 ec2Client.DeleteSnapshot(deleteSnapshotRequest);
                    Console.WriteLine("Volume: " +
                              s.VolumeId +
                              " => Snapshot: " +
                              s.SnapshotId +
                              " Response: "
                              + response.HttpStatusCode.ToString());
                }
            }
            catch (AmazonEC2Exception ex)
            {
                Console.WriteLine(ex.Message);
            }
        }

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

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

Deleting Snapshots by Using the AWS Tools for Windows PowerShell

To delete many snapshots associated with a volume, you can use a programmatic approach. The example following demonstrates how to delete snapshots using the AWS Tools for Windows PowerShell. To use the example script, you should be familiar with running a PowerShell script. For more information, see Getting Started in the *AWS Tools for Windows PowerShell*. If you need to delete just a few snapshots, use the console as described in Deleting a Snapshot (p. 105).

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

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

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

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

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

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

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

#ListVolumes
$volumesResult = Get-SGVolume -GatewayARN $gatewayARN
$volumes = $volumesResult.VolumeInfos
Write-Output("`nVolume List")
foreach ($volumes in $volumesResult)
  { Write-Output("`nVolume Info:")
    Write-Output("ARN:  " + $volumes.VolumeARN)
    write-Output("Type: " + $volumes.VolumeType)
  }

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

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

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

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

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

Understanding Volume Status

Each volume has an associated status that tells you at a glance what the health of the volume is. Most of the time, the status indicates that the volume is functioning normally and that no action is needed on your part. In some cases, the status indicates a problem with the volume that might or might not require action on your part. You can find information following to help you decide when you need to take action.

Topics

You can see volume status on the AWS Storage Gateway console or by using one of the AWS Storage Gateway API operations, for example DescribeCachediSCSIVolumes or DescribeStorediSCSIVolumes. The following example shows volume status on the AWS Storage Gateway console. Volume status appears in the **Status** column for each storage volume on your gateway. A volume that is functioning normally has status as AVAILABLE.

The following table describes each storage volume status, and if and when you should take action based on each status. The AVAILABLE status is the normal status of a volume, and a volume should have this status all or most of the time it is in use.

Status	Meaning
AVAILABLE	The volume is available for use. The AVAILABLE status is the normal running status for a volume.
BOOTSTRAPPING	The gateway is synchronizing data locally with a copy of the data stored in AWS. You typically do not need to take any action for this status, because the storage volume will automatically see the AVAILABLE status in most cases. The following are three scenarios when a volume status is BOOTSTRAPPING: • A gateway has unexpectedly shut down. • A gateway's upload buffer has been exceeded. In this scenario, bootstrapping occurs when your volume has the PASS THROUGH status and the amount of free upload buffer increases sufficiently. You can provide additional upload buffer space as one way to increase the percentage of free upload buffer space. In this particular scenario, the storage volume goes from PASS THROUGH to BOOTSTRAPPING to AVAILABLE status. You can continue to use this volume during this bootstrapping period. However, you cannot take snapshots of the volume at this point. • You are creating a gateway-stored volume and preserving existing local disk data. In this scenario, your gateway starts uploading all of the data up to AWS and the volume has the BOOTSTRAPPING status until all of the data from the local disk is copied to AWS. You can use the volume during this bootstrapping period. However, you cannot take snapshots of the volume at this point.
CREATING	The volume is currently being created and is not ready to be used. The CREATING status is transitional. No action is required.
DELETING	The volume is currently being deleted. The DELETING status is transitional. No action is required.

Status	Meaning
IRRECOVERABLE	An error occurred from which the volume cannot recover. For information on taking action in this situation, see Troubleshooting Volume Issues (p. 198).
PASS THROUGH	Data maintained locally is out of sync with data stored in AWS. Note that data written to a volume while the volume is in PASS THROUGH status remains in the cache until the volume status is BOOTSTRAPPING, and starts to upload to AWS when BOOTSTRAPPING status begins. The PASS THROUGH status can occur for several reasons, listed following: • The PASS THROUGH status occurs if your gateway has run out of upload buffer space. Your applications can continue to read from and write data to your storage volumes while the volumes have the PASS THROUGH status. However, the gateway is not writing any of your volume data to its upload buffer or uploading any of this data to AWS. The gateway will continue to upload any data written to the volume before the volume entered the PASS THROUGH status. Any pending or scheduled snapshots of a storage volume fail while the volume has the PASS THROUGH status. For information about what action to take when your storage volume has the PASS THROUGH status because the upload buffer has been exceeded, see Troubleshooting Volume Issues (p. 198). • The PASS THROUGH status occurs when there is more than one storage volume bootstrapping at once. Only one gateway storage volume can bootstrap at a time. For example, if you create two storage volumes and choose to preserve existing data on both of them, then the second storage volume has the PASS THROUGH status until the first storage volume finishes bootstrapping. In this scenario, you do not need to take action. Each storage volume will change to the AVAILABLE status automatically when it is finished being created. You can read and write to the storage volume while it has the PASS THROUGH or BOOTSTRAPPING status. • Infrequently, the PASS THROUGH status can indicate that a disk allocated for upload buffer use has failed. For information about what action to take in this scenario, see Troubleshooting Volume Issues (p. 198).
RESTORING	The volume is being restored from an existing snapshot. This status applies only for stored volumes. For more information, see How AWS Storage Gateway Works (Architecture) (p. 3). If you restore two storage volumes at the same time, both storage volumes show RESTORING as their status. Each storage volume will change to the AVAILABLE status automatically when it is finished being created. You can read and write to a storage volume and take a snapshot of it while it has the RESTORING status.

Status	Meaning
RESTORING PASS THROUGH	The volume is being restored from an existing snapshot and has encountered an upload buffer issue. This status applies only for stored volumes. For more information, see How AWS Storage Gateway Works (Architecture) (p. 3). One reason that can cause the RESTORING PASS THROUGH status is if your gateway has run out of upload buffer space. Your applications can continue to read from and write data to your storage volumes while they have the RESTORING PASS THROUGH status. However, no snapshots of a storage volume can occur during the RESTORING PASS THROUGH status period. For information about what action to take when your storage volume has the RESTORING PASS THROUGH status because upload buffer capacity has been exceeded, see Troubleshooting Volume Issues (p. 198). Infrequently, the RESTORING PASS THROUGH status can indicate that a disk allocated for an upload buffer has failed. For information about what action to take in this scenario, see Troubleshooting Volume Issues (p. 198).
UPLOAD BUFFER NOT CONFIGURED	The volume cannot be created or used, because the gateway does not have an upload buffer configured. For information on how to add upload buffer capacity for volumes in a gateway-cached setup, see Adding and Removing Upload Buffer (p. 144). For information on how to add upload buffer capacity for volumes in a gateway-stored setup, see Adding and Removing Upload Buffer (p. 144).

Cached Volume Status Transitions

The following state diagram describes the most common transitions between gateway-cached volume statuses. You don't need to understand the diagram in detail to use your gateway effectively. Rather, the diagram provides detailed information if you are interested in understanding more about how AWS Storage Gateway works.

The diagram shows neither the UPLOAD BUFFER NOT CONFIGURED status nor the DELETING status. Volume states in the diagram are represented by green, yellow, and red boxes. The colors are interpreted as follows.

Color	Volume Status
Green	The gateway is operating normally. The volume status is AVAILABLE or will eventually become AVAILABLE.
Yellow	The volume has the PASS THROUGH status, which indicates there is a potential issue with the storage volume. If this status appears because the upload buffer space is filled, then in some cases buffer space will become available again. At that point, the storage volume self-corrects to the AVAILABLE status. In other cases, you might have to add more upload buffer space to your gateway to allow the storage volume status to become AVAILABLE. For information on how to troubleshoot a case when upload buffer capacity has been exceeded, see Troubleshooting Volume

Color	Volume Status
	Issues (p. 198). For information on how to add upload buffer capacity, see Adding and Removing Upload Buffer (p. 144).
Red	The storage volume has the IRRECOVERABLE status. In this case, you should delete the volume. For information on how to do this, see To remove a volume (p. 103).

In the diagram, a transition between two states is depicted with a labeled line. For example, the transition from the CREATING status to the AVAILABLE status is labeled as *Create Basic Volume or Create Volume from Snapshot* and represents creating a cached volume. For more information about creating storage volumes, see Adding a Volume (p. 103).

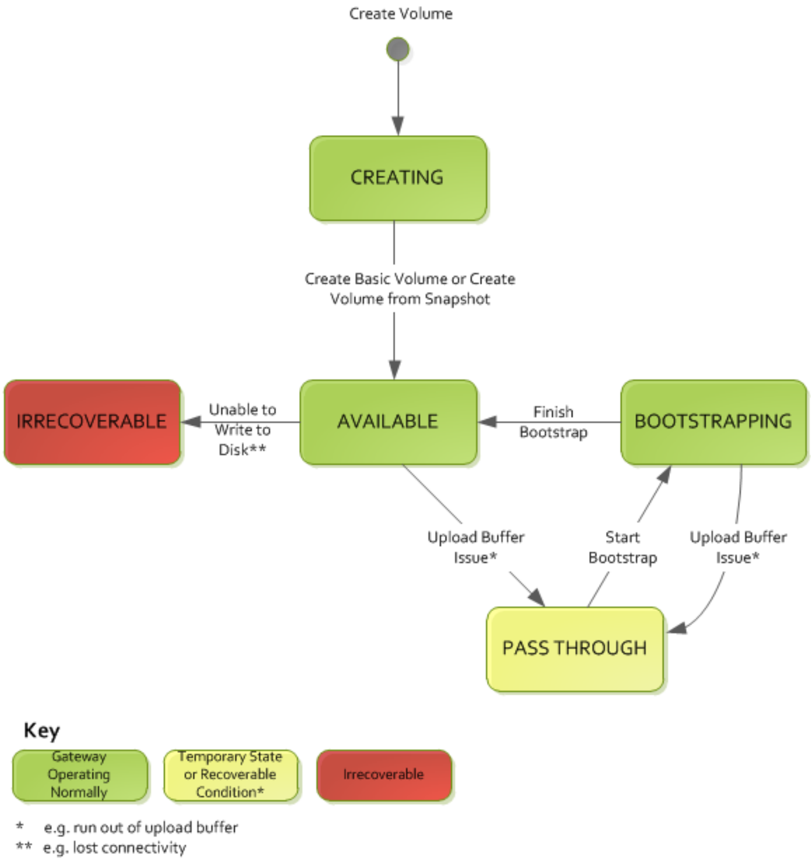
Create Volume
CREATING
Create Basic Volume or Create Volume from Snapshot
IRRECOVERABLE
Unable to Write to Disk**
AVAILABLE
Finish Bootstrap
BOOTSTRAPPING
Upload Buffer Issue*
Start Bootstrap
Upload Buffer Issue*
PASS THROUGH
Key
Gateway Operating Normally
Temporary State or Recoverable Condition*
Irrecoverable
* e.g. run out of upload buffer
** e.g. lost connectivity

> **Note**
> The volume status of PASS THROUGH is depicted as yellow in this diagram and does not match the color of this status icon in the **Status** box of the AWS Storage Gateway console.

Stored Volume Status Transitions

The following state diagram describes the most common transitions between gateway-stored volume statuses. You don't need to understand the diagram in detail to use your gateway effectively. Rather, the diagram provides detailed information if you are interested in understanding more about how AWS Storage Gateway works.

The diagram shows neither the UPLOAD BUFFER NOT CONFIGURED status nor the DELETING status. Volume states in the diagram are represented by green, yellow, and red boxes. The colors are interpreted as follows.

Color	Volume Status
Green	The gateway is operating normally. The volume status is AVAILABLE or will eventually become AVAILABLE.
Yellow	When you are creating a storage volume and preserving data, then the path from the CREATING status to the PASS THROUGH status occurs if another volume is bootstrapping. In this case, the volume with the PASS THROUGH status goes to the BOOTSTRAPPING status and then to the AVAILABLE status when the first volume is finished bootstrapping. Other than the specific scenario mentioned, yellow (PASS THROUGH status) indicates that there is a potential issue with the storage volume, the most common one being an upload buffer issue. If upload buffer capacity has been exceeded, then in some cases buffer space will become available again. At that point, the storage volume self-corrects to the AVAILABLE status. In other cases, you might have to add more upload buffer capacity to your gateway to return the storage volume to the AVAILABLE status. For information on how to troubleshoot a case when upload buffer capacity has been exceeded, see Troubleshooting Volume Issues (p. 198). For information on how to add upload buffer capacity, see Adding and Removing Upload Buffer (p. 144).
Red	The storage volume has the IRRECOVERABLE status. In this case, you should delete the volume. For information on how to do this, see Deleting a Volume (p. 103) or To remove the underlying local disk (Microsoft Hyper-V) (p. 270).

In the following diagram, a transition between two states is depicted with a labeled line. For example, the transition from the CREATING status to the AVAILABLE status is labeled as *Create Basic Volume* and represents creating a storage volume without preserving data or creating the volume from a snapshot. For more information about creating a storage volume, see To create a volume using the console (p. 266).

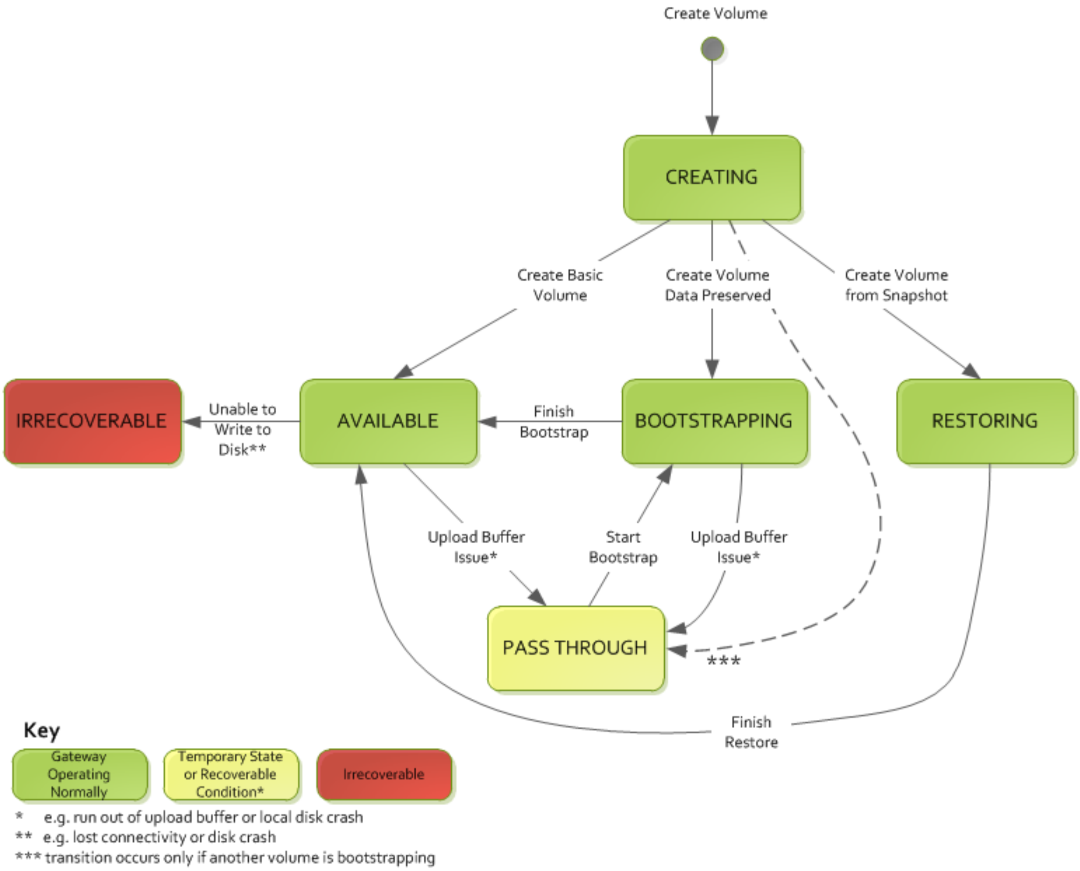
Create Volume
CREATING
Create Basic Volume
Create Volume Data Preserved
Create Volume from Snapshot
IRRECOVERABLE
Unable to Write to Disk**
AVAILABLE
Finish Bootstrap
BOOTSTRAPPING
RESTORING
Upload Buffer Issue*
Start Bootstrap
Upload Buffer Issue*
PASS THROUGH
***
Finish Restore
Key
Gateway Operating Normally
Temporary State or Recoverable Condition*
Irrecoverable
* e.g. run out of upload buffer or local disk crash
** e.g. lost connectivity or disk crash
*** transition occurs only if another volume is bootstrapping

Note
The volume status of PASS THROUGH is depicted as yellow in this diagram and does not match the color of this status icon in the **Status** box of the AWS Storage Gateway console.

Managing Your Tape Gateway

Following, you can find information about how to manage your tape gateway resources.

Topics

Adding Virtual Tapes

You can add tapes in your tape gateway when you need them. For more information, see Adding Virtual Tapes (p. 116). For information about tape gateway tape limits, see AWS Storage Gateway Limits (p. 296).

Retrieving Archived Tapes

To access data stored on an archived virtual tape, you must first retrieve the archived virtual tape you want to your tape gateway. AWS Storage Gateway provides one VTL per gateway. You can restore a tape to a tape gateway.

If you have multiple tape gateways in a 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**
> It takes up to 24 hours for the tape to be available in your tape gateway.

> **Note**
> There is a charge for retrieving tapes from archive. For detailed pricing information, see AWS Storage Gateway Pricing.

To retrieve an archived tape to your gateway

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Tapes**. You can search for all tapes that are archived to displays all virtual tapes that have been archived by all your gateways.
3. Choose the virtual tape you want to retrieve, and choose **Retrieve Tape** from the **Actions** menu.

   > **Note**
   > The status of the virtual tape 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 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 Amazon Glacier) 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.

Deleting Tapes

You can delete virtual tapes from your tape gateway by using the AWS Storage Gateway console.

> **Note**
> If the tape you want to delete from your tape gateway has a status of RETRIEVING, you must first cancel the retrieval before deleting the tape. For more information, see Canceling Tape Retrieval (p. 277). After the tape retrieval is canceled, the tape's status changes back to ARCHIVED. You can then delete the tape.
> 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 (p. 90). After the tape is ejected, the tape status changes back to ARCHIVED. You can then delete the tape.

To delete a virtual tape

> **Caution**
> This procedure will permanently delete the selected virtual tape.

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Tapes**.
3. Choose the virtual tape you want to delete.
4. On the **Action** menu, choose **Delete tape** to open the confirmation box.
5. 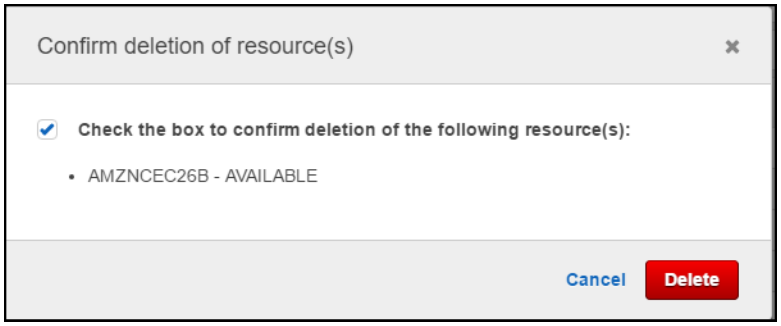

   Confirm deletion of resource(s)

   ☑ Check the box to confirm deletion of the following resource(s):

   - AMZNCEC26B - AVAILABLE

6. Make sure the tape listed is the tape you intend to delete, select the confirmation check box, and then choose **Delete**.

   After the tape is deleted, it disappears from the tape 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 or might not require action on your part. You can find information following to help you decide when you need to take action.

Topics

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 cannot be loaded into a tape drive, because it 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 be seen. When the upload is completed, the status changes to ARCHIVING.	Amazon S3
ARCHIVING	The virtual tape is being moved by AWS Storage Gateway to the archive, which is backed by Amazon	Data is being moved from Amazon S3 to Amazon Glacier

Status	Description	Tape Data Is Stored In
	Glacier. This process happens after the data upload to AWS is completed.	
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. **Note** The virtual tape can be retrieved only to a tape gateway.	Data is being moved from Amazon Glacier 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. You must disable the inaccessible tape gateway before you can recover the virtual tapes.	Amazon S3
IRRECOVERABLE	The virtual tape cannot be read from or written to. This status indicates an error in your tape gateway.	Amazon S3

Tape Status Information in the Archive

To determine the status of a virtual tape

1. Open the AWS 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.

The following table lists and describes 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. **Note** The virtual tape can be retrieved only to a tape gateway.
RETRIEVED	The virtual tape has been retrieved from the archive. The retrieved tape is read-only.

Monitoring Your Gateway

In this section, you can find information about how to monitor a gateway, including monitoring the volumes or tapes associated with the gateway and monitoring the upload buffer. You use the AWS Management Console to view metrics for your gateway. For example, you can view the number of bytes used in read and write operations, the time spent in read and write operations, and the time taken to retrieve data from the AWS 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.

Topics

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

Monitoring Your Volume Gateway

In this section, you can find information about how to monitor a gateway in a gateway-cached or gateway-stored setup, including monitoring the volumes associated with the gateway and monitoring the upload buffer. You use the AWS Management Console to view metrics for your gateway. For example, you can view the number of bytes used in read and write operations, the time spent in read and write operations, and the time taken to retrieve data from the AWS 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.

Topics

- Monitoring Cache Storage (p. 131)

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

Using Amazon CloudWatch Metrics

You can get monitoring data for your gateway using either the AWS Management Console or the CloudWatch API. The console displays a series of graphs based on the raw data from the CloudWatch API. You can also use the CloudWatch API through one of the 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`, `GatewayName`, and `VolumeId`. In the CloudWatch console, you can use the `Gateway Metrics` and `Volume Metrics` views to easily select gateway-specific and volume-specific dimensions. For more information about dimensions, see Dimensions in the *Amazon CloudWatch User Guide*>.
- The metric name, such as `ReadBytes`.

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

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

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

- Viewing Available Metrics
- Getting Statistics for a Metric
- Creating CloudWatch Alarms

Measuring Performance Between Your Application and Gateway

Data throughput, data latency, and operations per second are three measures that you can use to understand how your application storage that is using your gateway is performing. When you use the correct aggregation statistic, you can use Storage Gateway metrics to measure these values.

A *statistic* is an aggregation of a metric over a specified period of time. When you view the values of a metric in CloudWatch, use the `Average` statistic for data latency (milliseconds), use the `Sum` statistic for data throughput (bytes per second), and use the `Samples` statistic for input/output operations per second (IOPS). For more information, see Statistics in the *Amazon CloudWatch User Guide*.

The following table summarizes the metrics and corresponding statistic you can use to measure the throughput, latency, and IOPS between your applications and gateways.

Item of Interest	How to Measure
Throughput	Use the `ReadBytes` and `WriteBytes` metrics with the `Sum` CloudWatch statistic. For example, the `Sum` value of the `ReadBytes` metric over a sample period of 5 minutes divided by 300 seconds gives you the throughput as a rate in bytes per second.
Latency	Use the `ReadTime` and `WriteTime` metrics with the `Average` CloudWatch statistic. For example, the `Average` value of the `ReadTime` metric gives you the latency per operation over the sample period of time.
IOPS	Use the `ReadBytes` and `WriteBytes` metrics with the `Samples` CloudWatch statistic. For example, the `Samples` value of the `ReadBytes` metric over a sample period of 5 minutes divided by 300 seconds gives you IOPS.

For the average latency graphs and average size graphs, the average is calculated over the total number of operations (read or write, whichever is applicable to the graph) that completed during the period.

To measure the data throughput from an application to a volume

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

The following image shows the `ReadBytes` and `WriteBytes` metrics for a volume with the `Sum` statistic. In the image, the cursor over a data point displays information about the data point including its value and the number of bytes. Divide the bytes value by the **Period** value (5 minutes) to get the data throughput at that sample point. For the point highlighted, the read throughput is 2,384,199,680 bytes divided by 300 seconds, which is 7.6 megabytes per second.

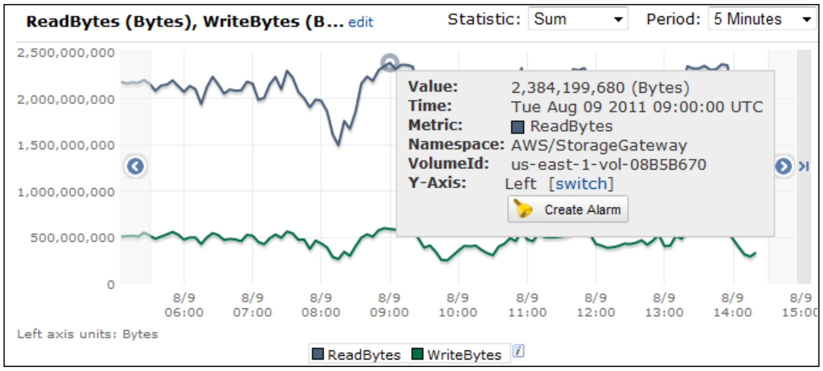

ReadBytes (Bytes), WriteBytes (B... edit

Statistic: Sum

Period: 5 Minutes

Value: 2,384,199,680 (Bytes)
Time: Tue Aug 09 2011 09:00:00 UTC
Metric: ReadBytes
Namespace: AWS/StorageGateway
VolumeId: us-east-1-vol-08B5B670
Y-Axis: Left [switch]

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

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

The following image shows the `ReadBytes` and `WriteBytes` metrics for a storage volume with the `Samples` statistic. In the image, the cursor over a data point displays information about the data point, including its value and the number of samples. Divide the samples value by the **Period** value (5 minutes) to get the operations per second at that sample point. For the point highlighted, the number of write operations is 24,373 bytes divided by 300 seconds, which is 81 write operations per second.

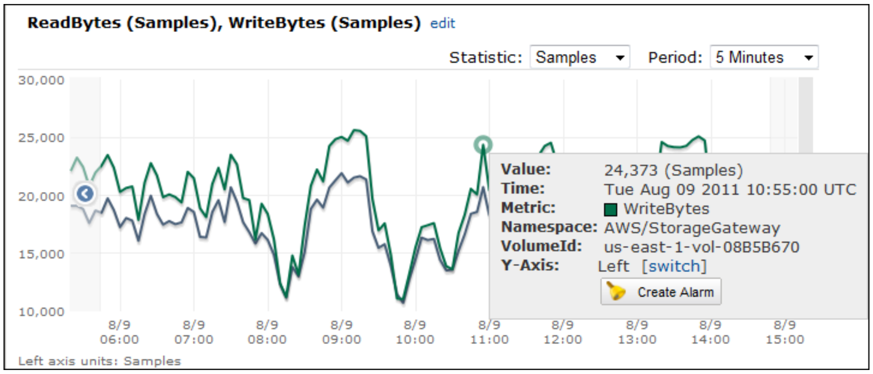

ReadBytes (Samples), WriteBytes (Samples) edit

Statistic: Samples

Period: 5 Minutes

Value: 24,373 (Samples)
Time: Tue Aug 09 2011 10:55:00 UTC
Metric: WriteBytes
Namespace: AWS/StorageGateway
VolumeId: us-east-1-vol-08B5B670
Y-Axis: Left [switch]

Measuring Performance Between Your Gateway and AWS

Data throughput, data latency, and operations per second are three measures that you can use to understand how your application storage using the Storage Gateway is performing. These three values can be measured using the Storage Gateway metrics provided for you when you use the correct aggregation statistic. The following table summarizes the metrics and corresponding statistic to use to measure the throughput, latency, and input/output operations per second (IOPS) between your gateway and AWS.

Item of Interest	How to Measure
Throughput	Use the `ReadBytes` and `WriteBytes` metrics with the `Sum` CloudWatch statistic. For example, the `Sum` value of the `ReadBytes` metric over a sample period of 5 minutes divided by 300 seconds gives you the throughput as a rate in bytes per second.
Latency	Use the `ReadTime` and `WriteTime` metrics with the `Average` CloudWatch statistic. For example, the `Average` value of the `ReadTime` metric gives you the latency per operation over the sample period of time.
IOPS	Use the `ReadBytes` and `WriteBytes` metrics with the `Samples` CloudWatch statistic. For example, the `Samples` value of the `ReadBytes` metric over a sample period of 5 minutes divided by 300 seconds gives you IOPS.
Throughput to AWS	Use the `CloudBytesDownloaded` and `CloudBytesUploaded` metrics with the `Sum` CloudWatch statistic. For example, the `Sum` value of the `CloudBytesDownloaded` metric over a sample period of 5 minutes divided by 300 seconds gives you the throughput from AWS to the gateway as bytes per second.
Latency of data to AWS	Use the `CloudDownloadLatency` metric with the `Average` statistic. For example, the `Average` statistic of the `CloudDownloadLatency` metric gives you the latency per operation.

To measure the upload data throughput from a gateway to AWS

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

The following image shows the `CloudBytesUploaded` metric for a gateway volume with the `Sum` statistic. In the image, the cursor over a data point displays information about the data point, including its value and bytes uploaded. Divide this value by the **Period** value (5 minutes) to get the throughput at that sample point. For the point highlighted, the throughput from the gateway to AWS is 555,544,576 bytes divided by 300 seconds, which is 1.7 megabytes per second.

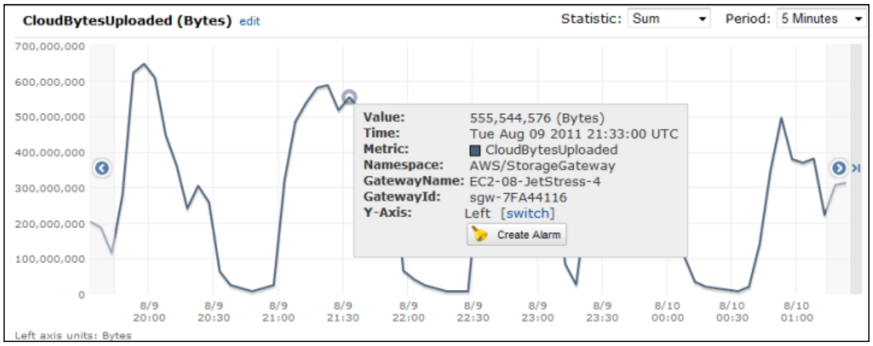

Statistic: Sum

Period: 5 Minutes

Value: 555,544,576 (Bytes)
Time: Tue Aug 09 2011 21:33:00 UTC
Metric: CloudBytesUploaded
Namespace: AWS/StorageGateway
GatewayName: EC2-08-JetStress-4
GatewayId: sgw-7FA44116
Y-Axis: Left [switch]

To measure the latency per operation of a gateway

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

To measure the data latency from a gateway to AWS

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

The resulting time-ordered set of data points contains the latency in milliseconds.

To set an upper threshold alarm for a gateway's throughput to AWS

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

To set an upper threshold alarm for reading data from AWS

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

Monitoring Your Tape Gateway

In this section, you can find information about 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

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	GatewayId, GatewayName	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 GatewayId and the GatewayName dimensions. Throughput and latency data of a tape gateway is based on all the virtual tapes in the tape gateway. 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

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 `ReadTime` and `WriteTime` metrics with the `Average` CloudWatch statistic. For example, the `Average` value of the `ReadTime` metric gives you the latency per operation over the sample period of time.
Throughput to AWS	Use the `CloudBytesDownloaded` and `CloudBytesUploaded` metrics with the `Sum` CloudWatch statistic. For example, the `Sum` value of the `CloudBytesDownloaded` 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 `CloudDownloadLatency` metric with the `Average` statistic. For example, the `Average` statistic of the `CloudDownloadLatency` 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.

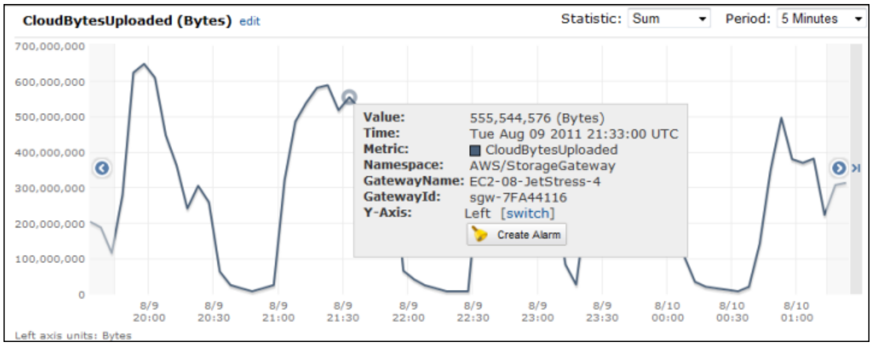

CloudBytesUploaded (Bytes)

Statistic: Sum

Period: 5 Minutes

Value: 555,544,576 (Bytes)
Time: Tue Aug 09 2011 21:33:00 UTC
Metric: ■ CloudBytesUploaded
Namespace: AWS/StorageGateway
GatewayName: EC2-08-JetStress-4
GatewayId: sgw-7FA44116
Y-Axis: Left [switch]

Left axis units: Bytes

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**.

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 proactively 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 gateway-cached and gateway-stored architectures. For more information, see How AWS Storage Gateway Works (Architecture) (p. 3).

> **Note**
> The `WorkingStoragePercentUsed`, `WorkingStorageUsed`, and `WorkingStorageFree` metrics represent the upload buffer for the gateway-stored volume setup only before the release of the cached-volume feature in Storage Gateway. Now you should 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 `UploadBufferPercentUsed`, `UploadBufferUsed`, and `UploadBufferFree` metrics with the `Average` statistic. For example, use the `UploadBufferUsed` with the `Average` statistic to analyze the storage usage over a time period.

To measure upload buffer percent 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.

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. In 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 Simple Notification Service (Amazon SNS) topic. For more information about Amazon SNS, see Set Up Amazon SNS.
   c. For **Topic**, type a descriptive name for the topic.
   d. Choose **Add Action**.
   e. Choose **Continue**.
6. Review the alarm settings, and then create the alarm.
   a. In the **Review** page of the Create Alarm Wizard, review the alarm definition, metric, and associated actions from this step. Associated actions include, 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 Simple Notification Service (Amazon SNS) email topic that is sent to the email address that you specified when creating the topic.

      The following image shows a notification.

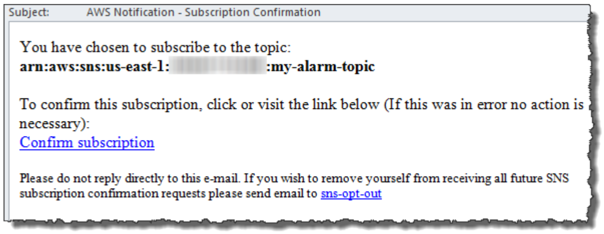

Subject: AWS Notification - Subscription Confirmation

You have chosen to subscribe to the topic:
arn:aws:sns:us-east-1:[redacted]:my-alarm-topic

To confirm this subscription, click or visit the link below (If this was in error no action is necessary):
Confirm subscription

Please do not reply directly to this e-mail. If you wish to remove yourself from receiving all future SNS subscription confirmation requests please send email to sns-opt-out

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 proactively add cache storage to a gateway.

You only monitor cache storage in the gateway-cached architecture. For more information, see How AWS Storage Gateway Works (Architecture) (p. 3).

Item of Interest	How to Measure
Total usage of cache	Use the `CachePercentUsed` and `TotalCacheSize` metrics with the `Average` statistic. For example, use the `CachePercentUsed` with the `Average` statistic to analyze the cache usage over a period of time. The `TotalCacheSize` metric changes only when you add cache to the gateway.
Percentage of read requests that are served from the cache	Use the `CacheHitPercent` metric with the `Average` statistic. Typically, you want `CacheHitPercent` to remain high.
Percentage of cache that is dirty—that is, it contains content that has not been uploaded to AWS	Use the `CachePercentDirty` metrics with the `Average` statistic. Typically, you want `CachePercentDirty` to remain low.

To measure the cache's percentage 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 cache's percentage 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 Storage Gateway Metrics

You can find information following about metrics you can use to monitor your gateway and volumes.

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 (p. 121).

The following table describes the Storage Gateway metrics that you can use to get information about your gateway. The entries in the table are grouped functionally by measure.

> **Note**
> The reporting period for these metrics is 5 minutes.

Metric	Description	Gateway-Cached	Gateway-Stored	Tape Gateway
`CacheHitPercent`	Percent of application reads served from the cache. The sample is taken at the end of the reporting period. Units: Percent	yes	no	yes
`CachePercentUsed`	Percent use of the gateway's cache storage. The sample is taken at the end of the reporting period. Units: Percent	yes	no	yes
`CachePercentDirty`	Percent of the gateway's cache that has not been persisted to AWS. The sample is taken at the end of the reporting period.	yes	no	yes

Metric	Description	Gateway-Cached	Gateway-Stored	Tape Gateway
	Units: Percent			
`CloudBytesDownloaded`	The total number of compressed bytes that the gateway downloaded from AWS during the reporting period. Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure input/ output operations per second (IOPS). Units: Bytes	yes	yes	yes
`CloudDownloadLatency`	The total number of milliseconds spent reading data from AWS during the reporting period. Use this metric with the `Average` statistic to measure latency. Units: Milliseconds	yes	yes	yes
`CloudBytesUploaded`	The total number of compressed bytes that the gateway uploaded to AWS during the reporting period. Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS. Units: Bytes	yes	yes	yes
`UploadBufferFree`	The total amount of unused space in the gateway's upload buffer. The sample is taken at the end of the reporting period. Units: Bytes	yes	no	yes
`CacheFree`	The total amount of unused space in the gateway's cache storage. The sample is taken at the end of the reporting period. Units: Bytes	yes	no	yes
`UploadBufferPercentUsed`	Percent use of the gateway's upload buffer. The sample is taken at the end of the reporting period. Units: Percent	yes	no	yes
`UploadBufferUsed`	The total number of bytes being used in the gateway's upload buffer. The sample is taken at the end of the reporting period. Units: Bytes	yes	no	yes

Metric	Description	Gateway-Cached	Gateway-Stored	Tape Gateway
`CacheUsed`	The total number of bytes being used in the gateway's cache storage. The sample is taken at the end of the reporting period. Units: Bytes	yes	no	yes
`QueuedWrites`	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. Units: Bytes	yes	yes	yes
`ReadBytes`	The total number of bytes read from your on-premises applications in the reporting period for all volumes in the gateway. Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS. Units: Bytes	yes	yes	yes
`ReadTime`	The total number of milliseconds spent to do read operations from your on-premises applications in the reporting period for all volumes in the gateway. Use this metric with the `Average` statistic to measure latency. Units: Milliseconds	yes	yes	yes
`TotalCacheSize`	The total size of the cache in bytes. The sample is taken at the end of the reporting period. Units: Bytes	yes	no	yes
`WriteBytes`	The total number of bytes written to your on-premises applications in the reporting period for all volumes in the gateway. Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS. Units: Bytes	yes	yes	yes

Metric	Description	Gateway-Cached	Gateway-Stored	Tape Gateway
`WriteTime`	The total number of milliseconds spent to do write operations from your on-premises applications in the reporting period for all volumes in the gateway. Use this metric with the `Average` statistic to measure latency. Units: Milliseconds	yes	yes	yes
`TimeSinceLastRecoveryPoint`	The time since the last available recovery point. For more information, see Using Recovery Snapshots for Your Gateway-Cached Setup (p. 201). Units: Seconds	yes	yes	no
`WorkingStorageFree`	The total amount of unused space in the gateway's working storage. The sample is taken at the end of the reporting period. **Note** Working storage applies only to the gateway-stored volume setup. The upload buffer applies to both the gateway-stored and gateway-cached volume setups. If you are working with both types of gateway setups, you might find it more convenient to use just the corresponding upload buffer metric, `UploadBufferFree`. Units: Bytes	no	yes	no
`WorkingStoragePercentUsed`	Percent used of the gateway's upload buffer. The sample is taken at the end of the reporting period. **Note** Working storage applies only to the gateway-stored volume setup. The upload buffer applies to both the gateway-stored and gateway-cached volume setups. If you are working with both types of gateway setups, you might find it more convenient to use just the corresponding upload buffer metric, `UploadBufferPercentUsed`. Units: Percent	no	yes	no

Metric	Description	Gateway-Cached	Gateway-Stored	Tape Gateway
WorkingStorageUsed	The total number of bytes being used in the gateway's upload buffer. The sample is taken at the end of the reporting period. **Note** Working storage applies only to the gateway-stored volume setup. The upload buffer applies to both the gateway-stored and gateway-cached volume setups. If you are working with both types of gateway setups, you might find it more convenient to use just the corresponding upload buffer metric, UploadBufferUsed. Units: Bytes	no	yes	no

Volume Metrics

You can find information following about the Storage Gateway metrics that cover a volume of a gateway. Each volume of a gateway has a set of metrics associated with it. Note that some volume-specific metrics have the same name as certain gateway-specific metrics. These metrics represent the same kinds of measurements but are scoped to the volume instead of the gateway. You must always specify whether you want to work with either a gateway or a volume metric before working with a metric. Specifically, when working with volume metrics, you must specify the VolumeId that identifies the storage volume for which you are interested in viewing metrics. For more information, see Using Amazon CloudWatch Metrics (p. 121).

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

Metric	Description	Gateway-Cached	Gateway-Stored
CacheHitPercent	Percent of application read operations from the volume that are served from cache. The sample is taken at the end of the reporting period. When there are no application read operations from the volume, this metric reports 100 percent. Units: Percent	yes	no
CachePercentUsed	The volume's contribution to the	yes	no

Metric	Description	Gateway-Cached	Gateway-Stored
	overall percent use of the gateway's cache storage. The sample is taken at the end of the reporting period. Use the `CachePercentUsed` metric of the gateway to view overall percent use of the gateway's cache storage. For more information, see Gateway Metrics (p. 132). Units: Percent		
`CachePercentDirty`	The volume's contribution to the overall percentage of the gateway's cache that has not been persisted to AWS. The sample is taken at the end of the reporting period. Use the `CachePercentDirty` metric of the gateway to view the overall percentage of the gateway's cache that has not been persisted to AWS. For more information, see Gateway Metrics (p. 132). Units: Percent	yes	no
`ReadBytes`	The total number of bytes read from your on-premises applications in the reporting period. Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS. Units: Bytes	yes	yes

Metric	Description	Gateway-Cached	Gateway-Stored
`ReadTime`	The total number of milliseconds spent to do read operations from your on-premises applications in the reporting period. Use this metric with the `Average` statistic to measure latency. Units: Milliseconds	yes	yes
`WriteBytes`	The total number of bytes written to your on-premises applications in the reporting period. Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS. Units: Bytes	yes	yes
`WriteTime`	The total number of milliseconds spent to do write operations from your on-premises applications in the reporting period. Use this metric with the `Average` statistic to measure latency. Units: Milliseconds	yes	yes
`QueuedWrites`	The number of bytes waiting to be written to AWS, sampled at the end of the reporting period. Units: Bytes	yes	yes

Logging AWS Storage Gateway API Calls by Using AWS CloudTrail

Storage Gateway is integrated with AWS CloudTrail, a service that captures API calls made by or on behalf of Storage Gateway in your AWS account and delivers the log files to an Amazon S3 bucket that you specify. CloudTrail captures API calls from the Storage Gateway console or from the Storage Gateway API. Using the information collected by CloudTrail, you can determine what request was made to Storage Gateway, the source IP address from which the request was made, who made the request, when it was made, and so on. To learn more about CloudTrail, including how to configure and enable it, see the *AWS CloudTrail User Guide*.

Storage Gateway Information in CloudTrail

When CloudTrail logging is enabled in your AWS account, API calls made to Storage Gateway actions are tracked in log files. Storage Gateway records are written together with other AWS service records in a log file. CloudTrail determines when to create and write to a new file based on a time period and file size.

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 log entry contains information about who generated the request. The user identity information in the log helps you determine whether the request was made with root or IAM user credentials, with temporary security credentials for a role or federated user, or by another AWS service. For more information, see the **userIdentity** field in the CloudTrail Event Reference in the *AWS CloudTrail User Guide*.

You can store your log files in your bucket for as long as you want, but you can also define Amazon S3 lifecycle rules to archive or delete log files automatically. By default, your log files are encrypted by using Amazon S3 server-side encryption (SSE).

You can choose to have CloudTrail publish Amazon Simple Notification Service (Amazon SNS) notifications when new log files are delivered if you want to take quick action upon log file delivery. For more information, see Configuring Amazon SNS Notifications.

You can also aggregate Storage Gateway log files from multiple AWS regions and multiple AWS accounts into a single Amazon S3 bucket. For more information, see Aggregating CloudTrail Log Files to a Single Amazon S3 Bucket.

Understanding Storage Gateway Log File Entries

CloudTrail log files can contain one or more log entries where each entry is made up of multiple JSON-formatted events. A log entry represents a single request from any source and includes information about the requested action, any parameters, the date and time of the action, and so on. The log entries are not guaranteed to be in any particular order. That is, they are not an ordered stack trace of the public API calls.

The following example shows a CloudTrail log entry that demonstrates the ActivateGateway action.

```
        {"Records":[{"eventVersion":"1.02","userIdentity":
{"type":"IAMUser","principalId":"AIDAII5AUEPBH2M7JTNVC","arn":"arn:aws:iam::111122223333:us
StorageGateway-team/JohnDoe",

 "accountId":"111122223333","accessKeyId":"AKIAIOSFODNN7EXAMPLE","userName":"JohnDoe"},"eve

 "eventSource":"storagegateway.amazonaws.com","eventName":"ActivateGateway","awsRegion":"us
east-1","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-1","activationKey":"EHFBX-1NDD0-P0IVU-PI259-DHK88","gatewayType":"VTL"},
        "responseElements":{"gatewayARN":"arn:aws:storagegateway:us-
east-1:111122223333:gateway/
cloudtrailgatewayvtl"},"requestID":"54BTFGNQI71987UJD2IHTCT8NF1Q8GLLE1QEU3KPGG6F0KSTAUU0",
        "eventID":"635f2ea2-7e42-45f0-
bed1-8b17d7b74265","eventType":"AwsApiCall","apiVersion":"20130630","recipientAccountId":"4
```

The following example shows a CloudTrail log entry that demonstrates the ListGateways action.

```
        {"Records":[{"eventVersion":"1.02","userIdentity":
{"type":"IAMUser","principalId":"AIDAII5AUEPBH2M7JTNVC","arn":"arn:aws:iam::111122223333:us
StorageGateway-team/JohnDoe",
        "accountId:"111122223333",
 "accessKeyId":"AKIAIOSFODNN7EXAMPLE","userName":"JohnDoe"},"eventTime":"2014-12-03T19:41:5

 "eventSource":"storagegateway.amazonaws.com","eventName":"ListGateways","awsRegion":"us-
east-1","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":"6U2N42CU37KAO8BG6V1I23FRSJ1Q
        "eventID":"f76e5919-9362-48ff-a7c4-
d203a189ec8d","eventType":"AwsApiCall","apiVersion":"20130630",
        "recipientAccountId":"444455556666"}]}
```

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 (p. 18).

Topics

Starting and Stopping Your Gateway

You can find information in the following about starting and stopping a gateway. For example, you might need to stop your gateway VM for maintenance, such as when applying a patch to your hypervisor. Before you power off the gateway VM, you should stop the gateway.

> **Note**
> Although this section focuses on starting and stopping your gateway using the AWS Storage Gateway Management Console, keep in mind that you can also start and stop your gateway by using your local console. If your gateway is hosted on an Amazon EC2 instance, see Logging In to Your Amazon EC2 Gateway Local Console (p. 180). If your gateway is hosted on a VM, see Logging In to Your AWS Storage Gateway Local Console (p. 157)

> **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.

To start a gateway

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose the gateway to start.
3. Choose **Details**. and then choose **Start gateway**.

To stop a gateway

1. Open the AWS 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.
3. On the **Action** menu, 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.

Managing Local Disks for Your AWS Storage Gateway

The gateway virtual machine (VM) uses the local disks that you allocate on-premises for buffering and storage. For cached volumes and tape gateways, you allocate two disks, one disk for the upload buffer and the other for cache storage. For stored volumes, you allocate one disk for the upload buffer.

Topics

The following table describes the types of local storage and the gateways that require each.

Local Storage	Description	Gateway Type
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.	• cached volumes • stored volumes • Tape Gateways
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	• cached volumes • Tape Gateways

Local Storage	Description	Gateway Type
	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 that use the same underlying physical storage resource (that is, 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, or that is backed by a less-performant RAID configuration such as RAID 1, can lead to poor performance in some situations when used to back both the cache storage and upload buffer.

After the initial configuration and deployment of your gateway, you might find that you need to adjust the local storage by adding or removing disks for an upload buffer or adding disks for cache storage.

Configuring Local Storage for Your Gateway

When you created your gateway, you allocated disks for your gateway to use as upload buffer or cache storage. The upload buffer and cache storage are created from local disks you provisioned for your gateway VM when you first created your gateway. After your gateway is up and running, you might decide to configure additional upload buffer or cache storage for your gateway. You use the suggested sizing formula in deciding the disk sizes. For more information on sizing storage, see Adding and Removing Upload Buffer (p. 144) or Adding Cache Storage (p. 146). If you are configuring local storage for the first time, see Configuring Local Disks (p. 57) for instructions.

Configuring an Upload Buffer or Cache Storage

After your gateway is activated, you might need to add additional disks and configure them as local storage. The following procedure shows you how to configure an upload buffer or cache storage for your gateway.

To configure upload buffer or cache storage

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Gateways**.
3. In the **Action** menu, choose **Edit local disks**.
4. In the Edit local disks dialog box, identify the disks you provisioned and decide which one you want to use for upload buffer or cached storage.

   > **Note**
   > For gateway-stored volumes, only the upload buffer is displayed because gateway-stored volumes have no cache disks.

5. In the drop-down list box, in the **Allocate to** column, choose **Upload Buffer** for the disk to use as upload buffer.

6. For gateways created with cached volumes and tape gateway, choose **Cache** for the disk you want to use as a cache storage.

   If you don't see your disks, choose the **Refresh** button.
7. Choose **Save** to save your configuration settings.

For stored volumes, you configure one of the two disks for use by your application's data and the other disk as an upload buffer.

Adding and Removing Upload Buffer

After you configure your initial gateway, you can allocate and configure additional upload buffer capacity or reduce the capacity as your application needs change. To learn more about how to size your upload buffer based on your application needs, see Sizing the Upload Buffer (p. 145).

Topics

Adding Upload Buffer Capacity

As your application needs change and you add more volume capacity, you might need to increase the gateway's upload buffer capacity as well. You can add more buffer capacity to your gateway without interrupting existing gateway functions. Note that when you add more upload buffer capacity, you do so with the gateway VM turned on. However, when you reduce the amount of upload buffer capacity, you must first turn off the VM. You can add more upload buffer capacity by using the Storage Gateway console or the Storage Gateway API:

- For information on adding buffer capacity with the console, see To configure upload buffer or cache storage (p. 143). This procedure assumes that your gateway has at least one local disk available on its VM that you can allocate as an upload buffer to the gateway.
- For information on adding buffer capacity with the API, see AddUploadBuffer.

Removing Upload Buffer Capacity

As your application needs change and you change the volume configuration for a gateway, you might need to decrease the gateway's upload buffer capacity. Or, a local disk allocated as upload buffer space might fail and you might need to remove that disk from your upload buffer and assign a new local disk. In both cases, you can remove buffer capacity using the Storage Gateway console.

The following procedure assumes that your activated gateway has at least one local disk allocated as an upload buffer for the gateway. In the procedure, you start on the Storage Gateway console, leave the console and use the VMware vSphere client or the Microsoft Hyper-V Manager to remove the disk, and then return to the console.

To find the ID of a disk allocated as an upload buffer

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Gateways**.
3. On the **Action** menu, choose **Edit local Disks**.
4. In the **Edit local disks** dialog box, note the value of the virtual device node for the local disk to be removed. You can find the node value in the **Disk ID** column.

You use the disk's virtual device node in the vSphere client to help ensure that you remove the correct disk.

5. Stop the gateway by following the steps in the Starting and Stopping Your Gateway (p. 141) procedure.

   Note
   Before you stop the gateway, ensure that no application is writing data to it and that no snapshots are in progress. You can check the snapshot schedule of volumes on the **Snapshot Schedules** tab of the Storage Gateway console. For more information, see Editing a Snapshot Schedule (p. 104).

6. To remove the underlying local disk, do one of the following procedures.

For a Gateway Hosted In	Do This
VMware ESXi	Follow the steps in To remove a disk allocated for the upload buffer (VMware ESXi) (p. 251).
Microsoft Hyper-V	Follow the steps in To remove an underlying disk allocated for the upload buffer (Microsoft Hyper-V) (p. 253).

7. On the Storage Gateway console, turn on the gateway.

   Important
   After removing a disk used as an upload buffer, you must turn the gateway back on before adding new disks to the VM.

8. On the **Volumes** tab of the Storage Gateway console, check that all volumes have a status of AVAILABLE (p.).

   After a gateway restart, a storage volume might go through the PASS THROUGH (p.) and BOOTSTRAPPING (p.) states as the gateway adjusts to the upload buffer disk that you removed. A volume that passes through these two states will eventually come to the AVAILABLE (p.) state. You can use a volume during the PASS THROUGH and BOOTSTRAPPING states. However, you cannot take snapshots of the volume in these states.

Sizing the Upload Buffer

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

Note
For volume gateways, when the upload buffer reaches its capacity, your volume goes to PASS THROUGH status. In this status, new data that your application writes is persisted locally but not uploaded to AWS immediately. Thus, you cannot take new snapshots. When the upload buffer capacity frees up, the volume goes through BOOTSTRAPPING status. In this status, any new data that was persisted locally is uploaded to AWS. Finally, the volume returns to ACTIVE status. Storage Gateway then resumes normal synchronization of the data stored locally with the copy stored in AWS, and you can start taking new snapshots. For more information about volume status, see Understanding Volume Status (p. 111).
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, 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 enabled bandwidth throttling. This rate should be adjusted for compression. When uploading data to AWS, the gateway applies data compression where possible. For example, if your application data is text-only, you might get 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.

If your incoming rate is higher than the outgoing rate, or if the formula returns a value less than 150 GiB, we strongly recommend that you allocate at least 150 GiB of upload buffer space.

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

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

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

Note that you can initially use this approximation to determine the disk size that you want to allocate to the gateway as upload buffer space. Add more upload buffer space as needed using the Storage Gateway console. Also, you can use the Amazon CloudWatch operational metrics to monitor upload buffer usage and determine additional storage requirements. For information on metrics and setting the alarms, see Monitoring the Upload Buffer (p. 129).

If you decide that you need to change your upload buffer capacity, take one of the following actions.

To	Do This
Add more upload buffer capacity to your gateway.	Follow the steps in Adding Upload Buffer Capacity (p. 144).
Remove a disk allocated as upload buffer space.	Follow the steps in Removing Upload Buffer Capacity (p. 144).

Adding 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.

Important
Only gateways created with cached volumes and tape gateways require cache storage.

Topics

- Sizing Cache Storage (p. 148)
- Adding Cache Storage for Your Gateway (p. 148)

The following diagram highlights the cache storage in the larger picture of the gateway-cached architecture. For more information, see How AWS Storage Gateway Works (Architecture) (p. 3).

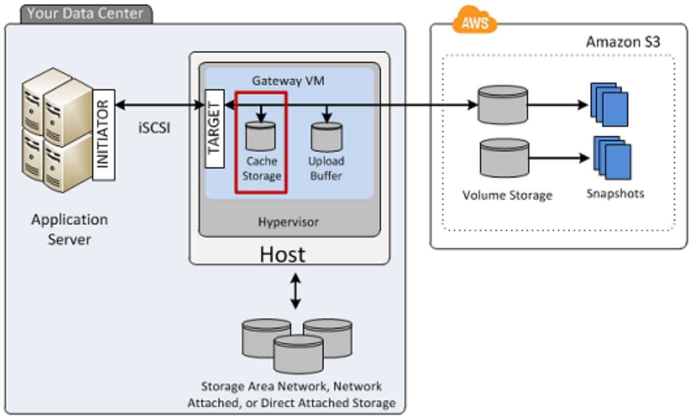

The following diagram highlights the cache storage in the larger picture of the tape gateway architecture. For more information, see How AWS Storage Gateway Works (Architecture) (p. 3).

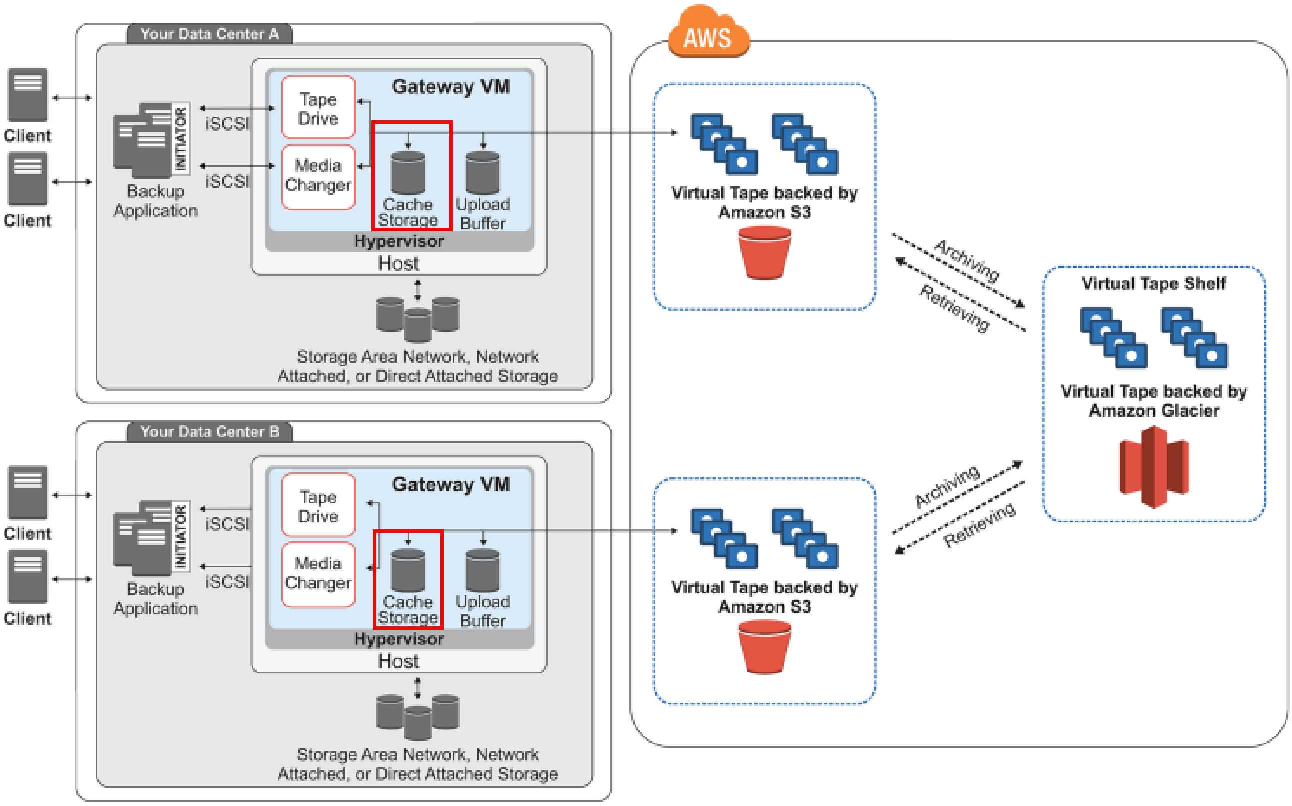

The amount of cache storage your gateway requires depends on how much of your application data you want to provide low-latency access to. The cache storage must be at least the size of the upload buffer. This guideline helps ensure that the cache storage is large enough to persistently hold all data that has not yet been uploaded to Amazon S3. When your cache storage has filled up with dirty data (that is, data that has not been uploaded to AWS), application write operations to your volumes or tapes are blocked until more cache storage becomes available. However, application read operations from the volume or tapes are still allowed.

Here are some guidelines you can follow to help ensure you have adequate cache storage allocated for your gateway.

- **Use the sizing formula.** – As your application needs change, you should periodically review the recommended formula for sizing cache storage. For more information, see Sizing Cache Storage (p. 148).
- **Use Amazon CloudWatch metrics.** – You can proactively avoid filling up cache storage with dirty data by monitoring how cache storage is being used—particularly, by reviewing cache misses. CloudWatch provides usage metrics such as the `CachePercentDirty` and `CacheHitPercent` metrics for monitoring how much of the gateway's cache storage has not been uploaded to Amazon S3. You can set an alarm to trigger a notification to you when the percentage of the cache that is dirty exceeds a threshold or the cache hit percentage falls below a threshold. Both of these can indicate that the cache storage size is not adequate for the gateway. For a full list of Storage Gateway metrics, see Understanding Storage Gateway Metrics (p. 132).

Sizing Cache Storage

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 Sizing the Upload Buffer (p. 145).

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 (p. 131).

If you decide that you need to increase your gateway's cache storage capacity, follow the steps in Adding Cache Storage for Your Gateway (p. 148).

Adding Cache Storage for Your Gateway

After you configure your initial gateway cache storage as described in Configuring an Upload Buffer or Cache Storage (p. 143), you can add cache storage to your gateway as your application needs change. To learn more about how to size your cache storage based on your application needs, see Adding Cache Storage (p. 146).

You can add more cache storage to your gateway without interrupting existing gateway functions and with the gateway VM turned on.

> **Note**
> Removing a disk allocated as cache storage is currently not supported.

You can add more cache storage by using the Storage Gateway console or the Storage Gateway API:

- For information on adding cache storage using the console, To configure upload buffer or cache storage (p. 143). This procedure assumes that your activated gateway has at least one local disk available on its VM that you can allocate as cache storage for the gateway.
- For information on adding cache storage by using the API, see AddCache.

Optimizing Gateway Performance

You can find information following about how to optimize the performance of your gateway. The guidance is based on adding resources to your gateway and adding resources to your application server.

Add Resources to Your Gateway

Use higher-performance disks

To optimize gateway performance, you can add high performance disks such as serial attached SCSI (SAS) disks and solid-state drives (SSDs). You can also attach virtual disks to your VM directly from a storage area network (SAN) instead of through the VMware Virtual Machine File System (VMFS) layer or the Microsoft Hyper-V New Technology File System (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 (p. 127).

> **Note**
> CloudWatch metrics are not available for all gateways. For information about gateway metrics, see Monitoring Your Gateway (p. 120)

Add CPU resources to your gateway host

The minimum requirement for a gateway host server is four virtual processors. To optimize gateway performance, you should confirm that the four virtual processors that are assigned to the gateway VM are backed by four cores and that you are not oversubscribing the CPUs of the host server. When you add additional CPUs to your gateway host server, you increase the processing capability of the gateway to deal with, in parallel, both storing data from your application to your local storage and uploading this data to Amazon S3. Additional CPUs also 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.

AWS Storage Gateway supports using eight CPUs in your gateway host server. You can use eight CPUs to significantly improve the performance of your gateway. We recommend the following gateway configuration for your gateway host server:

- Eight CPUs
- 7.5 GiB RAM
- Disk 1 attached to paravirtual controller 1, to be used as the gateway cache as follows:
  - Use RAID 10 backed by 15,000 RPM 6 Gbps disks with a stripe size of 256 K each
  - Use a Virtual Machine File System version 5 (VMFS5) data store created on the RAID 10
  - Use a virtual disk created on VMFS5 as the cache for the gateway
- Disk 2 attached to paravirtual controller 2, to be used as the gateway upload buffer as follows:
  - Use RAID 10 backed by 15,000 RPM 6 Gbps disks with a stripe size of 512 K each
  - Use a VMFS5 data store created on the RAID 10
  - Use a virtual disk created on VMFS5 as the upload buffer for the gateway
- Network adapter 1 configured on VM network 1:
  - Use VM network 1 and add a 1 Gbps network adapter 1 to be used for ingestion
- Network adapter 2 configured on VM network 2:
  - Use VM network 2 and add a 1 Gbps network adapter 2 to be used to connect to AWS

Back gateway virtual disks with separate physical disks

When you provision disks in a gateway setup, we strongly recommend that you do not 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, to use as an 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, 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, or that is backed by a less-performant RAID configuration such as RAID 1, can lead to poor performance—for example, when used to back both the cache storage and upload buffer in a gateway setup.

Change the volumes configuration

For volumes gateways, if you find that adding more volumes to a gateway reduces the throughput to the gateway, consider adding the volumes to a separate gateway. In particular, if a volume is used for a high-throughput application, consider creating a separate gateway for the high-throughput application. However, as a general rule, you should not use one gateway for all of your high-throughput applications and another gateway for all of your low-throughput applications. To measure your volume throughput, use the `ReadBytes` and `WriteBytes` metrics. For more information on these metrics, see Measuring Performance Between Your Application and Gateway (p. 122).

Add Resources to Your Application Environment

Increase the bandwidth between your application server and your gateway

To optimize gateway performance, ensure that the network bandwidth between your application and the gateway can sustain your application needs. You can use the `ReadBytes` and `WriteBytes` metrics of the gateway to measure the total data throughput (for more information on these metrics, see Measuring Performance Between Your Tape Gateway and AWS (p. 127)). 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 make use of additional CPU resources, then adding more CPUs can help your application to scale its I/O load.

Managing Bandwidth for Your Gateway

You can limit (or throttle) the upload throughput from the gateway to AWS or the download throughput from your AWS to your gateway. Using bandwidth throttling helps you to control the amount of network bandwidth used by your gateway. 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 AWS 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. As described directly following, you can change these limits by using the AWS Storage Gateway console. Or, for information about changing bandwidth rate limits programmatically, see the following topics.

Topics

- Updating Gateway Bandwidth Rate Limits Using the AWS SDK for .NET (p. 153)
- Updating Gateway Bandwidth Rate Limits Using the AWS Tools for Windows PowerShell (p. 155)

To change a gateway's bandwidth throttling using the console

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Gateways**, and then choose the gateway you want to manage.
3. On the **Action** menu, choose **Edit Bandwidth Rate Limit**.
4. In the **Edit Rate Limits** dialog box, type new limit values, and then choose **Save**. Your changes appear in the **Details** tab for your gateway.

Updating Gateway Bandwidth Rate Limits Using the AWS SDK for Java

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 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*.

The following Java code example updates a gateway's bandwidth rate limits. You need to update the code and provide the service endpoint, your gateway Amazon Resource Name (ARN), and the upload and download limits. For a list of AWS service endpoints you can use with AWS Storage Gateway, see Regions and Endpoints in the *AWS General Reference*.

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

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


public class UpdateBandwidthExample {

    public static AWSStorageGatewayClient sgClient;

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

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

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

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

        // Create a storage gateway client
        sgClient = new AWSStorageGatewayClient(new PropertiesCredentials(

 ListDeleteVolumeSnapshotsExample.class.getResourceAsStream("AwsCredentials.properties")));

        sgClient.setEndpoint(serviceURL);

        UpdateBandwidth(gatewayARN, uploadRate, downloadRate);

    }

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

            UpdateBandwidthRateLimitResult updateBandwidthRateLimitResult =
 sgClient.updateBandwidthRateLimit(updateBandwidthRateLimitRequest);
            String returnGatewayARN =
 updateBandwidthRateLimitResult.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 bandwith.\n" +
 ex.toString());
        }
    }
}
```

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

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 by using the AWS Software Development Kit (SDK) for .NET. To use the example code, you should be familiar with running a .NET console application. For more information, see Getting Started in the *AWS SDK for .NET Developer Guide*.

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

The following C# code example updates a gateway's bandwidth rate limits. You need to update the code and provide the service endpoint, your gateway Amazon Resource Name (ARN), and the upload and download limits. For a list of AWS service endpoints you can use with AWS Storage Gateway, see Regions and Endpoints in the *AWS General Reference*.

```
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 bandwith.\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 Limits by Using the AWS Tools for Windows PowerShell

The following PowerShell script example updates a gateway's bandwidth rate limits. You need to update the script and provide your gateway Amazon Resource Name (ARN), and the upload and download limits.

```
<#
.DESCRIPTION
    Update Gateway bandwidth limits.

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

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

$UploadBandwidthRate = 51200
$DownloadBandwidthRate = 102400
$gatewayARN = "*** provide gateway ARN ***"

#Update Bandwidth Rate Limits
Update-SGBandwidthRateLimit -GatewayARN $gatewayARN `
                            -AverageUploadRateLimitInBitsPerSec
 $UploadBandwidthRate `
                            -AverageDownloadRateLimitInBitsPerSec
 $DownloadBandwidthRate

$limits =  Get-SGBandwidthRateLimit -GatewayARN $gatewayARN

Write-Output("`nGateway: " + $gatewayARN);
Write-Output("`nNew Upload Rate: " +
 $limits.AverageUploadRateLimitInBitsPerSec)
Write-Output("`nNew Download Rate: " +
 $limits.AverageDownloadRateLimitInBitsPerSec)
```

Managing Gateway Updates Using the AWS Storage Gateway Console

AWS Storage Gateway periodically deploys important updates and patches to your gateway that must be applied. AWS notifies you with a message on the AWS Storage Gateway console and by email in advance of any updates to your gateway. Software updates force a restart of your gateway, which typically takes a few minutes to complete. You don't have take any action during this update.

You can choose to let AWS Storage Gateway apply updates according to the maintenance schedule for your gateway, or you can apply the updates yourself. When you deploy and activate your gateway, a default weekly maintenance schedule is set. You can modify this schedule at any time by choosing **edit time** next to **Maintenance Start Time** on the **Action** menu in the console.

When updates are available, the **Details** tab displays a maintenance message and the **Apply Update Now** button. The date and time the last successful software update was applied to the gateway appears in the **Details** tab.

> **Important**
> A software update forces your gateway to restart your gateway. You can minimize the chance of any disruption to your applications by increasing your iSCSI initiators' timeouts. For more information about increasing iSCSI initiator timeouts for Windows and Linux, see Customizing Your Windows iSCSI Settings (p. 280) and Customizing Your Linux iSCSI Settings (p. 45).

To edit the maintenance time window

1. In the navigation menu, choose **Gateways**, and then choose the gateway you want to edit the time for.
2. On the **Action** menu, choose **Edit maintenance window**.
3. Edit the values for **Day of the week** and **Time**. Your changes appear in the **Details** tab for the gateway.

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.

Topics

- Performing Maintenance Tasks on the VM Local Console (p. 156)
- Performing Maintenance Tasks on the Amazon EC2 Gateway Local Console (p. 180)

Performing Maintenance 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.

Topics

- Logging In to Your AWS Storage Gateway Local Console (p. 157)

- Routing Your On-Premises Gateway Through a Proxy (p. 161)
- Configuring Your Gateway Network (p. 162)
- Testing Your Gateway Connection to the Internet (p. 165)
- Synchronizing Your Gateway VM Time (p. 167)
- Running Storage Gateway Commands on the Local Console (p. 168)
- Viewing Your Gateway System Resource Status (p. 170)
- Configuring Network Adapters for Your Gateway (p. 171)
- Creating a Volume on a Gateway with Multiple Network Adapters (p. 179)

Logging In to Your AWS Storage Gateway Local Console

Some gateway maintenance tasks require that you log in to your gateway's local console. You can access the local console through your hypervisor client software.

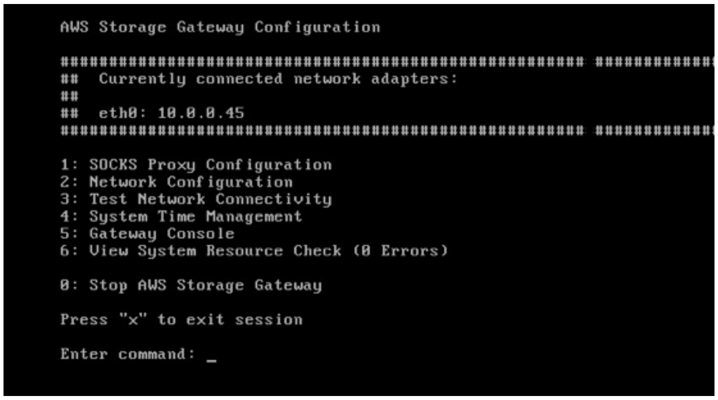

```
AWS Storage Gateway Configuration

###################################################### ############
##  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: _
```

In this topic, we show you how to access the local console of a gateway hosted in these hypervisors:

- VMware ESXi—for more information, see To access your gateway's local console (VMware ESXi) (p. 157).
- Microsoft Hyper-V—for more information, see To access your gateway's local console (Microsoft Hyper-V) (p. 158).

To access your gateway's local console (VMware ESXi)

1. In the VMware vSphere client, select your gateway VM.
2. Ensure 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.

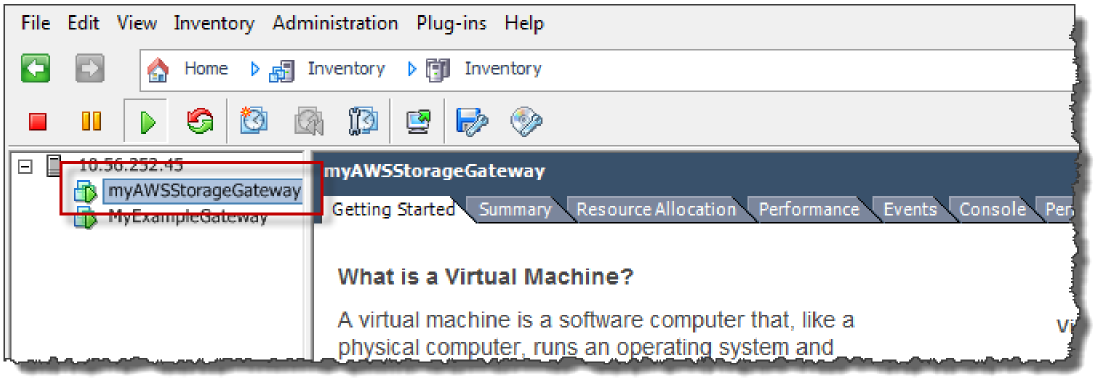

What is a Virtual Machine?

A virtual machine is a software computer that, like a physical computer, runs an operating system and

3. Choose the **Console** tab.

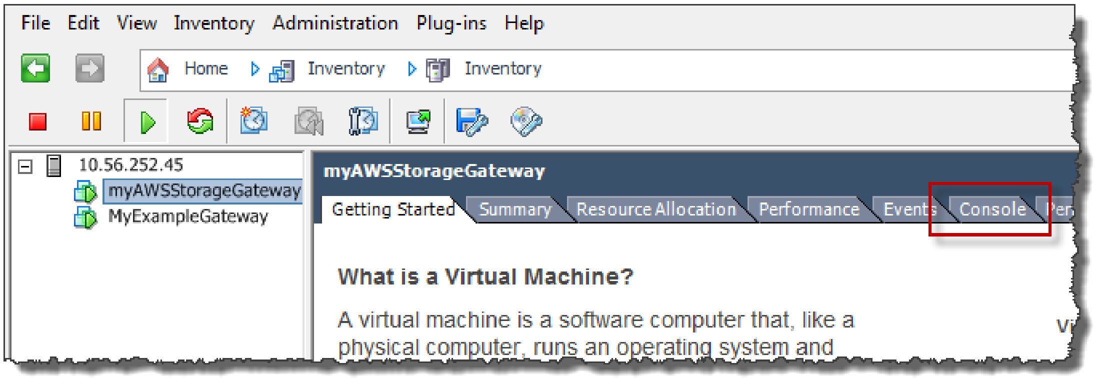

What is a Virtual Machine?

A virtual machine is a software computer that, like a physical computer, runs an operating system and

4. 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**.

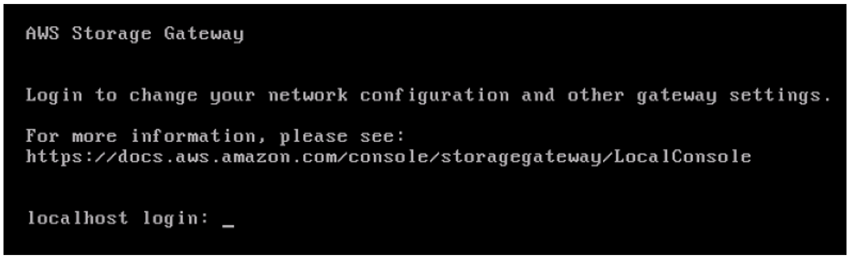

```
AWS Storage Gateway

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: _
```

5. To log in using the default credentials, continue to the procedure Logging in to the Local Console Using Default Credentials (p. 159).

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. Ensure 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.

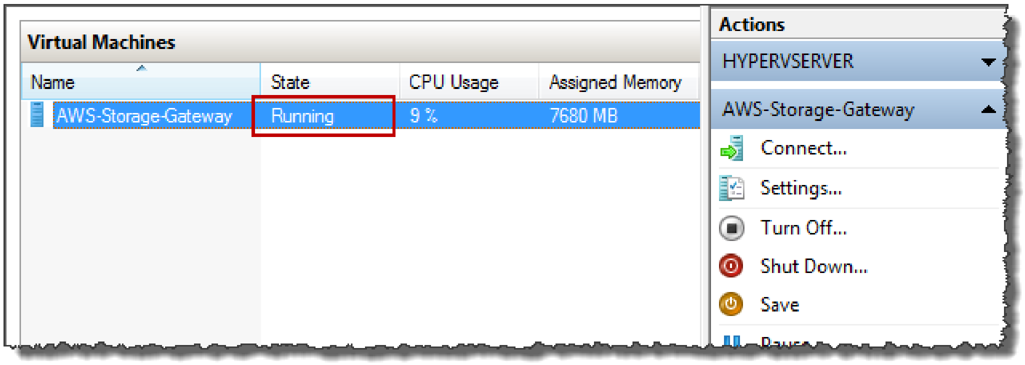

Virtual Machines

Name	State	CPU Usage	Assigned Memory
AWS-Storage-Gateway	Running	9 %	7680 MB

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

   The **Virtual Machine Connection** window appears. If an authentication window appears, type the user name and password provided to you by the hypervisor administrator.

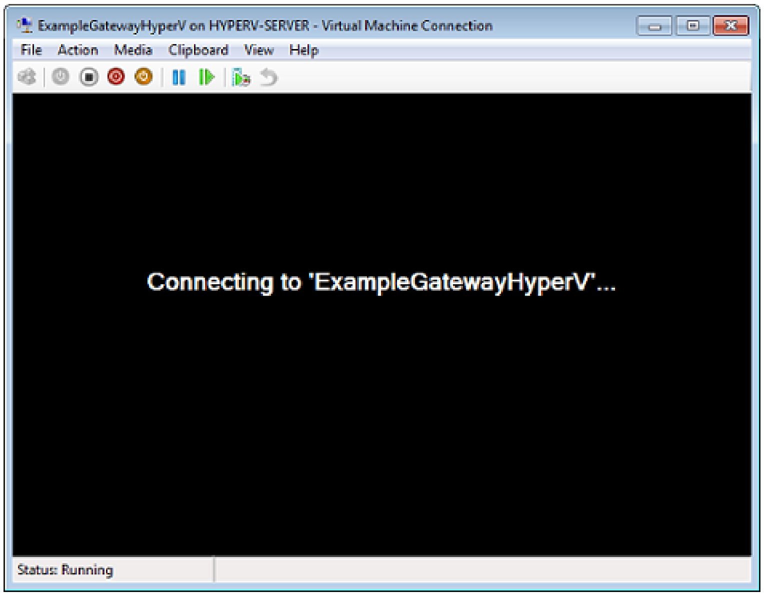

ExampleGatewayHyperV on HYPERV-SERVER - Virtual Machine Connection
File
Action
Media
Clipboard
View
Help

Connecting to 'ExampleGatewayHyperV'...

Status: Running

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

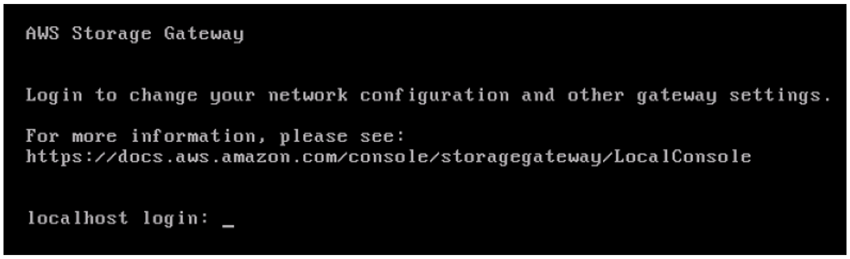

```
AWS Storage Gateway

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: _
```

5. To log in default credentials, continue to the procedure Logging in to the Local Console Using Default Credentials (p. 159).

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 user name and password 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 enables 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 (p. 161).

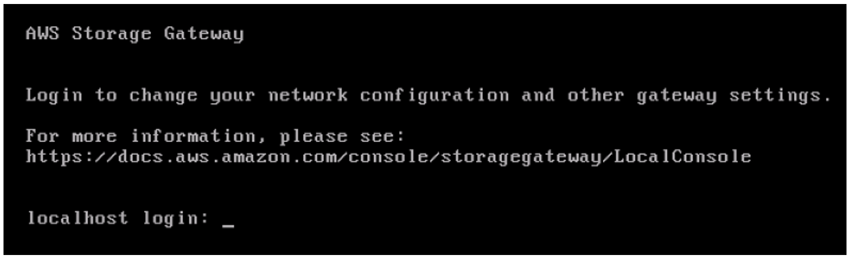

```
AWS Storage Gateway

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: _
```

To log in to the gateway's local console

- If this is your first time logging in to the local console, log in to the VM with the user name *sguser* and password *sgpassword*. Otherwise, use your credentials to log in.

After you log in, you see the **Storage Gateway Configuration** main menu, as shown in the following screenshot.

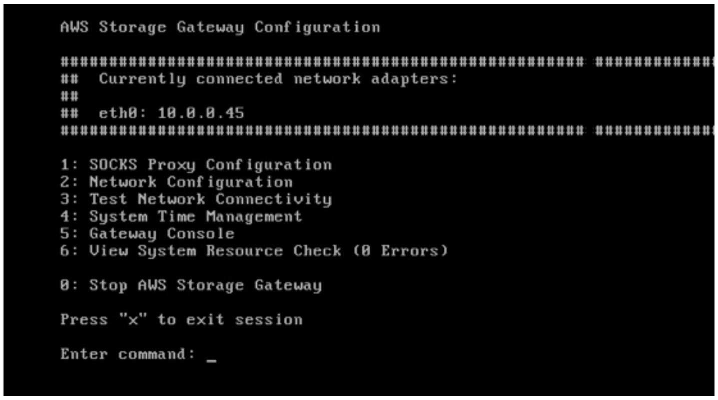

```
AWS Storage Gateway Configuration

###################################################### #############
##  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: _
```

Note

We recommend changing the default password. You do this by running the `passwd` command from the Gateway Console menu (item 5 on the main menu). For information about how to run the command, see Running Storage Gateway Commands on the Local Console (p. 168). 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 (p. 161).

To	See
Configure a SOCKS proxy for your gateway	Routing Your On-Premises Gateway Through a Proxy (p. 161).
Configure your network	Configuring Your Gateway Network (p. 162).
Test network connectivity	Testing Your Gateway Connection to the Internet (p. 165).
Manage VM time	Synchronizing Your Gateway VM Time (p. 167).

To	See
Run Storage Gateway console commands	Running Storage Gateway Commands on the Local Console (p. 168).
View system resource check	Viewing Your Gateway System Resource Status (p. 170).

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

To exit the configuration session, type **x** to exit the menu.

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 *sguser* and the password *sgpassword*. We recommend that you set a new password. 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 AWS 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. On the **Action** menu, 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. AWS 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

AWS Storage Gateway supports the configuration of a Socket Secure version 5 (SOCKS5) proxy between your on-premises gateway and AWS.

> **Note**
> The only proxy configuration AWS Storage Gateway supports 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, AWS Storage Gateway routes all HyperText Transfer Protocol Secure (HTTPS) traffic through your proxy server. For information about network requirements for your gateway, see Network and Firewall Requirements (p. 12).

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 AWS Storage Gateway Local Console (p. 157).
2. On the **AWS Storage Gateway Configuration** main menu, type **1** to begin configuring the SOCKS proxy.

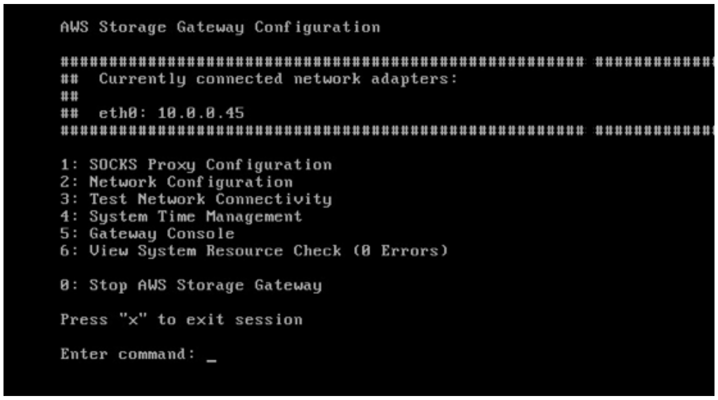

```
AWS Storage Gateway Configuration

####################################################### ############
##  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. Choose one of the following options on the **AWS Storage Gateway SOCKS Proxy Configuration** menu:

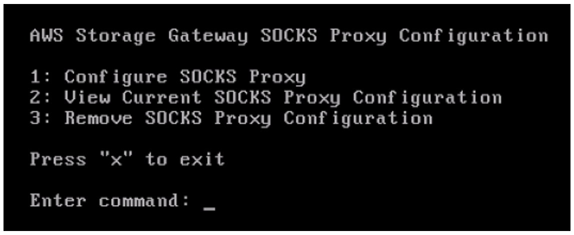

```
AWS Storage Gateway SOCKS Proxy Configuration

1: Configure SOCKS Proxy
2: View Current SOCKS Proxy Configuration
3: Remove SOCKS Proxy Configuration

Press "x" to exit

Enter command: _
```

To	Do This
Configure a SOCKS proxy	Type option **1**. You will need to supply a host name and port to complete configuration.
View the current SOCKS proxy configuration	Type option **2**. If a SOCKS proxy is not configured, the message `SOCKS Proxy not configured` is displayed. If a SOCKS proxy is configured, the host name and port of the proxy are displayed.
Remove a SOCKS proxy configuration	Type option **3**. The message `SOCKS Proxy Configuration Removed` is displayed.

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. For instructions, see Logging In to Your AWS Storage Gateway Local Console (p. 157).
2. On the **AWS Storage Gateway Configuration** main menu, type option **2** to begin configuring a static IP address.

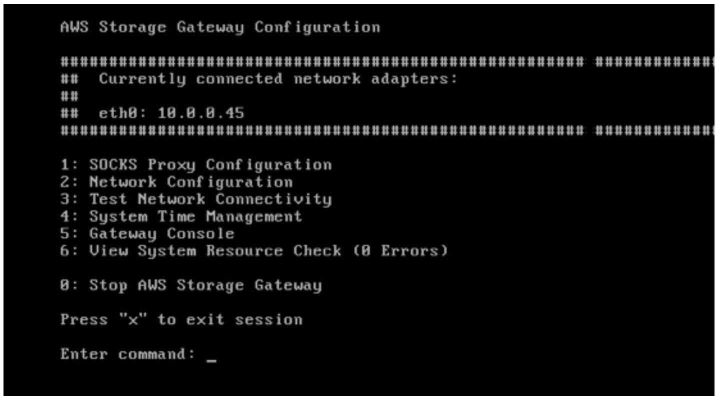

```
AWS Storage Gateway Configuration

####################################################### ############
##  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. Choose one of the following options on the **AWS Storage Gateway Network Configuration** menu:

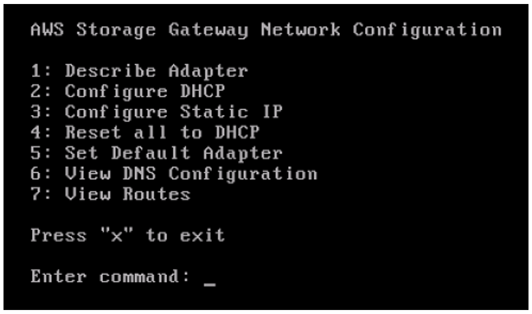

```
AWS Storage Gateway Network Configuration

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: _
```

To	Do This
Describe network adapter	Type option **1**. 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: • Media access control (MAC) address • IP address • Netmask • Gateway IP address • DHCP enabled status You use the same adapter name when you configure a static IP address (option **3**) as when you set your gateway's default route adapter (option **5**).
Configure DHCP	Type option **2**. You are prompted to configure network interface to use DHCP.

To	Do This

	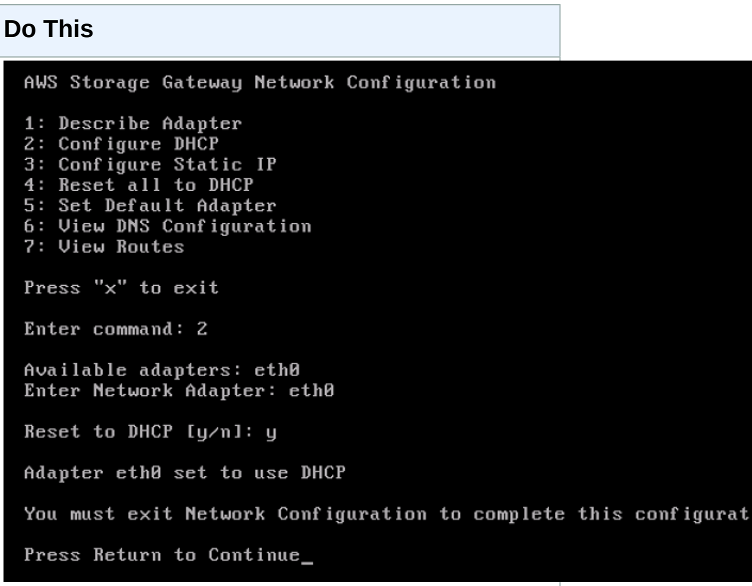
	`AWS Storage Gateway Network Configuration`
	`1: Describe Adapter`
	`2: Configure DHCP`
	`3: Configure Static IP`
	`4: Reset all to DHCP`
	`5: Set Default Adapter`
	`6: View DNS Configuration`
	`7: View Routes`
	`Press "x" to exit`
	`Enter command: 2`
	`Available adapters: eth0`
	`Enter Network Adapter: eth0`
	`Reset to DHCP [y/n]: y`
	`Adapter eth0 set to use DHCP`
	`You must exit Network Configuration to complete this configurat`
	`Press Return to Continue_`

Configure a static IP address for your gateway	Type option **3**. You are prompted to type the following information to configure a static IP: • Network adapter name • IP address • Netmask • Default gateway address • Primary Domain Name Service (DNS) address • Secondary DNS address **Important** If your gateway has already been activated, you must shut it down and restart it from the AWS Storage Gateway console for the settings to take effect. For more information, see Starting and Stopping Your Gateway (p. 141). If your gateway uses more than one network interface, you must set all enabled interfaces to use DHCP or static IP addresses. 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 disabled. To enable the interface in this case, you must set it to a static IP. If both interfaces are initially set to use static IP addresses and you then set the gateway to use DHCP, both interfaces will use DHCP.

To	Do This
Reset all your gateway's network configuration to DHCP	Type option **4**. 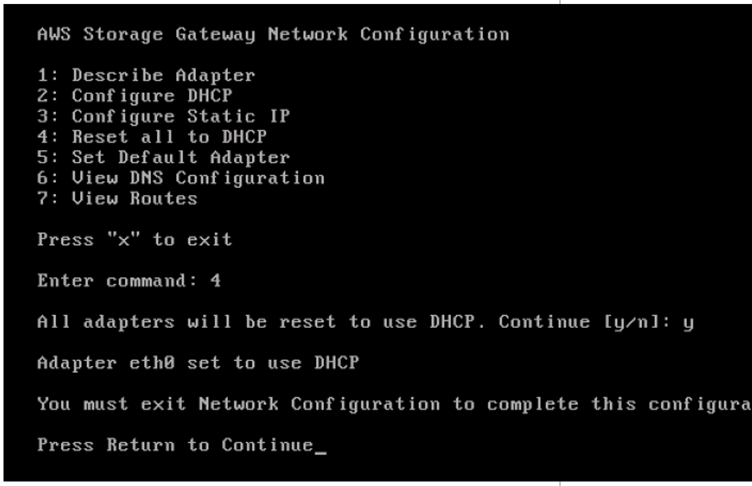 AWS Storage Gateway Network Configuration 1: Describe Adapter 2: Configure DHCP 3: Configure Static IP 4: Reset all to DHCP 5: Set Default Adapter 6: View DNS Configuration 7: View Routes Press "x" to exit Enter command: 4 All adapters will be reset to use DHCP. Continue [y/n]: y Adapter eth0 set to use DHCP You must exit Network Configuration to complete this configura Press Return to Continue_ All network interfaces are set to use DHCP. **Important** If your gateway has already been activated, you must shut down and restart your gateway from the AWS Storage Gateway console for the settings to take effect. For more information, see Starting and Stopping Your Gateway (p. 141).
Set your gateway's default route adapter	Type option **5**. The available adapters for your gateway are shown, and you are prompted to select one of the adapters—for example, **`eth0`**.
View your gateway's DNS configuration	Type option **6**. The IP addresses of the primary and secondary DNS name servers are displayed.
View routing tables	Type option **7**. The default route of your gateway is displayed.

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. For instructions, see Logging In to Your AWS Storage Gateway Local Console (p. 157).

2. On the **AWS Storage Gateway Configuration** main menu, type option **3** to begin testing network connectivity.

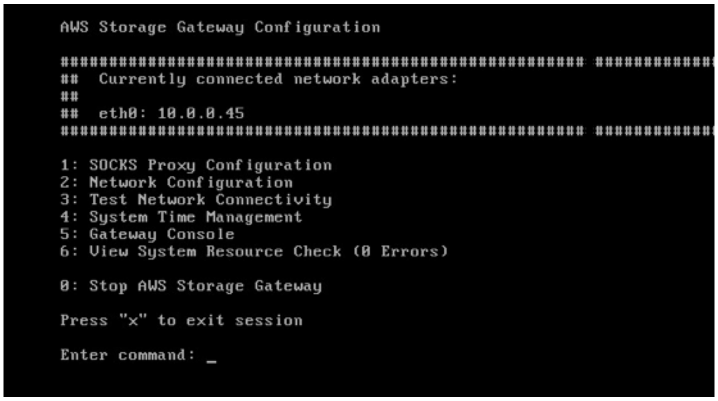

```
AWS Storage Gateway Configuration

###################################################### ############
##  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: _
```

The console displays the available regions.

3. Select the region you want to test. Following are the available regions for gateways deployed on-premises.

Note
Only cached volumes and tape gateway are available on Amazon EC2.

Region Name	Region String	File Gateway	Volume Gateway	Tape Gateway
US East (N. Virginia)	`us-east-1`	yes	yes	yes
US East (Ohio)	`us-east-2`	yes	yes	yes
US West (N. California)	`us-west-1`	yes	yes	yes
US West (Oregon)	`us-west-2`	yes	yes	yes
Canada (Central)	`ca-central-1`	yes	yes	yes
EU (Ireland)	`eu-west-1`	yes	yes	yes
EU (Frankfurt)	`eu-central-1`	yes	yes	yes
Asia Pacific (Tokyo)	`ap-northeast-1`	yes	yes	yes
Asia Pacific (Seoul)	`ap-northeast-2`	yes	yes	yes
Asia Pacific (Singapore)	`ap-southeast-1`	yes	yes	no
Asia Pacific (Sydney)	`ap-southeast-2`	yes	yes	yes
South America (São Paulo)	`sa-east-1`	yes	yes	no

Each endpoint in the selected region displays either a PASSED or FAILED message, as shown following.

Message	Description
[PASSED]	AWS Storage Gateway has Internet connectivity.
[FAILED]	AWS Storage Gateway does not have Internet connectivity.

For information about network and firewall requirements, see Network and Firewall Requirements (p. 12).

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 (p. 215). 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 Hyper-V gateway VM to an NTP server

1. Log in to your gateway's local console. For instructions, see Logging In to Your AWS Storage Gateway Local Console (p. 157).
2. On the **AWS Storage Gateway Configuration** main menu, type **4** for **System Time Management**.

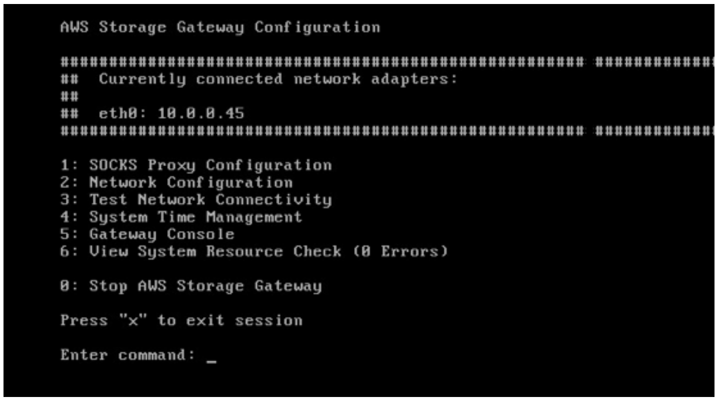

```
AWS Storage Gateway Configuration

###################################################### #############
##  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, type **1** for **View and Synchronize System Time**.

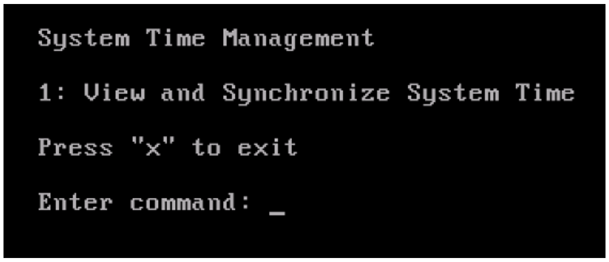

```
System Time Management

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 Network Time Protocol (NTP) time, type **y**. Otherwise, type **n**.

   If you type **y** to synchronize, the synchronization might take a few moments.

   The following screenshot shows a VM that does not require time synchronization.

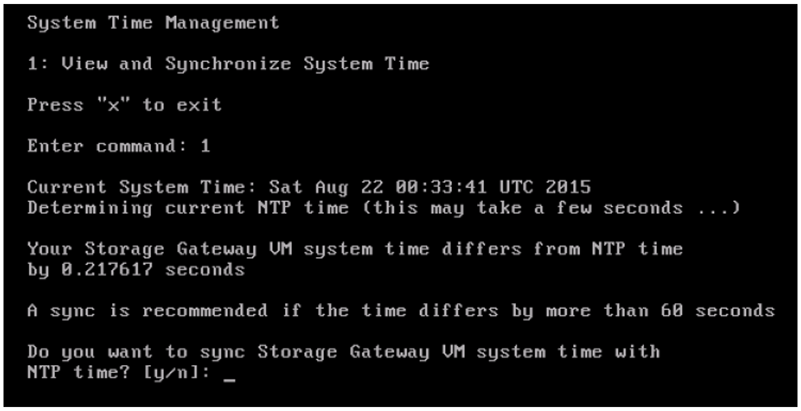

```
System Time Management

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.

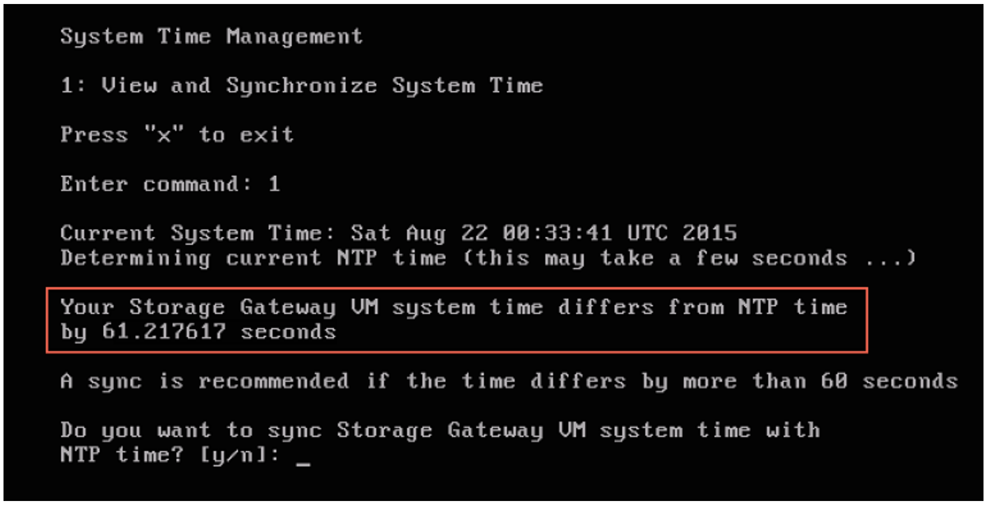

```
System Time Management

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]: _
```

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 AWS Storage Gateway Local Console (p. 157).
2. On the **AWS Storage Gateway Configuration** main menu, type option **5** for **Gateway Console**.

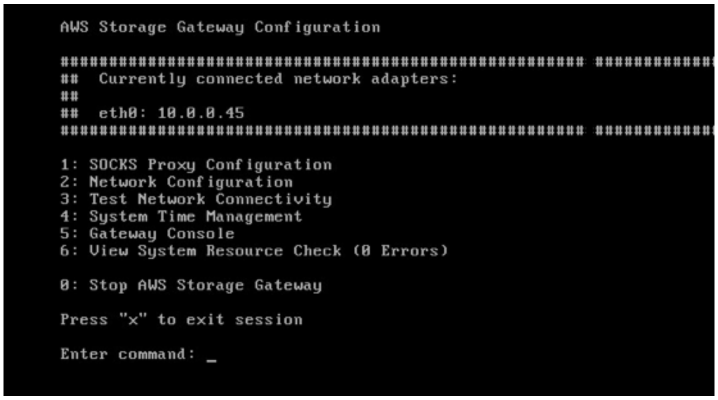

```
AWS Storage Gateway Configuration

####################################################### ############
##  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 AWS Storage Gateway console, type **h**, and then press the **Return** key.

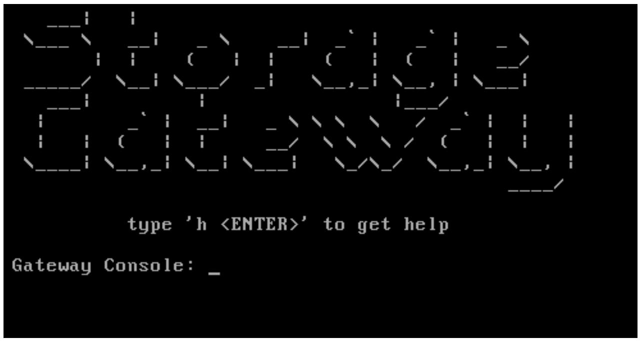

```
        type 'h <ENTER>' to get help

Gateway Console: _
```

The console displays the **Available Commands** menu with the available commands and after the menu a **Gateway Console** prompt, as shown in the following screenshot.

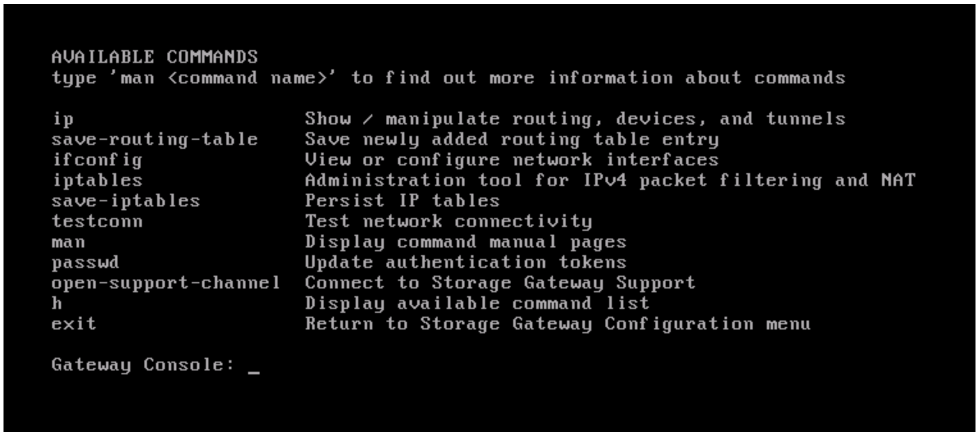

```
AVAILABLE COMMANDS
type 'man <command name>' to find out more information about commands

ip                      Show / manipulate routing, devices, and tunnels
save-routing-table      Save newly added routing table entry
ifconfig                View or configure network interfaces
iptables                Administration tool for IPv4 packet filtering and NAT
save-iptables           Persist IP tables
testconn                Test network connectivity
man                     Display command manual pages
passwd                  Update authentication tokens
open-support-channel    Connect to Storage Gateway Support
h                       Display available command list
exit                    Return to Storage Gateway Configuration menu

Gateway Console: _
```

4. To learn about a command, type `man` + *`command name`* at the **Gateway Console** prompt.

Viewing Your Gateway System Resource Status

When your gateway starts, it checks its virtual CPU cores, root volume size, and RAM and 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 AWS Storage Gateway Local Console (p. 157).
2. In the **AWS Storage Gateway Configuration** main menu, type **6** to view the results of a system resource check.

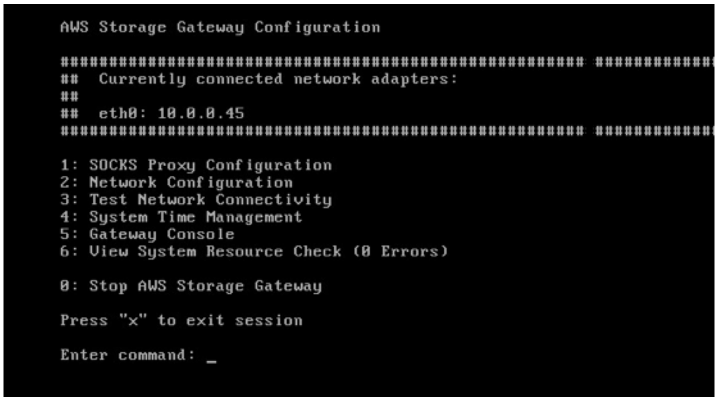

```
AWS Storage Gateway Configuration

###################################################### ############
##  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: _
```

The console displays an **[OK**], **[WARNING]**, or **[FAIL]** message for each resource as described in the table following.

Message	Description
[OK]	The resource has passed the system resource check.
[WARNING]	The resource does not meet the recommended requirements, but your gateway will continue to function. AWS Storage Gateway displays a message that describes the results of the resource check.
[FAIL]	The resource does not meet the minimum requirements. Your gateway might not function properly. AWS 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.

The following screenshot shows a **[FAIL]** message and the accompanying error message.

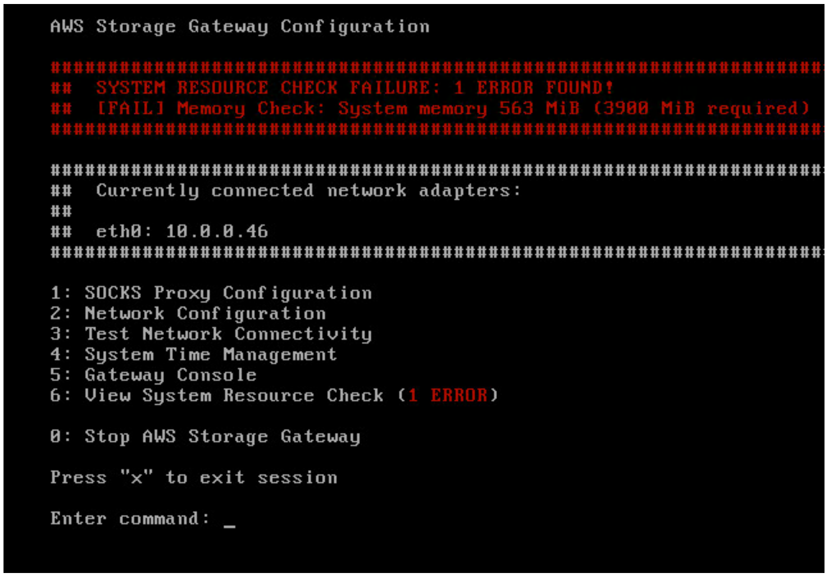

```
AWS Storage Gateway Configuration

###################################################################
##  SYSTEM RESOURCE CHECK FAILURE: 1 ERROR FOUND!
##  [FAIL] Memory Check: System memory 563 MiB (3900 MiB required)
###################################################################

###################################################################
##  Currently connected network adapters:
##
##  eth0: 10.0.0.46
###################################################################

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

0: Stop AWS Storage Gateway

Press "x" to exit session

Enter command: _
```

Configuring Network Adapters for Your Gateway

By default, AWS 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

AWS 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 this adapter, see the 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.

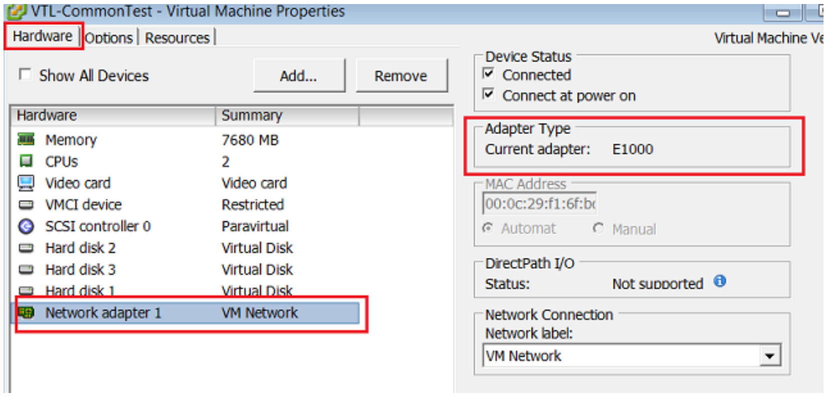

Hardware	Summary
Memory	7680 MB
CPUs	2
Video card	Video card
VMCI device	Restricted
SCSI controller 0	Paravirtual
Hard disk 2	Virtual Disk
Hard disk 3	Virtual Disk
Hard disk 1	Virtual Disk
Network adapter 1	VM Network

Device Status
☑ Connected
☑ Connect at power on

Adapter Type
Current adapter: E1000

MAC Address
00:0c:29:f1:6f:bc
● Automat ○ Manual

DirectPath I/O
Status: Not supported

Network Connection
Network label:
VM Network

4. Choose the E1000 network adapter, and then choose **Remove**. In this example, the E1000 network adapter is **Network adapter 1**.

   Note
   Although you can run the E1000 and VMXNET3 network adapters in your gateway at the same time, we don't recommend doing so because it can cause network problems.

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

After your gateway restarts, reconfigure the adapter you just added to make sure that network connectivity to the Internet is established.

To configure the adapter for the network

1. In the VSphere client, choose the **Console** tab to start the local console. You will 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 (p. 159).

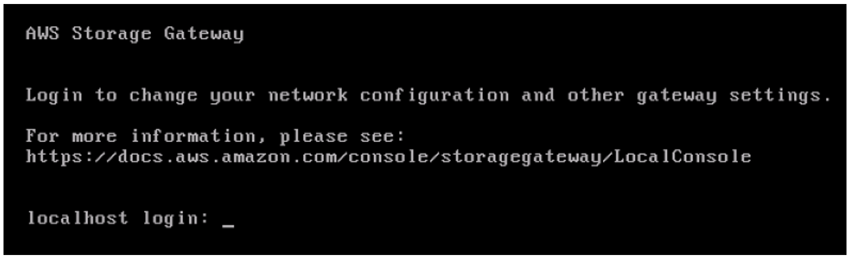

```
AWS Storage Gateway


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: _
```

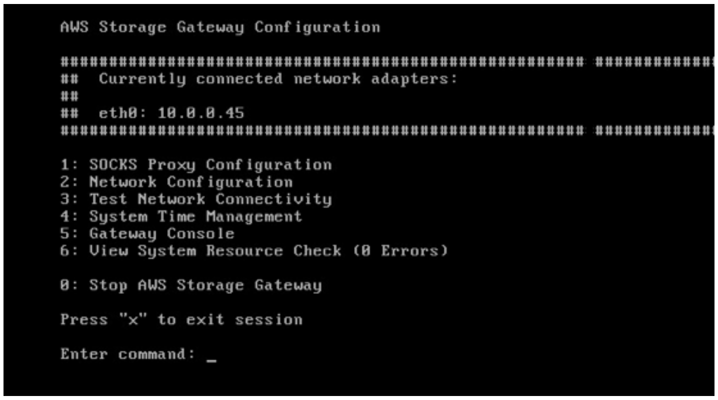

```
AWS Storage Gateway Configuration

###################################################### ############
##  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: _
```

2. At the prompt, type **2** to select **Network Configuration**, and then press **`Enter`** to open the network configuration menu.
3. At the prompt, type **4** to select **Reset to DHCP**, and then type **`y`** (for yes) at the prompt to reset the adapter you just added to use Dynamic Host Configuration Protocol (DHCP). You can type **5** to set all adapters to DHCP.
4. At the **Enter the adapter** prompt, type **`eth0`**, and then press **`Enter`** to continue. The only adapter available is **`eth0`**.

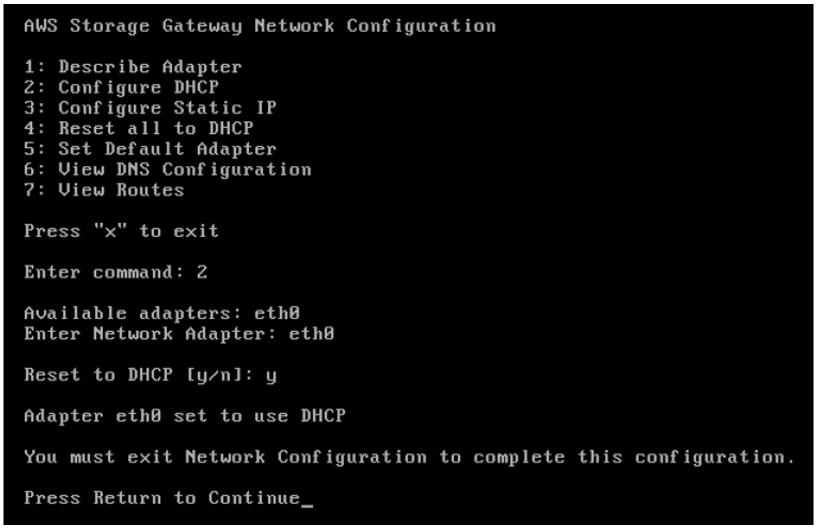

```
AWS Storage Gateway Network Configuration

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: 2

Available adapters: eth0
Enter Network Adapter: eth0

Reset to DHCP [y/n]: y

Adapter eth0 set to use DHCP

You must exit Network Configuration to complete this configuration.

Press Return to Continue_
```

If your gateway is already activated, you must shut it down and restart it from the AWS 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 (p. 165).

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.

The following procedure assumes that your gateway VM already has one network adapter defined and that you are adding a second adapter. The first procedure shows how to add an adapter for VMware ESXi, and the second shows how for Microsoft Hyper-V.

To configure your gateway to use an additional network adapter (VMware ESXi)

1. Shut down the gateway. For instructions, see To stop a gateway (p. 142).
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**.

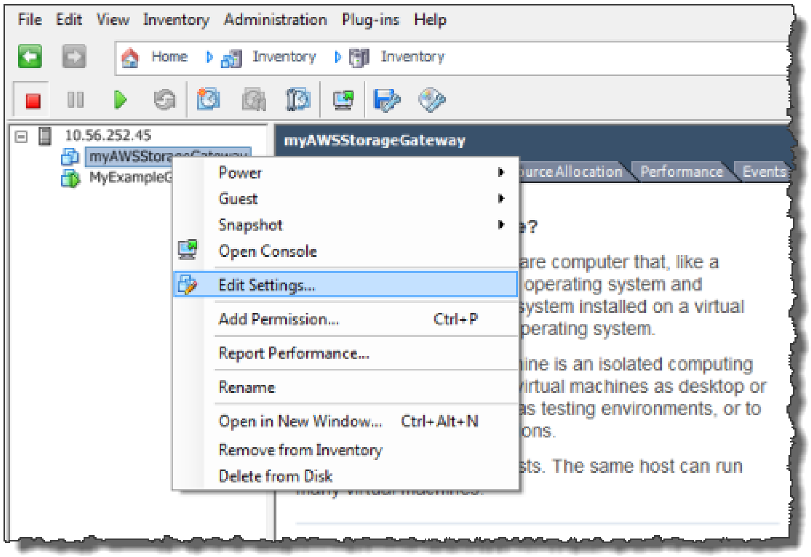

4. On the **Hardware** tab of the **Virtual Machine Properties** dialog box, choose **Add** to add a device.

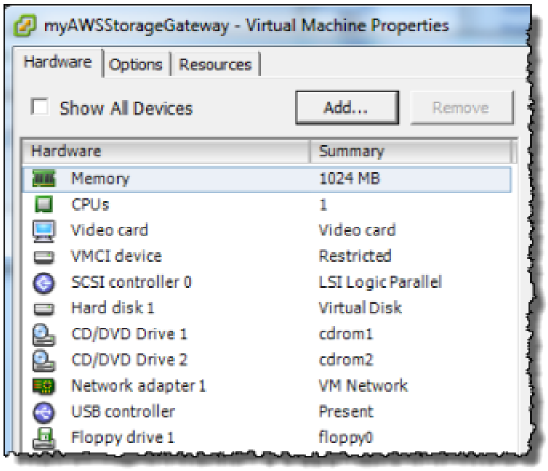

Hardware	Summary
Memory	1024 MB
CPUs	1
Video card	Video card
VMCI device	Restricted
SCSI controller 0	LSI Logic Parallel
Hard disk 1	Virtual Disk
CD/DVD Drive 1	cdrom1
CD/DVD Drive 2	cdrom2
Network adapter 1	VM Network
USB controller	Present
Floppy drive 1	floppy0

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**.

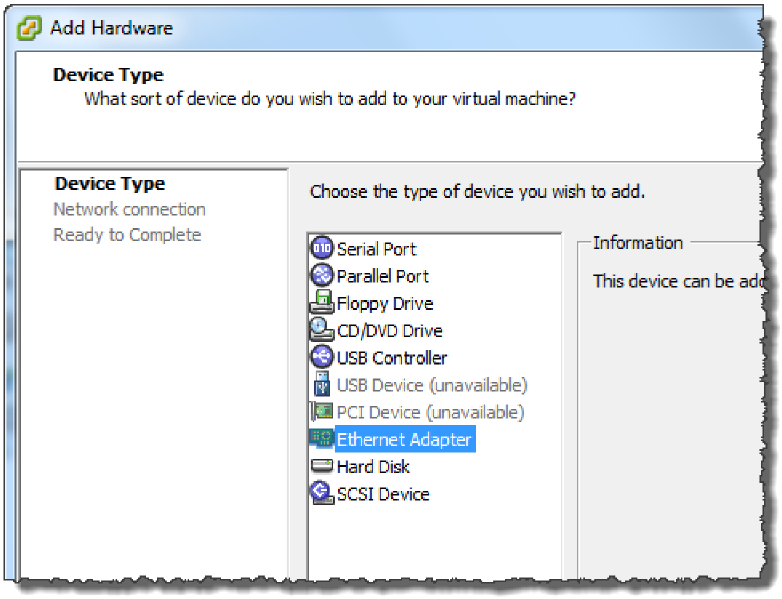

Device Type
What sort of device do you wish to add to your virtual machine?

Choose the type of device you wish to add.

This device can be add

   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 E1000 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.

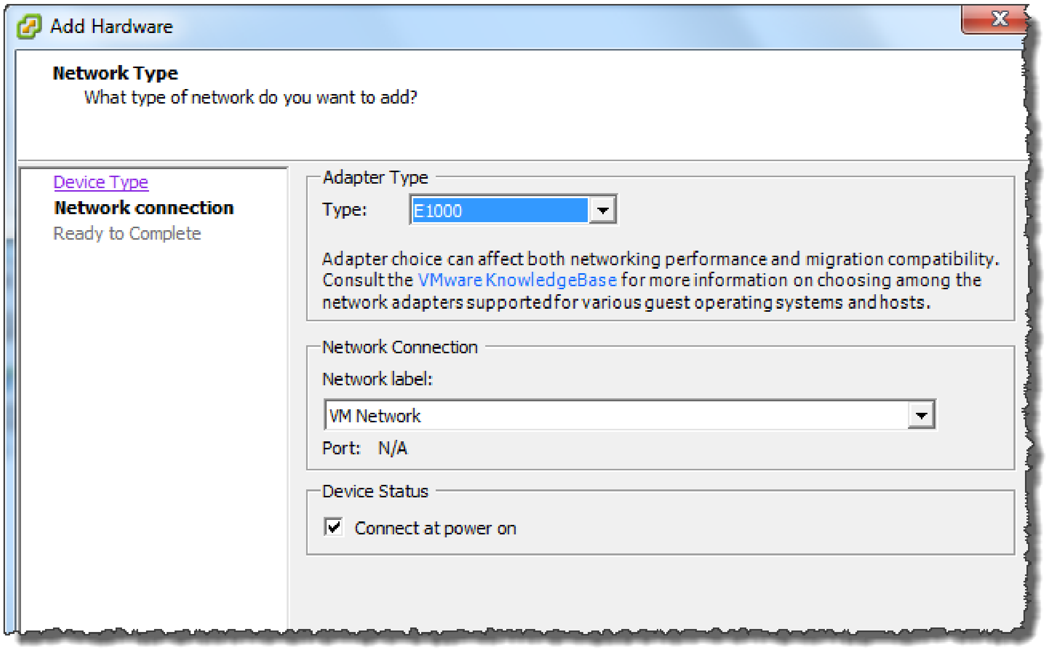

Add Hardware

Network Type

What type of network do you want to add?

Adapter Type

Type: E1000

Adapter choice can affect both networking performance and migration compatibility. Consult the VMware KnowledgeBase for more information on choosing among the network adapters supported for various guest operating systems and hosts.

Network Connection

Network label: VM Network

Port: N/A

Device Status

☑ Connect at power on

c. In the **Ready to Complete** pane, review the information, and then choose **Finish**.

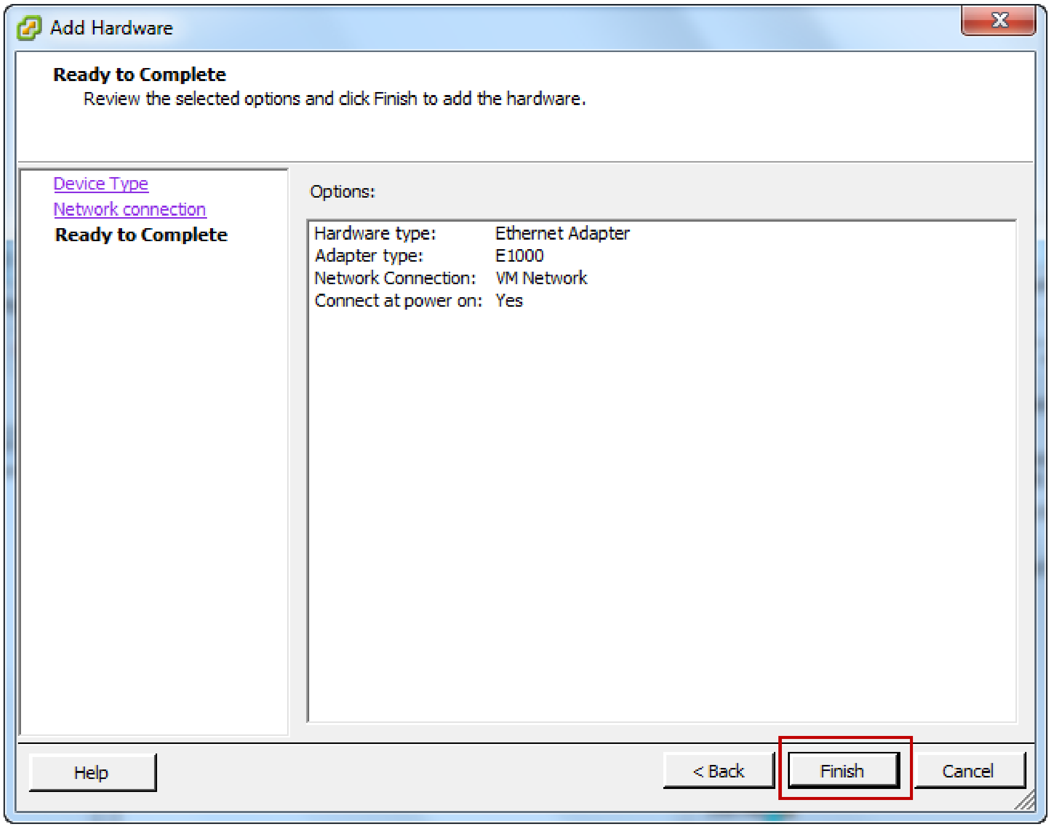

Add Hardware

Ready to Complete

Review the selected options and click Finish to add the hardware.

Options:

Hardware type: Ethernet Adapter
Adapter type: E1000
Network Connection: VM Network
Connect at power on: Yes

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

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

The following image is for illustration only. In practice, one of the IP addresses will be the address by which the gateway communicates to AWS and the other will be an address in a different subnet.

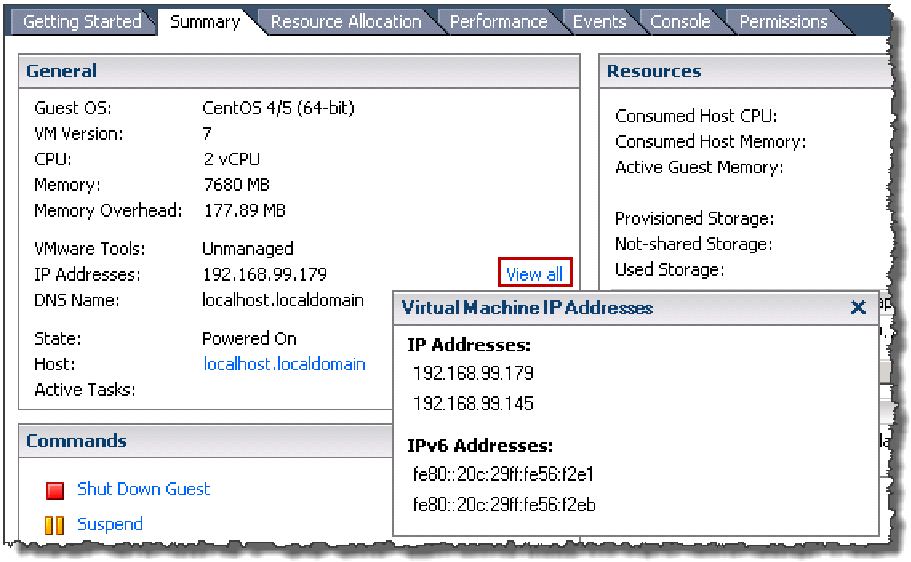

General

Guest OS: CentOS 4/5 (64-bit)
VM Version: 7
CPU: 2 vCPU
Memory: 7680 MB
Memory Overhead: 177.89 MB

VMware Tools: Unmanaged
IP Addresses: 192.168.99.179 View all
DNS Name: localhost.localdomain

State: Powered On
Host: localhost.localdomain
Active Tasks:

Resources

Consumed Host CPU:
Consumed Host Memory:
Active Guest Memory:

Provisioned Storage:
Not-shared Storage:
Used Storage:

Commands

Shut Down Guest
Suspend

Virtual Machine IP Addresses

IP Addresses:
192.168.99.179
192.168.99.145

IPv6 Addresses:
fe80::20c:29ff:fe56:f2e1
fe80::20c:29ff:fe56:f2eb

7. On the Storage Gateway console, turn on the gateway. For instructions, see To start a gateway (p. 142).
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.

To configure your gateway to use an additional network adapter (Microsoft Hyper-V)

1. On the Storage Gateway console, turn off the gateway. For instructions, see To stop a gateway (p. 142).
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**.

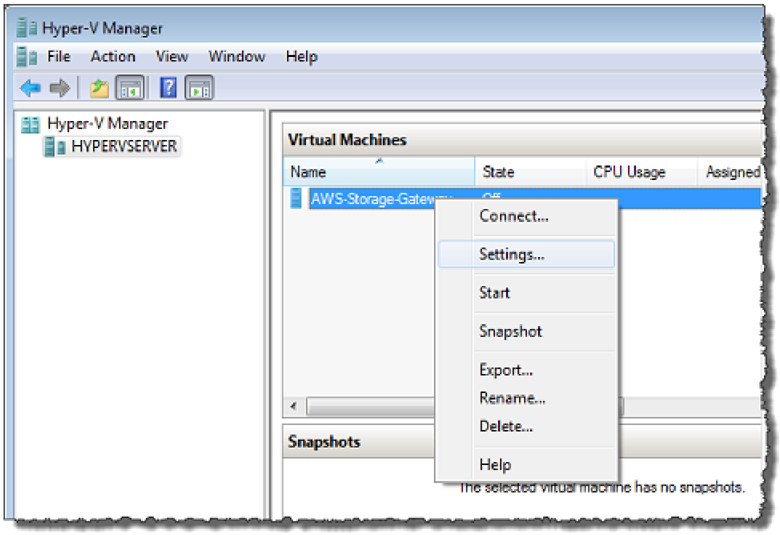

The selected virtual machine has no snapshots.

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.

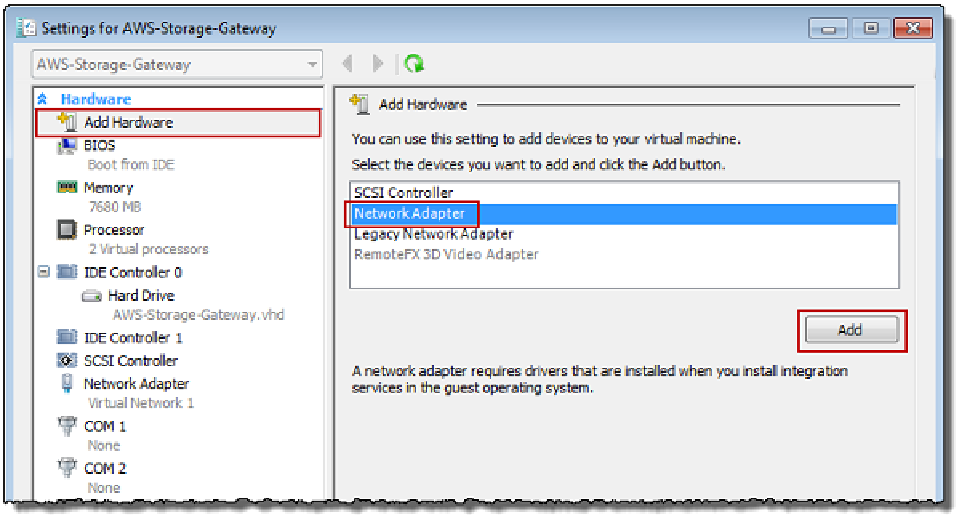

BIOS: Boot from IDE
Memory: 7680 MB
Processor: 2 Virtual processors
Hard Drive: AWS-Storage-Gateway.vhd
Network Adapter: Virtual Network 1
COM 1: None
COM 2: None

You can use this setting to add devices to your virtual machine.

Select the devices you want to add and click the Add button.

A network adapter requires drivers that are installed when you install integration services in the guest operating system.

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.

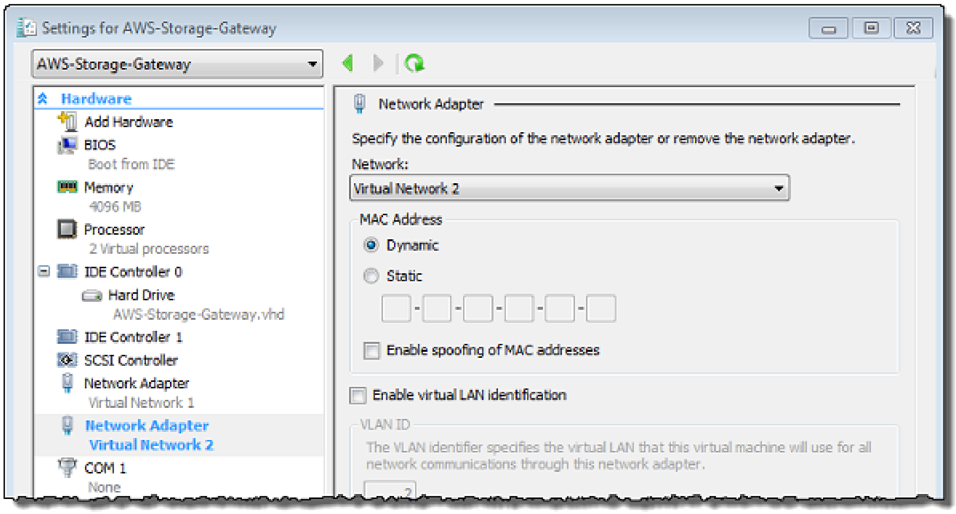

Settings for AWS-Storage-Gateway

AWS-Storage-Gateway

Hardware
- Add Hardware
- BIOS
Boot from IDE
- Memory
4096 MB
- Processor
2 Virtual processors
- IDE Controller 0
  - Hard Drive
AWS-Storage-Gateway.vhd
- IDE Controller 1
- SCSI Controller
- Network Adapter
Virtual Network 1
- Network Adapter
Virtual Network 2
- COM 1
None

Network Adapter

Specify the configuration of the network adapter or remove the network adapter.

Network: Virtual Network 2

MAC Address
- ● Dynamic
- ○ Static

□ Enable spoofing of MAC addresses

□ Enable virtual LAN identification

VLAN ID

The VLAN identifier specifies the virtual LAN that this virtual machine will use for all network communications through this network 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 gateway (p. 142).
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.

Creating a Volume on a Gateway with Multiple Network Adapters

If you have defined your gateway to use multiple network adapters, then when you create a volume for the gateway you must choose which IP address your storage applications will use to access the volume. Each network adapter defined for a gateway represents one IP address that you can choose. For information about how to add a network adapter to your gateway, see Configuring Your Gateway for Multiple NICs (p. 173).

To create a volume using a specified network adapter

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose the gateway you want to work with, and choose the **Volumes** tab.
3. Choose **Create New Volume**.
4. Configure the volume. If you have a gateway-cached volume, use the procedure described in Adding a Volume (p. 103).
5. Select an IP address to use to access the volume.

   Note that the **Create volume** dialog box displays a drop-down list for **Network interface**, with one IP address per adapter configured for the gateway VM. If the gateway VM is configured for only one network adapter, a drop-down list does not appear because there is only one IP address.

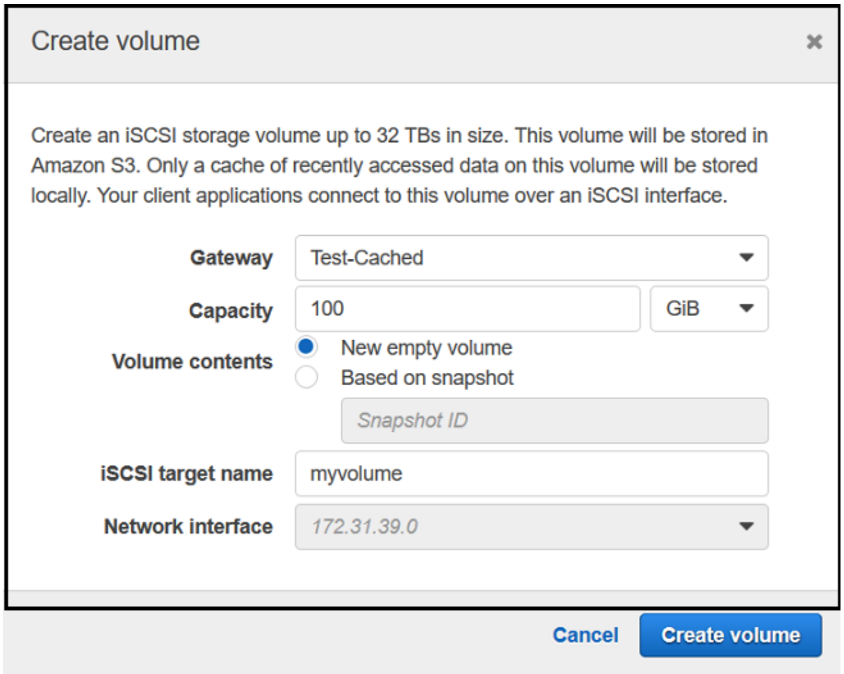

Create volume

Create an iSCSI storage volume up to 32 TBs in size. This volume will be stored in Amazon S3. Only a cache of recently accessed data on this volume will be stored locally. Your client applications connect to this volume over an iSCSI interface.

Gateway: Test-Cached

Capacity: 100 GiB

Volume contents: ● New empty volume ○ Based on snapshot

Snapshot ID

iSCSI target name: myvolume

Network interface: 172.31.39.0

6. Choose **Create volume**.

To create a connection to the storage volume, see Connecting to Volumes on Your Volume Gateway (p. 279).

Performing Maintenance Tasks on the Amazon EC2 Gateway Local Console

Some maintenance tasks require that you log in to the local console when running a gateway deployed on an Amazon EC2 instance. In this section, you can find information about how to log in to the local console and perform maintenance tasks.

Topics

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 from a Windows computer, log in as *sguser*.

2. After you log in, you see the **AWS Storage Gateway Configuration** main menu, as shown in the following screenshot.

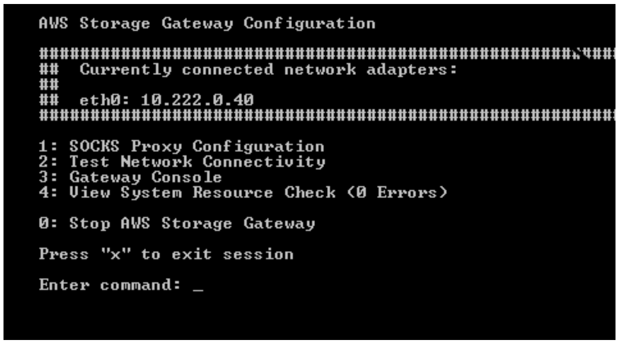

```
AWS Storage Gateway Configuration

########################################################
##  Currently connected network adapters:
##
##  eth0: 10.222.0.40
########################################################

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

0: Stop AWS Storage Gateway

Press "x" to exit session

Enter command: _
```

To	See
Configure a SOCKS proxy for your gateway	Routing Your Gateway Deployed on Amazon EC2 Through a Proxy (p. 181)
Test network connectivity	Testing Your Gateway Connectivity to the Internet (p. 182)
Run Storage Gateway console commands	Running Storage Gateway Commands on the Local Console (p. 184)
View a system resource check	Logging In to Your Amazon EC2 Gateway Local Console (p. 180).

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

To exit the configuration session, type **x** to exit the menu.

Routing Your Gateway Deployed on Amazon EC2 Through a Proxy

AWS Storage Gateway supports the configuration of a Socket Secure version 5 (SOCKS5) proxy between your gateway deployed on Amazon EC2 and AWS.

> **Note**
> The only proxy configuration AWS Storage Gateway supports 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, AWS Storage Gateway will route all HyperText Transfer Protocol Secure (HTTPS) traffic through your proxy server.

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 (p. 180).
2. On the **AWS Storage Gateway Configuration** main menu, type **1** to begin configuring the SOCKS proxy.

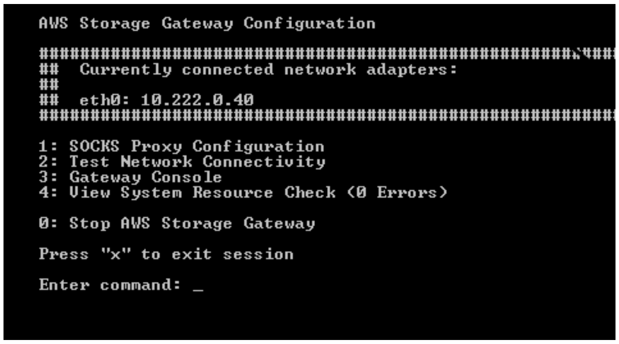

```
AWS Storage Gateway Configuration

#################################################################
##  Currently connected network adapters:
##
##  eth0: 10.222.0.40
#################################################################

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

0: Stop AWS Storage Gateway

Press "x" to exit session

Enter command: _
```

3. Choose one of the following options in the **AWS Storage Gateway SOCKS Proxy Configuration** menu:

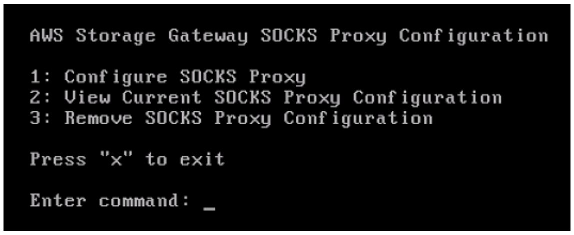

```
AWS Storage Gateway SOCKS Proxy Configuration

1: Configure SOCKS Proxy
2: View Current SOCKS Proxy Configuration
3: Remove SOCKS Proxy Configuration

Press "x" to exit

Enter command: _
```

To	Do This
Configure a SOCKS proxy	Type **1**. You need to supply a host name and port to complete configuration.
View the current SOCKS proxy configuration	Type **2**. If a SOCKS proxy is not configured, the message `SOCKS Proxy not configured` is displayed. If a SOCKS proxy is configured, the host name and port of the proxy are displayed.
Remove a SOCKS proxy configuration	Type **3**. The message `SOCKS Proxy Configuration Removed` is displayed.
Exit this menu and return to the previous menu	Type **x**.

Testing Your Gateway Connectivity 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. For instructions, see Logging In to Your Amazon EC2 Gateway Local Console (p. 180).
2. In the **AWS Storage Gateway Configuration** main menu, type **2** to begin testing network connectivity.

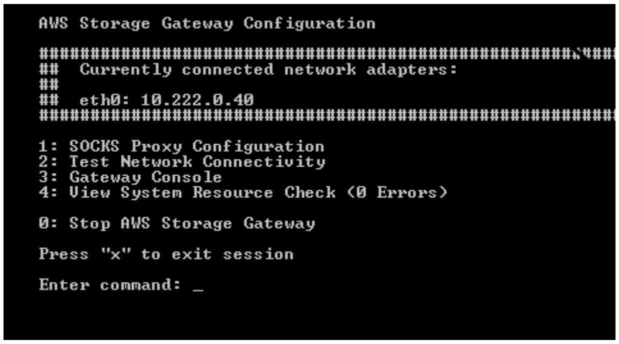

```
AWS Storage Gateway Configuration

#####################################################################
##  Currently connected network adapters:
##
##  eth0: 10.222.0.40
#####################################################################

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

0: Stop AWS Storage Gateway

Press "x" to exit session

Enter command: _
```

The console displays the available regions.

3. Select the region you want to test. Following are the available regions for gateways deployed an EC2 instance.

   Note
   Only cached volumes and tape gateway are available on Amazon EC2.

Region Name	Region String	File Gateway	Volume Gateway	Tape Gateway
US East (N. Virginia)	`us-east-1`	yes	yes	yes
US East (Ohio)	`us-east-2`	yes	yes	yes
US West (N. California)	`us-west-1`	yes	yes	yes
US West (Oregon)	`us-west-2`	yes	yes	yes
Canada (Central)	`ca-central-1`	yes	yes	yes
EU (Ireland)	`eu-west-1`	yes	yes	yes
EU (Frankfurt)	`eu-central-1`	yes	yes	yes
Asia Pacific (Tokyo)	`ap-northeast-1`	yes	yes	yes
Asia Pacific (Seoul)	`ap-northeast-2`	yes	yes	yes
Asia Pacific (Singapore)	`ap-southeast-1`	yes	yes	no
Asia Pacific (Sydney)	`ap-southeast-2`	yes	yes	yes
South America (São Paulo)	`sa-east-1`	yes	yes	no

Each endpoint in the region you select displays either a **[PASSED]** or **[FAILED]** message, as shown following.

Message	Description
[PASSED]	AWS Storage Gateway has Internet connectivity.
[FAILED]	AWS Storage Gateway does not have Internet connectivity.

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 (p. 180).
2. In the **AWS Storage Gateway Configuration** main menu, type **3** for **Gateway Console**.

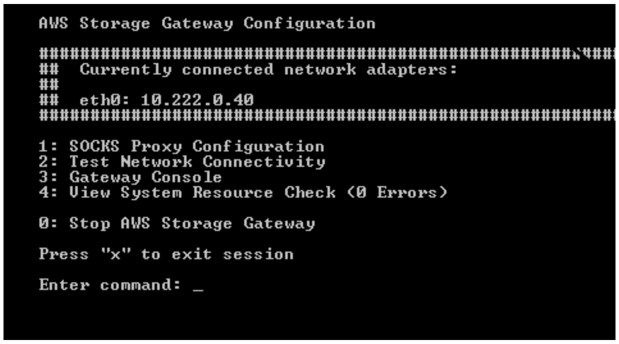

```
AWS Storage Gateway Configuration

##############################################################
##  Currently connected network adapters:
##
##  eth0: 10.222.0.40
##############################################################

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

0: Stop AWS Storage Gateway

Press "x" to exit session

Enter command: _
```

3. In the Storage Gateway console, type **h**, and then press the **Return** key.

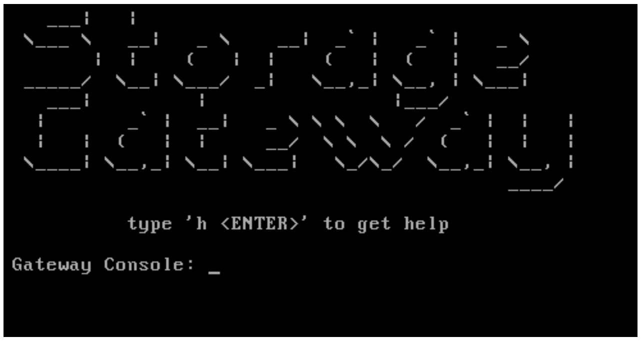

```
          type 'h <ENTER>' to get help

Gateway Console: _
```

The console displays the **Available Commands** menu with the available commands. After the menu, a **Gateway Console** prompt appears, as shown in the following screenshot.

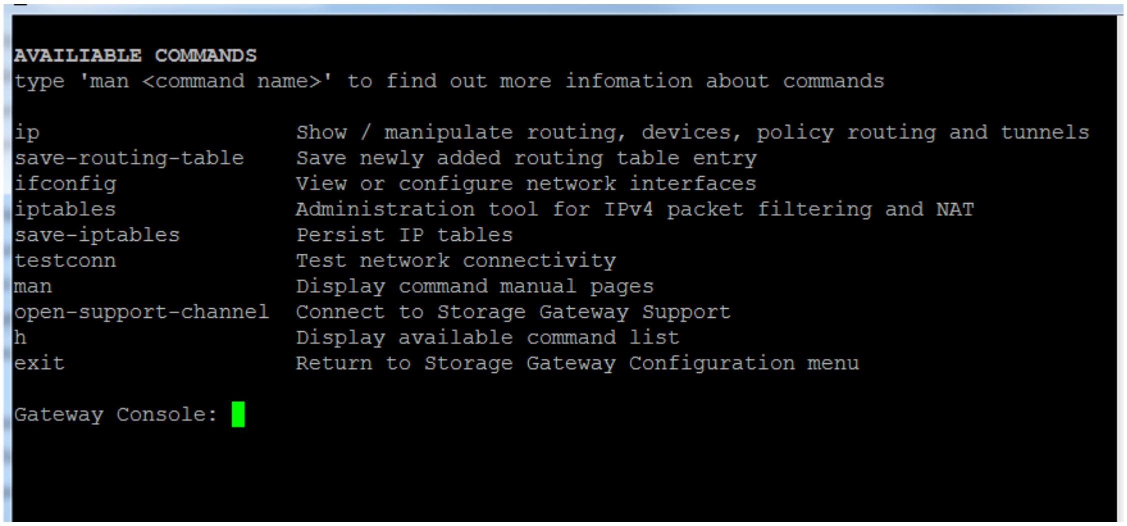

```
AVAILIABLE COMMANDS
type 'man <command name>' to find out more infomation about commands

ip                      Show / manipulate routing, devices, policy routing and tunnels
save-routing-table      Save newly added routing table entry
ifconfig                View or configure network interfaces
iptables                Administration tool for IPv4 packet filtering and NAT
save-iptables           Persist IP tables
testconn                Test network connectivity
man                     Display command manual pages
open-support-channel    Connect to Storage Gateway Support
h                       Display available command list
exit                    Return to Storage Gateway Configuration menu

Gateway Console:
```

4. To learn about a command, type **man** + *command name* at the **Gateway Console** prompt.

Viewing Your Gateway System Resource Status

When your gateway starts, it checks its virtual CPU cores, root volume size, and RAM and 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 (p. 180).
2. In the **AWS Storage Gateway Configuration** main menu, type **4** to view the results of a system resource check.

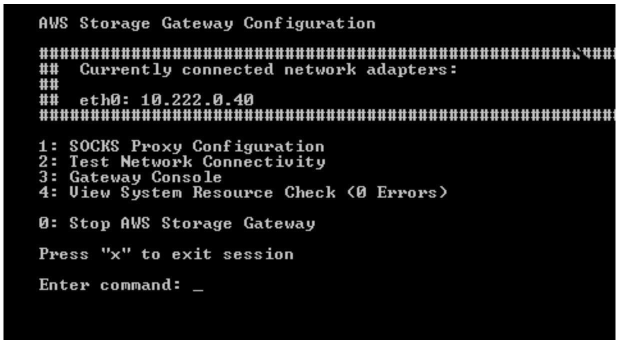

```
AWS Storage Gateway Configuration

##############################################################
##  Currently connected network adapters:
##
##  eth0: 10.222.0.40
##############################################################

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

0: Stop AWS Storage Gateway

Press "x" to exit session

Enter command: _
```

The console displays an **[OK**], **[WARNING]**, or **[FAIL]** message for each resource as described in the table following.

Message	Description
[OK]	The resource has passed the system resource check.
[WARNING]	The resource does not meet the recommended requirements, but your gateway will continue

Message	Description
	to function. AWS Storage Gateway displays a message that describes the results of the resource check.
[FAIL]	The resource does not meet the minimum requirements. Your gateway might not function properly. AWS 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.

The following screenshot shows a **[FAIL]** message and the accompanying error message.

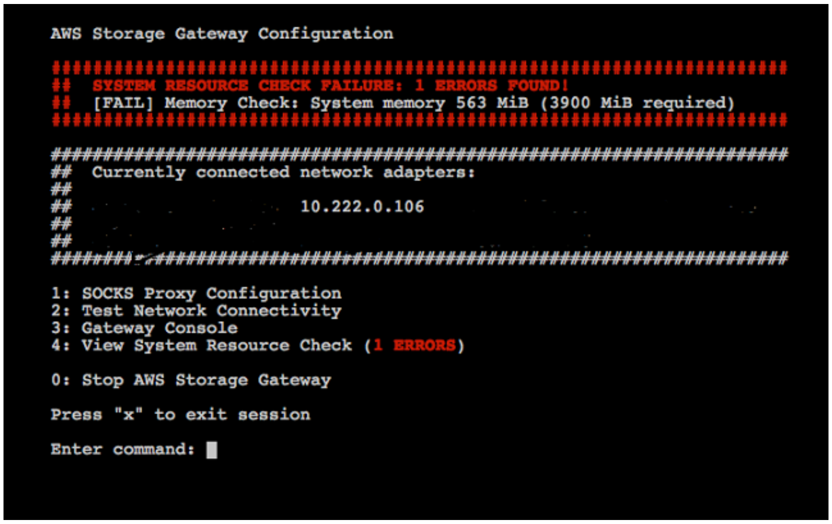

```
AWS Storage Gateway Configuration

########################################################################
##  SYSTEM RESOURCE CHECK FAILURE: 1 ERRORS FOUND!
##  [FAIL] Memory Check: System memory 563 MiB (3900 MiB required)
########################################################################

########################################################################
##  Currently connected network adapters:
##
##  [redacted]          10.222.0.106
##
##
########################################################################

1: SOCKS Proxy Configuration
2: Test Network Connectivity
3: Gateway Console
4: View System Resource Check (1 ERRORS)

0: Stop AWS Storage Gateway

Press "x" to exit session

Enter command: █
```

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.

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 AWS 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 a Amazon Elastic Compute Cloud (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 or Microsoft Hyper-V Manager 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 AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Gateways**, and then choose the gateway you want to delete.
3. On the **Action** menu, choose **Delete gateway**.
4. > **Important**
   > Before you do this step, be sure that there are no applications currently writing to the gateway's volumes. If you delete the gateway while it is in use, data loss can occur.

   > **Warning**
   > When a gateway is deleted, there is no way to get it back.

   In the confirmation dialog box that appears, select the check box to confirm your deletion. Make sure the gateway ID listed specifies the gateway you want to delete. and then choose **Delete**.

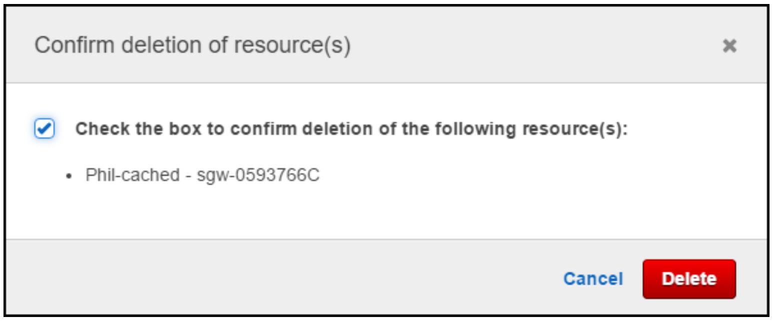

Confirm deletion of resource(s)

☑ Check the box to confirm deletion of the following resource(s):

- Phil-cached - sgw-0593766C

> **Important**
> You no longer pay software charges after you delete a gateway, but resources such as virtual tapes, Amazon Elastic Block Store (Amazon EBS) snapshots, and Amazon EC2 instances persist. You will continue to be billed for these resources. You can choose to remove Amazon EC2 instances and Amazon EBS snapshots by canceling your Amazon EC2 subscription.
> If you want to keep your Amazon EC2 subscription, you can delete your Amazon EBS snapshots using the Amazon EC2 console.

Removing Resources from a Gateway Deployed On-Premises

You can use the instructions following to remove resources from a gateway that is deployed on-premises.

Removing Resources from a Volume Gateway Deployed on a VM

If the gateway you want to delete are deployed on a virtual machine (VM), we suggest that you take the following actions to clean up resources:

- Delete the gateway. For instructions, see Deleting Your Gateway by Using the AWS Storage Gateway Console (p. 187).
- Delete all Amazon EBS snapshots you don't need. For instructions, see Deleting an Amazon EBS Snapshot in the *Amazon EC2 User Guide for Linux Instances*.

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 disable it and delete tapes that have the RETRIEVING or RETRIEVED status before you delete the tape gateway, as described following:

- For instructions on disabling a tape gateway, see Disabling Your Tape Gateway (p. 205).
- To delete tapes that have the RETRIEVING status from the library, cancel the retrieval process. For instructions, see Canceling Tape Retrieval (p. 277).
- To delete tapes that have the RETRIEVED status from the library, eject the tape using your backup software. For instructions, see Archiving the Tape (p. 90).

After disabling 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 AWS Storage Gateway Console (p. 187).

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 (p. 117).

> **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 Gateway-Cached Volumes Deployed on Amazon EC2

If you deployed a gateway with cached volumes on EC2, we suggest that you take the following actions to delete your gateway and clean up its resources:

1. In the Storage Gateway console, delete the gateway as shown in Deleting Your Gateway by Using the AWS Storage Gateway Console (p. 187).
2. In the Amazon EC2 console, stop your EC2 instance if you plan on using the instance again. Otherwise, terminate the instance. If you plan on deleting volumes, make note of the block devices that are attached to the instance and the devices' identifiers before terminating the instance. You will need these to identify the volumes you want to delete.
3. In the Amazon EC2 console, remove all Amazon EBS volumes that are attached to the instance if you don't plan on using them again. For more information, see Clean Up Your Instance and Volume in the *Amazon EC2 User Guide for Linux Instances*.

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 have RETRIEVING status from the tape gateway. For more information, see Canceling Tape Retrieval (p. 277).
2. Delete all virtual tapes that you have retrieved to your tape gateway. For more information, see Deleting Tapes (p. 117).
3. Delete all virtual tapes from the tape library. For more information, see Deleting Tapes (p. 117).
4. Delete the tape gateway. For more information, see Deleting Your Gateway by Using the AWS Storage Gateway Console (p. 187).
5. 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 for Linux Instances*.
6. Delete all archived virtual tapes. For more information, see Deleting Tapes (p. 117).

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.

Troubleshooting Your Gateway

Following, you can find information about troubleshooting issues related to gateways, volumes, virtual tapes, and snapshots. The gateway troubleshooting issues are split into two sections: gateways that are on-premises and gateways that are deployed on Amazon EC2. The on-premises gateway troubleshooting information covers gateways deployed on both the VMware ESXi and Microsoft Hyper-V clients. The troubleshooting information for volumes applies to volume gateway types. The troubleshooting information for tapes applies to tape gateways.

Topics

Troubleshooting On-Premises Gateway Issues

The following table lists typical issues that you might encounter working with your on-premises gateways.

Issue	Action to Take
You cannot find the IP address of your gateway.	Use the hypervisor client to connect to your host to find the gateway IP address. • For VMware ESXi, the VM's IP address can be found in the vSphere client on the **Summary** tab. For more information, see Activating Your Gateway (p. 55). • For Microsoft Hyper-V, the VM's IP address can be found by logging into the local console. For more information, see Activating Your Gateway (p. 55). If you are still having trouble finding the gateway IP address: • 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.

Issue	Action to Take
You're having network or firewall problems.	• Allow the appropriate ports for your gateway. • 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 (p. 12).
Your gateway's activation fails when you click the **Proceed to Activation** button in the AWS Storage Gateway console.	• 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 (p. 161). • 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 (p. 167). • After performing these steps, you can retry the gateway deployment using the AWS Storage Gateway console and the **Setup and Activate Gateway** wizard. • Check that your VM has at least 7.5 GB of RAM. Gateway allocation fails if there is less than 7.5 GB of RAM. For more information, see Requirements (p. 11).
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 Upload Buffer Capacity (p. 144)
You need to improve bandwidth between your gateway and AWS.	You can improve the bandwidth from your gateway to AWS by setting up your Internet connection to AWS on a 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 `CloudBytesDownloaded` and `CloudBytesUploaded` metrics of the gateway. For more on this subject, see Measuring Performance Between Your Gateway and AWS (p. 123). Improving your Internet connectivity helps to ensure that your upload buffer does not fill up.

Issue	Action to Take
Throughput to or from your gateway drops to zero.	• On the **Gateway** tab of the AWS 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 AWS Storage Gateway console, as shown in Starting and Stopping Your Gateway (p. 141). After the restart, the addresses in the **IP Addresses** list in the AWS Storage Gateway console's **Gateway** tab should match the IP addresses for your gateway, which you determine from the hypervisor client. • For VMware ESXi, the VM's IP address can be found in the vSphere client on the **Summary** tab. For more information, see Activating Your Gateway (p. 55). • For Microsoft Hyper-V, the VM's IP address can be found by logging into the local console. For more information, see Activating Your Gateway (p. 55). • Check your gateway's connectivity to AWS as described in Testing Your Gateway Connection to the Internet (p. 165). • Check your gateway's network adapter configuration, and ensure that all the interfaces you intended to be enabled for the gateway are enabled. To view the network adapter configuration for your gateway, follow the instructions in Configuring Your Gateway Network (p. 162) and select the option for viewing your gateway's network configuration. You can view the throughput to and from your gateway from the Amazon CloudWatch console. For more information about measuring throughput to and from your gateway to AWS, see Measuring Performance Between Your Gateway and AWS (p. 123).
You are having trouble importing (deploying) AWS Storage Gateway on Microsoft Hyper-V.	See Troubleshooting Your Microsoft Hyper-V Setup (p. 233), which discusses some of the common issues of deploying a gateway on Microsoft Hyper-V.

Enabling AWS Support Access to Your Gateway Hosted On-Premises

AWS Storage Gateway provides a local console you can use to perform several maintenance tasks, including enabling AWS Support to access your gateway to assist you with troubleshooting gateway issues. By default, AWS Support’s access to your gateway is disabled. You enable 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 enable AWS Support access to your gateway

1. Log in to your host's local console. For instructions, see Logging Into Your Gateway Local Console.

   The local console looks like the following.

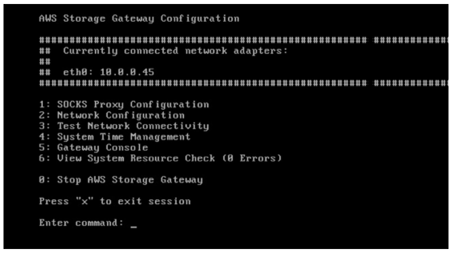

```
AWS Storage Gateway Configuration

############################################################ ############
##  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: _
```

2. At the prompt, type **5** to open the AWS Storage Gateway console.

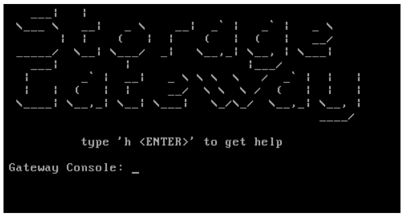

```
        type 'h <ENTER>' to get help

Gateway Console: _
```

3. Type **h** to open the **AVAILABLE COMMANDS** window.
4. In the **AVAILABLE COMMANDS** window, type **`open-support-channel`** to connect to customer support for AWS Storage Gateway. You must allow TCP port 22 to initiate 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.

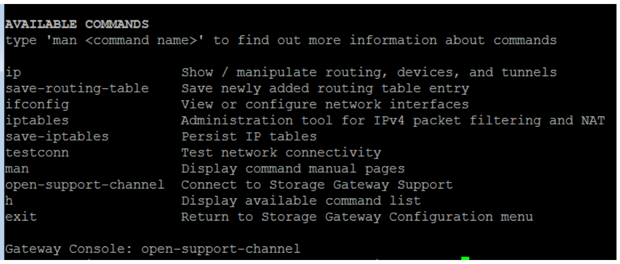

```
AVAILABLE COMMANDS
type 'man <command name>' to find out more information about commands

ip                      Show / manipulate routing, devices, and tunnels
save-routing-table      Save newly added routing table entry
ifconfig                View or configure network interfaces
iptables                Administration tool for IPv4 packet filtering and NAT
save-iptables           Persist IP tables
testconn                Test network connectivity
man                     Display command manual pages
open-support-channel    Connect to Storage Gateway Support
h                       Display available command list
exit                    Return to Storage Gateway Configuration menu

Gateway Console: open-support-channel
```

> **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. Once 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, type **`q`** to end it.
7. Type **`exit`** to log out of the AWS Storage Gateway console.
8. Follow the prompts to exit the local console.

Troubleshooting Amazon EC2 Gateway Issues

The following table lists typical issues that you might encounter working with your gateway deployed on Amazon Elastic Compute Cloud (Amazon EC2). For more information about the difference between an on-premises gateway and a gateway deployed in Amazon EC2, see Provisioning an Amazon EC2 Host (p. 53).

For more information about enabling AWS Support to access a gateway hosted on an EC2 instance, see the following section Enabling AWS Support to Access a Gateway Hosted on an Amazon EC2 Instance (p. 196).

Issue	Action to Take
Your Amazon EC2 gateway activation fails when you click the **Proceed to Activation** button in the AWS Storage Gateway console.	If activation has not occurred in a few moments, check the following in the Amazon EC2 console: - Port 80 is enabled in the security group 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 for Linux Instances*. - The gateway instance is marked as running. In the Amazon EC2 console, the **State** of the instance should be RUNNING. After correcting the problem, try activating the gateway again by going to the AWS Storage Gateway console, clicking **Deploy a new Gateway on Amazon EC2**, and re-entering the IP address of the instance.
You can't find your Amazon EC2 gateway instance in the list of instances.	If you did not give your instance a resource tag and you have many instances running, it can be hard to tell which instance you deployed the gateway in. In this case, you can take the following actions to find the gateway instance: - Check the name of the Amazon Machine Image (AMI) name on the **Description** tab of the instance. An instance based on the AWS Storage Gateway AMI should start with the text **`aws-storage-gateway-ami`**. - If you have several instances based on the AWS 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 Amazon EC2 gateway instance.	Check that the Amazon EBS volume in question is in the same Availability Zone as the gateway instance. If there is a discrepancy in Availability Zones, create a new Amazon EBS volume in the same Availability Zone as your instance.
You can't attach an initiator to a volume target of your Amazon EC2 gateway.	Check that the security group you launched the instance with includes a rule allowing the port that you are using for iSCSI access. The port is usually set as 3260. For more information on connecting to volumes, see Connecting to Volumes on Your Volume Gateway (p. 279).
You activated your Amazon EC2 gateway, but when you try to add storage volumes, you receive an error message indicating you have no disks available.	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 Provisioning an Amazon EC2 Host (p. 53). After attaching two or more Amazon EBS volumes, try creating volume storage on the gateway.

Issue	Action to Take
You need to remove a disk allocated as upload buffer space because you want to reduce the amount of upload buffer space.	Follow the steps in Adding and Removing Upload Buffer (p. 144).
Throughput to or from your Amazon EC2 gateway drops to zero.	• Verify the gateway instance is running. If the instance is starting due to a reboot, for example, wait for the instance to restart. • Verify that the gateway IP has not changed. If the instance was stopped and then restarted, the IP address of the instance 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 to AWS, see Measuring Performance Between Your Gateway and AWS (p. 123).

> **Warning**
> The elastic IP address of the Amazon EC2 instance cannot be used as the target address.

Enabling AWS Support to Access a Gateway Hosted on an Amazon EC2 Instance

AWS Storage Gateway provides a local console you can use to perform several maintenance tasks, including enabling AWS Support to access your gateway to assist you with troubleshooting gateway issues. By default, AWS Support’s access to your gateway is disabled. You enable this access through the Amazon EC2 local console. You log in to the Amazon EC2 local console via Secure Shell (SSH). To successfully log in via SSH, your instance's security group must have a rule allowing access via 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 enable 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 Instancein 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 sguser@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.

The local console looks like the following.

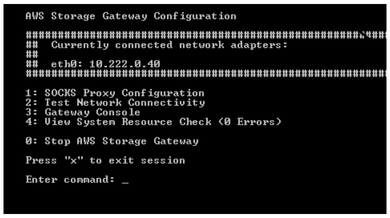

```
AWS Storage Gateway Configuration

##############################################################
##  Currently connected network adapters:
##
##  eth0: 10.222.0.40
##############################################################

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

0: Stop AWS Storage Gateway

Press "x" to exit session

Enter command: _
```

2. At the prompt, type **3** to open the AWS Storage Gateway console.

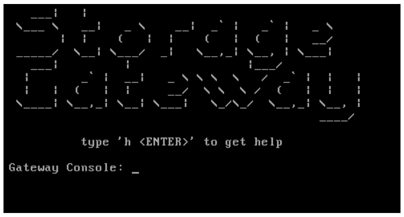

```
          type 'h <ENTER>' to get help

Gateway Console: _
```

3. Type **h** to open the **AVAILABLE COMMANDS** window.
4. In the **AVAILABLE COMMANDS** window, type **`open-support-channel`** to connect to customer support for AWS Storage Gateway. You must allow TCP port 22 to initiate 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.

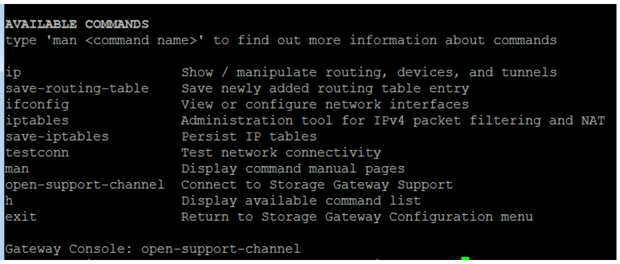

```
AVAILABLE COMMANDS
type 'man <command name>' to find out more information about commands

ip                     Show / manipulate routing, devices, and tunnels
save-routing-table     Save newly added routing table entry
ifconfig               View or configure network interfaces
iptables               Administration tool for IPv4 packet filtering and NAT
save-iptables          Persist IP tables
testconn               Test network connectivity
man                    Display command manual pages
open-support-channel   Connect to Storage Gateway Support
h                      Display available command list
exit                   Return to Storage Gateway Configuration menu

Gateway Console: open-support-channel
```

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. Once 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, type **`q`** to end it.
7. Type **`exit`** to exit the AWS Storage Gateway console.
8. Follow the console menus to log out of the AWS Storage Gateway instance.

Troubleshooting Volume, Snapshot, and Tape Issues

You can find information following about issues that you might face with your volumes, snapshots, and tapes, and how to troubleshoot these issues.

Topics

Troubleshooting Volume Issues

You can find information following about the most typical issues you might encounter when working with volumes, and actions that we suggest that you take to fix them.

Topics

The Console Says That Your Volume Is Not Configured

If the AWS Storage Gateway console indicates that your volume has a status of UPLOAD BUFFER NOT CONFIGURED (p.), add upload buffer capacity to your gateway. You cannot use a gateway to store your application data if the upload buffer for the gateway is not configured. For more information, see To configure upload buffer or cache storage (p. 143).

The Console Says That Your Volume Is Irrecoverable

If the AWS Storage Gateway console indicates that your volume has a status of IRRECOVERABLE (p.), you can no longer use this volume. You can try to delete the volume in the AWS Storage Gateway console. If there is data on the volume, then you can recover the data when you create a new volume based on the local disk of the VM that was initially used to create the volume. When you create the new volume, select **Preserve existing data**. For more information on this process, see Managing Volumes (p. 266). Make sure to delete pending snapshots of the volume before deleting the volume. For more information, see Deleting a Snapshot (p. 105).

If deleting the volume in the AWS Storage Gateway console does not work, then the disk allocated for the volume might have been improperly removed from the VM and cannot be removed from the appliance.

The Console Says That Your Volume Has PASS THROUGH Status

In some cases, the AWS Storage Gateway console might indicate that your volume has a status of PASS THROUGH (p.). A volume can have PASS THROUGH (p.) status for several reasons. Some reasons require action, and some do not.

An example of when you should take action if your volume has the PASS THROUGH status is when your gateway has run out of upload buffer space. To verify if your upload buffer was exceeded in the past, you can view the `UploadBufferPercentUsed` metric in the Amazon CloudWatch console; for more information, see Monitoring the Upload Buffer (p. 129). If your gateway has the PASS THROUGH status because it has run out of upload buffer space, you should allocate more upload buffer space to your gateway. Adding more buffer space will cause your volume to transition from PASS THROUGH to BOOTSTRAPPING (p.) to AVAILABLE (p.) automatically. While the volume has the BOOTSTRAPPING status, the gateway reads data off the volume's disk, uploads this data to Amazon S3, and catches up as needed. When the gateway has caught up and saved the volume data to Amazon S3, the volume status becomes AVAILABLE and snapshots can be started again. Note that when your volume has the PASS THROUGH or BOOTSTRAPPING status, you can continue to read and write data from the volume disk. For more information about adding more upload buffer space, see Adding and Removing Upload Buffer (p. 144).

To take action before the upload buffer is exceeded, you can set a threshold alarm on a gateway's upload buffer. For more information, see To set an upper threshold alarm for a gateway's upload buffer (p. 130).

In contrast, an example of not needing to take action when a volume has the PASS THROUGH status is when the volume is waiting to be bootstrapped because another volume is currently being bootstrapped. The gateway bootstraps volumes one at a time.

Infrequently, the PASS THROUGH status can indicate that a disk allocated for an upload buffer has failed. In this is the case, you should remove the disk. For more information, see Removing Upload Buffer Capacity (p. 144)

You Want to Verify Volume Integrity and Fix Possible Errors

If you want to verify volume integrity and fix possible errors, and your gateway uses Microsoft Windows initiators to connect to its volumes, you can use the Windows CHKDSK utility to verify the integrity of your volumes and fix any errors on the volumes. Windows can automatically run the CHKDSK tool when volume corruption is detected, or you can run it yourself.

Your Volume's iSCSI Target Doesn’t Appear in Windows Disk Management Console

If your volume's iSCSI target does not show up in the Disk Management Console in Windows, check that you have configured the upload buffer for the gateway. For more information, see To configure upload buffer or cache storage (p. 143).

You Want to Change Your Volume's iSCSI Target Name

If you want to change the iSCSI target name of your volume, you must delete the volume and add it again with a new target name. If you do so, you can preserve the data on the volume.

Your Scheduled Volume Snapshot Did Not Occur

If your scheduled snapshot of a volume did not occur, check whether your volume has the PASS THROUGH (p.) status, or if the gateway's upload buffer was filled just prior to the scheduled snapshot time. You can check the `UploadBufferPercentUsed` metric for the gateway in the Amazon CloudWatch console and create an alarm for this metric. For more information, see Monitoring the Upload Buffer (p. 129) and To set an upper threshold alarm for a gateway's upload buffer (p. 130).

You Need to Remove or Replace a Volume

If you need to remove a volume because it isn't needed, or you need to replace a volume disk that has failed, you should remove the volume first using the AWS Storage Gateway console. For more information, see To remove a volume (p. 103). You then use the hypervisor client to remove the backing storage:

- For VMware ESXi, remove the backing storage as described in Deleting a Volume (p. 103).
- For Microsoft Hyper-V, remove the backing storage as described in To remove the underlying local disk (Microsoft Hyper-V) (p. 270).

Throughput from Your Application to a Volume Has Dropped to Zero

If throughput from your application to a volume has dropped to zero, try the following:

- If you are using the VMware vSphere client, check that your volume's **Host IP** address matches one of the addresses that appears in the vSphere client on the **Summary** tab. You can find the **Host IP** address for a storage volume in the AWS Storage Gateway console in the **ISCSI Target Info** tab for the volume. A discrepancy in the IP address can occur, for example, when you assign a new static IP address to your gateway. If there is a discrepancy, restart your gateway from the AWS Storage Gateway console as shown in Starting and Stopping Your Gateway (p. 141). After the restart, the **Host IP** address in the **ISCSI Target Info** tab for a storage volume should match an IP address shown in the vSphere client on the **Summary** tab for the gateway.
- Check to see if **IPAddressNotFound** appears in the **Host IP** box for the volume. For example, this message can appear when you create a volume associated with an IP address of a network adapter of a gateway with two or more network adapters. When you remove or disable the network adapter that the volume is associated with, the **IPAddressNotFound** message is displayed. To address this issue, delete the volume and then re-create it preserving its existing data. For more information, see Managing Volumes (p. 266).
- Check that the iSCSI initiator your application uses is correctly mapped to the iSCSI target for the storage volume. For more information about connecting to storage volumes, see Connecting to Volumes on Your Volume Gateway (p. 279).

You can view the throughput for volumes and create alarms from the Amazon CloudWatch console. For more information about measuring throughput from your application to a volume, see Measuring Performance Between Your Application and Gateway (p. 122).

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.

Troubleshooting Snapshot Issues

You can find information following about issues you might encounter when working with snapshots, actions that we suggest that you take action to fix these issues, and using recovery snapshots for your gateway-cached setup.

Topics

A Volume Snapshot Has PENDING Status Longer Than Expected

If a volume snapshot remains in PENDING state longer than expected, the gateway VM might have crashed unexpectedly or the status of a volume might have changed to PASS THROUGH (p.) or IRRECOVERABLE (p.). If any of these are the case, the snapshot remains in PENDING status and the snapshot does not successfully complete.

If any of these situations is the case, we recommend that you delete the snapshot. For more information, see Deleting a Snapshot (p. 105). When the volume returns to AVAILABLE (p.) status, create a new snapshot of the volume.

Using Recovery Snapshots for Your Gateway-Cached Setup

AWS Storage Gateway provides recovery points for each volume in a gateway-cached volume architecture. A *volume recovery point* is a point in time at which all data of the volume is consistent and from which you can create a snapshot. You can use the snapshot to create a new volume in the event that your gateway becomes unreachable or one gateway-cached volume becomes irrecoverable.

Volume recovery points are maintained automatically for each gateway-cached volume. You can also take snapshots on a one-time, ad hoc basis or set up a snapshot schedule for the volume. For more information about snapshots, see Working with Snapshots (p. 253).

When the gateway becomes unreachable (such as when you shut it down), you have the option of creating a snapshot from a volume recovery point and using the snapshot.

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

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Gateways**.
3. Choose the unreachable gateway, and then choose the **Details** tab.

   A recovery snapshot message is displayed in the tab.

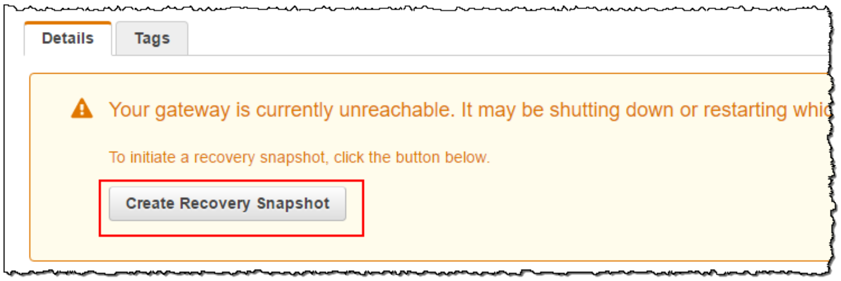

Your gateway is currently unreachable. It may be shutting down or restarting whic

To initiate a recovery snapshot, click the button below.

4. Choose **Create recovery snapshot** to open the **Create recovery snapshot** dialog box.
5. From the list of volumes displayed, choose the volume you want to recover, and then choose **Create snapshots**.

   AWS Storage Gateway initiates the snapshot process.
6. Find the snapshot using the steps in the procedure Finding a Snapshot (p. 255).

7. Restore the snapshot using one of the procedures in Restoring a Snapshot (p. 261).

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

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

If your tape gateway or the hypervisor host encounters an unrecoverable failure, you can recover the tapes from that tape gateway to another tape gateway.

AWS Storage Gateway periodically takes point-in-time snapshots of all the tapes in the library. These snapshots are called *recovery points.* You can recover tapes to another tape gateway from the latest recovery point. To recover tapes from a failed tape gateway to another tape gateway, use the following procedure.

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 Selecting a Gateway Type (p. 49).
2. Open the AWS 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 disable 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 Disabling Your Tape Gateway (p. 205).
7. From the tapes that the disabled 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 (p. 186).

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 a 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 AWS 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 Adding and Removing Upload Buffer (p. 144).

  > **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 Starting and Stopping Your Gateway (p. 141).
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 Adding and Removing Upload Buffer (p. 144).
3. Restart the gateway. For information about how to restart a gateway, see Starting and Stopping Your Gateway (p. 141).

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 Your Activated Gateway** page, configure the disk for cache storage. For information about how to do so, see Configuring Local Disks (p. 57).
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 AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Gateways**, and then choose your gateway.
3. On the **Action** menu, 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, AWS 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 Work with 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 Selecting a Gateway Type (p. 49).
- Disable 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 (p. 202).

  > **Note**
  > You have to reconfigure your iSCSI initiator and backup application to use the new tape gateway. For more information, see Connect Your Tape Gateway Devices to Your Windows Client (p. 59).

You Don't Need an IRRECOVERABLE Tape That Is Archived to archive and Retrieved to Tape Gateway

Suppose you have a virtual tape with the status IRRECOVERABLE, you don't need it, and the tape has previously been archived to archive and retrieved to your tape gateway. In this case, use your backup software to archive the tape in archive. For more information, see Archiving Tapes (p. 276).

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 (p. 117).

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.

Disabling Your Tape Gateway

If your tape gateway has failed and you want to recover the tapes to another gateway, you must first disable the failed gateway. Disabling a tape gateway locks down the virtual tapes in that gateway. That is, any data that you might write to these tapes after disabling the gateway will not be sent to AWS. For more information about recovering tapes, see You Need to Recover a Virtual Tape from a Malfunctioning Tape Gateway (p. 202).

Note that you can only disable a gateway on the Storage Gateway console if the gateway is no longer connected to AWS. If the gateway is connected to AWS, you won't see the **Disable Gateway** button in the console.

To disable your gateway

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Gateways** and choose the failed gateway.
3. Choose the **Details** tab for the gateway to display the disable gateway message.
4. Choose **Disable gateway**.

Additional AWS Storage Gateway Resources

In this section, you can find information about AWS and third-party software, tools, and resources that can help you set up or manage your gateway, and also about AWS Storage Gateway limits.

Topics

VM Setup

Topics

Components in Your VMware vSphere Environment for AWS Storage Gateway

You can use VMware to create an on-premises virtual machine to host a gateway using the AWS Storage Gateway service. You use a VMware client to interact with a VMware server and create your virtual machines. A gateway virtual machine definition—or template—that contains all the files and data for creating a new gateway is available from the AWS Storage Gateway detail page. The template is distributed as a single `.ova` file that is deployed on the VMware server. In this section, the components of the VMware vSphere environment that you need to know to use the AWS Storage Gateway service are discussed.

The following table describes the subset of vSphere components that you typically work with when using the AWS Storage Gateway service.

Component	Description
VMware vSphere	The VMware virtualization platform for managing its virtual computing infrastructure including the client and server.
VMware ESXi hypervisor OS (vSphere Server)	The VMware server OS that hosts the gateway virtual machine. You interact with the OS through the vSphere client GUI. To provision a gateway in AWS Storage Gateway, you only need to access the host during the activation of the gateway. For all other management and maintenance-related functions, you use the AWS Management Console.
VMware vSphere Client (vSphere Client)	The VMware software that you use on your computer to access and manage your VMware environment. You manage your virtual machine (which contains the gateway) using the client.
VMware High Availability	VMware High Availability (HA) is a component of vSphere that can provide protection from failures in your infrastructure layer supporting a gateway virtual machine (VM). VMware HA does this by using multiple hosts configured as a cluster so that if one host running a gateway VM fails, the gateway VM can be restarted automatically on another host within the cluster. AWS Storage Gateway can be used with VMware HA. For more information about VMware HA, see VMware HA: Concepts and Best Practices.
Virtual machine	The software implementation of a computer that contains the components of AWS Storage Gateway. The virtual machine runs on the VMware vSphere platform.
OVA, OVF	A template that represents a customized virtual machine. The AWS Storage Gateway appliance is an Open Virtualization Format (OVF) package that is distributed in an Open Virtualization Application (OVA). The OVA template contains all the information needed to configure and start a gateway. You deploy the template using the client connected to a VMware server. For instructions about downloading the OVA template for AWS Storage Gateway, see the AWS Storage Gateway detail page.
Data store	The storage on the vSphere server where the files that define a virtual machine are stored.

Component	Description
	These files come from the OVA file provided as part of the service. When you deploy the OVA, you select a data store on which to store the file, if there is more than one data store for the VMware server.

Configuring a VMware ESXi Host for AWS Storage Gateway

Following, you can find some basic information to help you set up your virtualization host and perform some optional host configuration.

The AWS Storage Gateway service includes an on-premises software appliance that communicates with the AWS cloud storage infrastructure. The appliance is packaged as a virtual machine that you deploy on a host running the VMware ESX/ESXi virtualization software. For more information on the VMware virtualization software, see VMware vSphere Hypervisor on the VMware website. For requirements that your VMware environment must meet to run AWS Storage Gateway, see Requirements (p. 11).

To install the VMware vSphere hypervisor OS on your host

1. Insert the VMware vSphere hypervisor disk in the disk drive.
2. Restart the computer.

   Depending on your computer bios settings, the computer might automatically boot off your disk. If not, check the relevant settings to boot the computer from the hypervisor disk.
3. Follow the instructions on the monitor to install the VMware hypervisor OS.

   This installation wipes any existing content on the disk and installs the hypervisor.

   > **Tip**
   > After a successful VMware hypervisor host installation, the monitor displays the IP address of the host computer. Note down this IP address. You use the IP address to connect to the host.
4. Set the time on the host.

   For instructions, see Synchronizing VM Time with Host Time (p. 215).

In the preceding steps, you provisioned a host with VMware hypervisor. The hypervisor is aware of host computer configuration, such as available processors, memory, and local hard disks. The host provides these resources to the AWS Storage Gateway.

You can optionally configure this host by adding more storage, such as additional direct-attached disks or SAN disks. The following steps illustrate how you can add one or more SAN disks to this host.

To connect to the hypervisor host

1. Start the VMware vSphere client and connect to the host using the host IP address.

   The VMware vSphere Client dialog box appears.

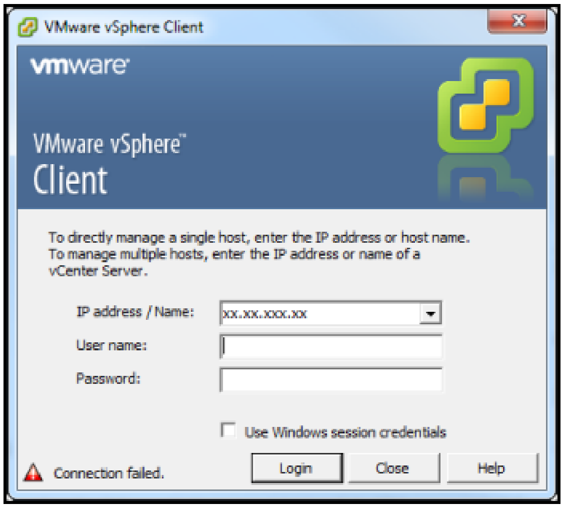

To directly manage a single host, enter the IP address or host name.
To manage multiple hosts, enter the IP address or name of a vCenter Server.

IP address / Name: xx.xx.xxx.xx

User name:

Password:

□ Use Windows session credentials

Connection failed.

2. For **IP Address**, type the IP address of the host.
3. For **User Name** and **Password**, type your host user name and password.
4. Choose **Login**.

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

To add a new iSCSI target

1. After you have connected to your remote device through the hypervisor, choose the **Configuration** tab of the host and choose **Storage** for **Hardware**.

   The **Datastores** pane shows the available data stores.

   For example, the following example shows that the host has two local hard drives available, datastore1 and datastore2.

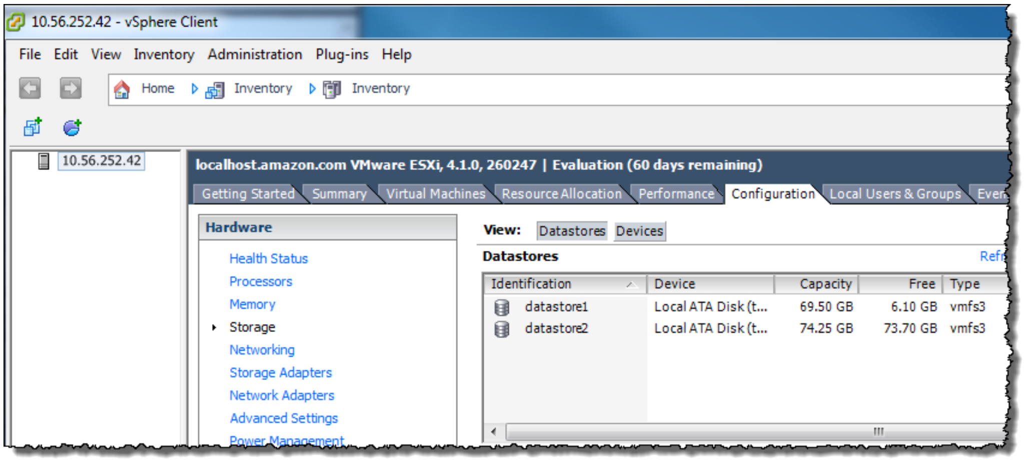

localhost.amazon.com VMware ESXi, 4.1.0, 260247 | Evaluation (60 days remaining)

Datastores

Identification	Device	Capacity	Free	Type
datastore1	Local ATA Disk (t...	69.50 GB	6.10 GB	vmfs3
datastore2	Local ATA Disk (t...	74.25 GB	73.70 GB	vmfs3

2. For **Hardware**, choose **Storage Adapters**.
3. In the **Storage Adapters** pane, choose **iSCSI Software Adapter**, and then choose **Properties** in the **Details** pane.

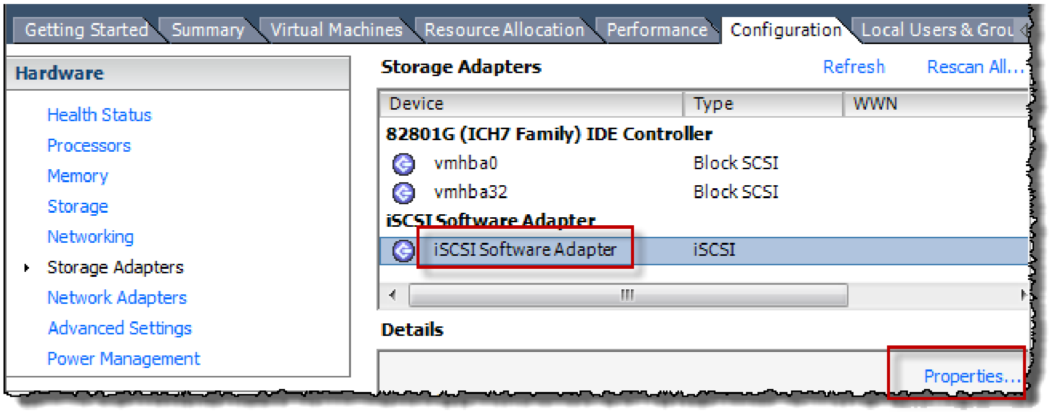

4. In the **iSCSI Initiator (iSCSI Software Adapter) Properties** dialog box, choose **Configure**.

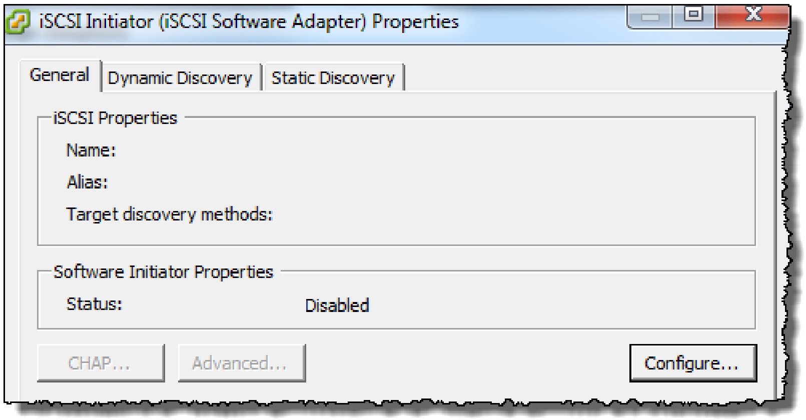

5. In the **General Properties** dialog box, select the **Enabled** check box to set the software initiator status to enabled, and then choose **OK**.

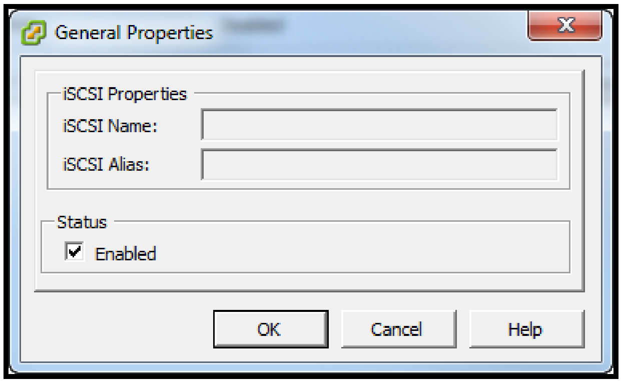

6. In the **iSCSI Initiator (iSCSI Software Adapter) Properties** dialog box, choose the **Dynamic Discovery** tab, and then choose **Add** to add an iSCSI target.

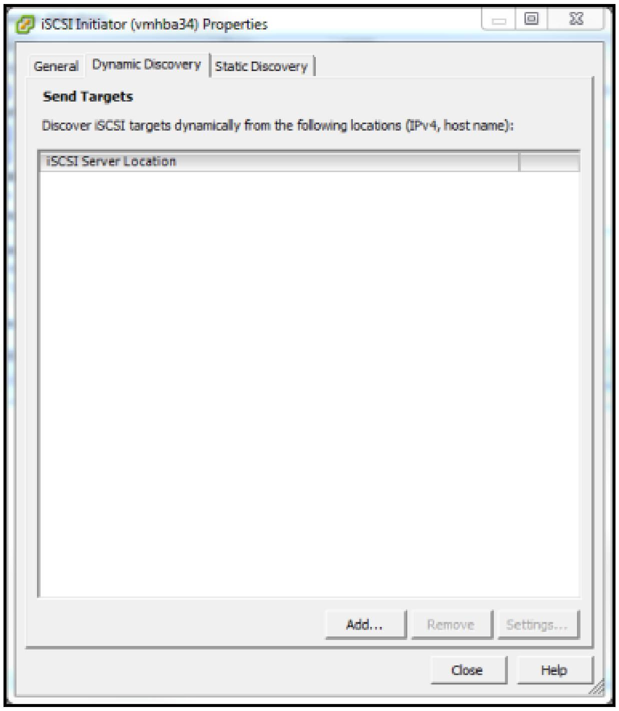

Discover iSCSI targets dynamically from the following locations (IPv4, host name):

7. In the **Add Send Target Server** dialog box, type a name for **iSCSI Server** and a port for **Port**, and then choose **OK**.

   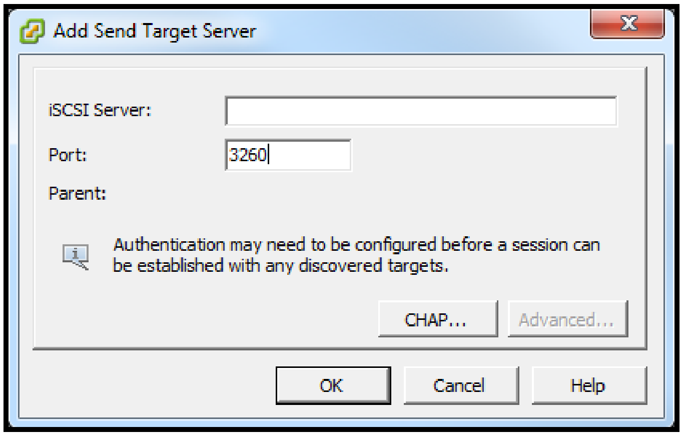

   Port: 3260

   Authentication may need to be configured before a session can be established with any discovered targets.

   Type the IP address or DNS name of the storage system.

   The new iSCSI server location that you type here appears in the **Sends Target** list on the **Dynamic Discovery** tab.

8. Choose **Close** to close the **iSCSI Initiator (iSCSI Software Adapter) Properties** dialog box.

Now you have added a new iSCSI target in the host configuration.

Deploying the AWS Storage Gateway VM to Your VMWare Host

You can choose to run AWS Storage Gateway on-premises either as a VMware ESXi Hypervisor or Microsoft Hyper-V Hypervisor virtual machine (VM) appliance. Following, you can find out how to deploy your gateway on a VMware ESXi Hypervisor. AWS Storage Gateway supports VMware ESXi Hypervisor versions 4.1, 5.0, 5.1, 5.5 and 6.0. For information about supported hosts, see Supported Hypervisors and Host Requirements (p. 16)

Deploying your gateway on VMWare hypervisor is a three-step process. When you complete all the tasks in this section, you can begin the activation process for your gateway.

Topics

Deploying Your Gateway on VMWare Host

To deploy the AWS Storage Gateway VM to your host

Note
Depending on the version of hypervisor you use, the graphical user interface (GUI) shown in these procedures may look slightly different. For more information, see Supported Hypervisors and Host Requirements (p. 16).

1. To connect to your hypervisor host, start the VMware vSphere client on Windows, type your host's IP and your credentials in the login dialog box, and then choose **Login**.
2. Deploy the AWS Storage Gateway VM on the host using the following steps.

   a. On the **File** menu of the vSphere client, choose **Deploy OVF Template**.

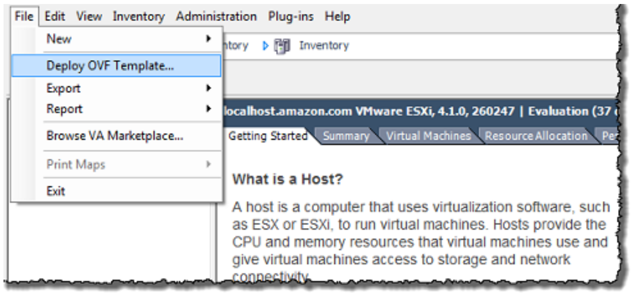

What is a Host?

A host is a computer that uses virtualization software, such as ESX or ESXi, to run virtual machines. Hosts provide the CPU and memory resources that virtual machines use and give virtual machines access to storage and network connectivity.

   Doing this opens the **Deploy OVF Template** wizard, a series of steps to provide information required to deploy the VM.

   b. In the **Source** pane, provide the file path to the AWS Storage Gateway `.ova` package, and then choose **Next**.

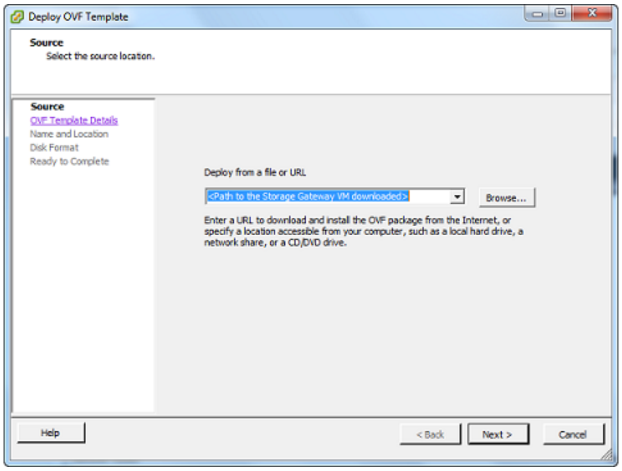

Deploy OVF Template

Source
Select the source location.

Source
OVF Template Details
Name and Location
Disk Format
Ready to Complete

Deploy from a file or URL

<Path to the Storage Gateway VM downloaded>

Enter a URL to download and install the OVF package from the Internet, or specify a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

c. In the **OVF Template Details** pane, choose **Next**.

d. In the **Name and Location** pane, type the VM name for **Name**, and then choose **Next**.

Although the vSphere client displays this VM name, AWS Storage Gateway does not use this name.

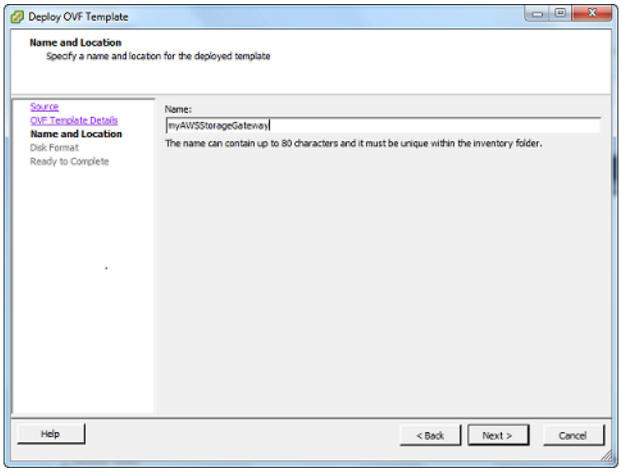

Deploy OVF Template

Name and Location
Specify a name and location for the deployed template

Source
OVF Template Details
Name and Location
Disk Format
Ready to Complete

Name:
myAWSStorageGateway

The name can contain up to 80 characters and it must be unique within the inventory folder.

e. The following **Datastore** pane appears only if your host has multiple data stores. In this pane, you select the data store where you want to deploy the VM, and then choose **Next**. If your host has only one data store, skip to the next step.

A *data store* is a virtual representation of underlying physical storage resources. The following example shows a host that has two data stores: datastore1 and datastore2.

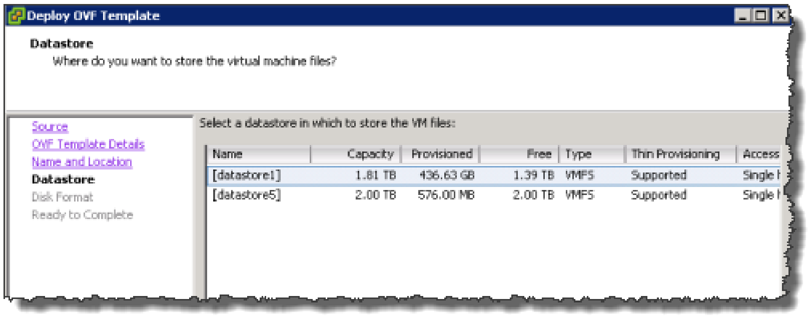

Deploy OVF Template

Datastore
Where do you want to store the virtual machine files?

Source
OVF Template Details
Name and Location
Datastore
Disk Format
Ready to Complete

Select a datastore in which to store the VM files:

Name	Capacity	Provisioned	Free	Type	Thin Provisioning	Access
[datastore1]	1.81 TB	436.63 GB	1.39 TB	VMFS	Supported	Single
[datastore5]	2.00 TB	576.00 MB	2.00 TB	VMFS	Supported	Single

f. In the **Disk Format** pane, select **Thick provisioned format**, and then choose **Next**.

   When you use thick provisioning, the disk storage is allocated immediately, resulting in better performance. In contrast, thin provisioning allocates storage on demand. On-demand allocation can affect the normal functioning of AWS Storage Gateway. For AWS Storage Gateway to function properly, the VM disks must be stored in thick-provisioned format.

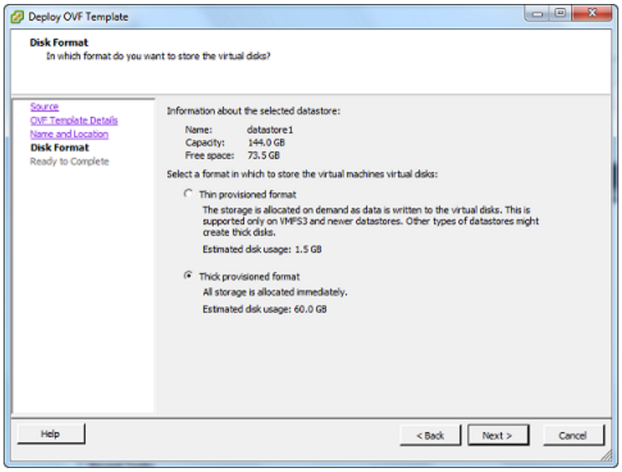

Disk Format
In which format do you want to store the virtual disks?

Information about the selected datastore:

Name: datastore1
Capacity: 144.0 GB
Free space: 73.5 GB

Select a format in which to store the virtual machines virtual disks:

○ Thin provisioned format
The storage is allocated on demand as data is written to the virtual disks. This is supported only on VMFS3 and newer datastores. Other types of datastores might create thick disks.
Estimated disk usage: 1.5 GB

● Thick provisioned format
All storage is allocated immediately.
Estimated disk usage: 60.0 GB

g. In the **Ready to Complete** pane, choose **Finish**.

   The AWS Storage Gateway VM starts deploying to your host.

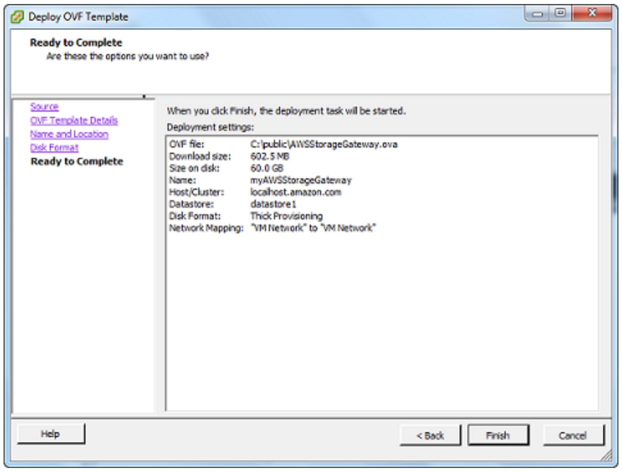

Ready to Complete
Are these the options you want to use?

When you click Finish, the deployment task will be started.

Deployment settings:

```
OVF file:         C:\public\AWSStorageGateway.ova
Download size:    602.5 MB
Size on disk:     60.0 GB
Name:             myAWSStorageGateway
Host/Cluster:     localhost.amazon.com
Datastore:        datastore1
Disk Format:      Thick Provisioning
Network Mapping:  "VM Network" to "VM Network"
```

h. View the details of the new VM to verify the information about the VM.

   i. Depending on the state of your vSphere client, you might need to choose the **Inventory** icon to view the host object that contains the new VM.

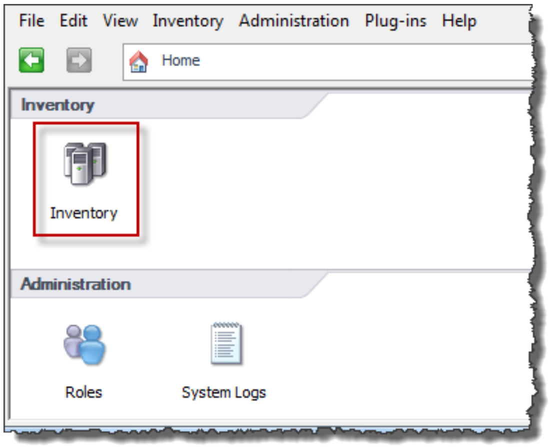

ii. Expand the host object to view the details of the new VM.

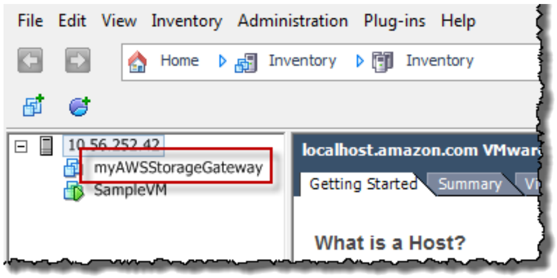

Synchronizing VM Time with Host Time

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.

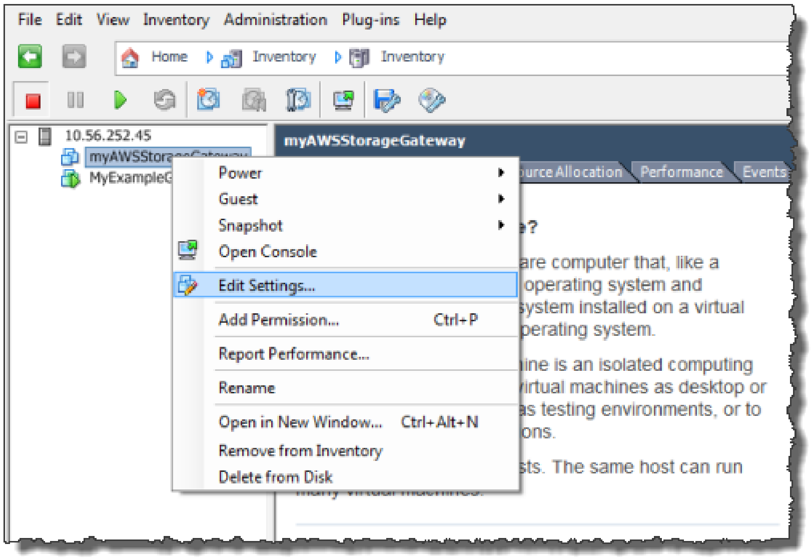

b. Choose the **Options** tab, and choose **VMware Tools** in the options list.

c. Check the **Synchronize guest time with host** option, and then choose **OK**.

The VM synchronizes its time with the host.

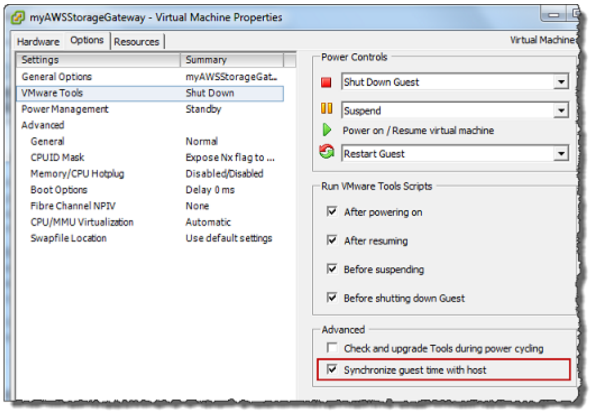

Settings	Summary
General Options	myAWSStorageGat...
VMware Tools	Shut Down
Power Management	Standby
Advanced	
General	Normal
CPUID Mask	Expose Nx flag to ...
Memory/CPU Hotplug	Disabled/Disabled
Boot Options	Delay 0 ms
Fibre Channel NPIV	None
CPU/MMU Virtualization	Automatic
Swapfile Location	Use default settings

Power Controls

Shut Down Guest

Suspend

Power on / Resume virtual machine

Restart Guest

Run VMware Tools Scripts

☑ After powering on

☑ After resuming

☑ Before suspending

☑ Before shutting down Guest

Advanced

□ Check and upgrade Tools during power cycling

☑ Synchronize guest time with 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.

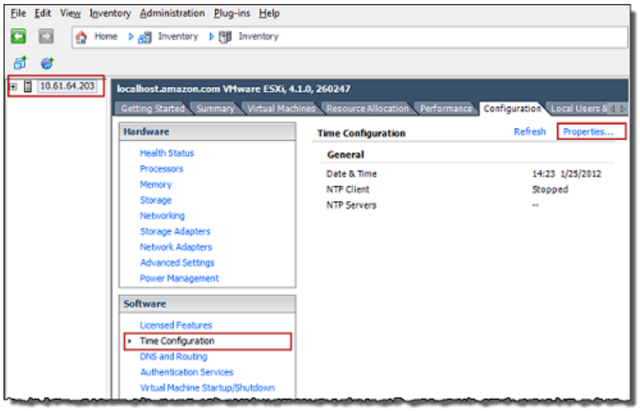

Date & Time 14:23 1/25/2012
NTP Client Stopped
NTP Servers --

c. In the **Date and Time** panel, set the date and time.

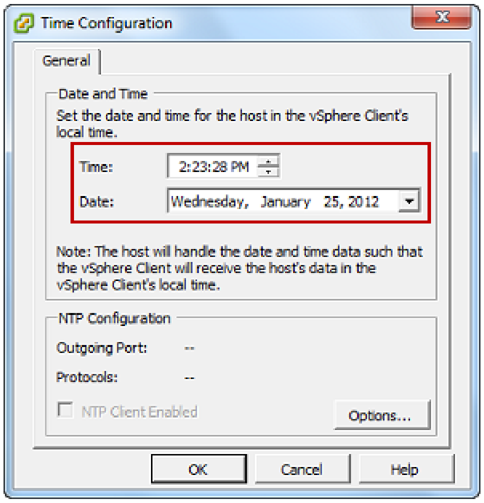

Set the date and time for the host in the vSphere Client's local time.

Time: 2:23:28 PM

Date: Wednesday, January 25, 2012

Note: The host will handle the date and time data such that the vSphere Client will receive the host's data in the vSphere Client's local time.

Outgoing Port: --

Protocols: --

□ NTP Client Enabled

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.

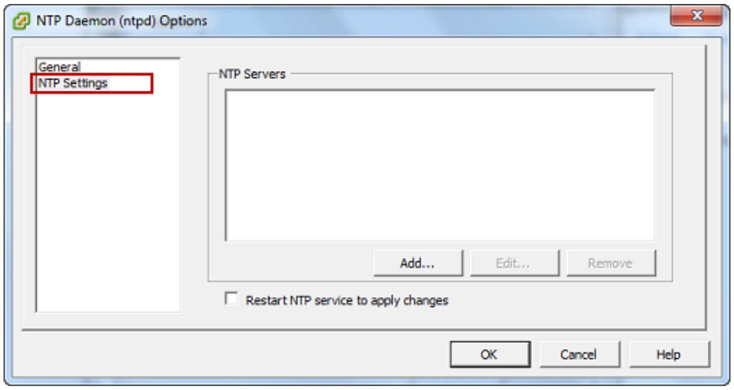

□ Restart NTP service to apply changes

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.

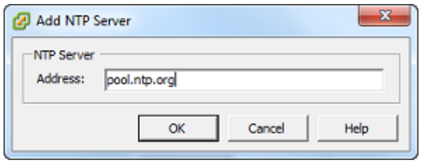

NTP Server

Address: pool.ntp.org

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.

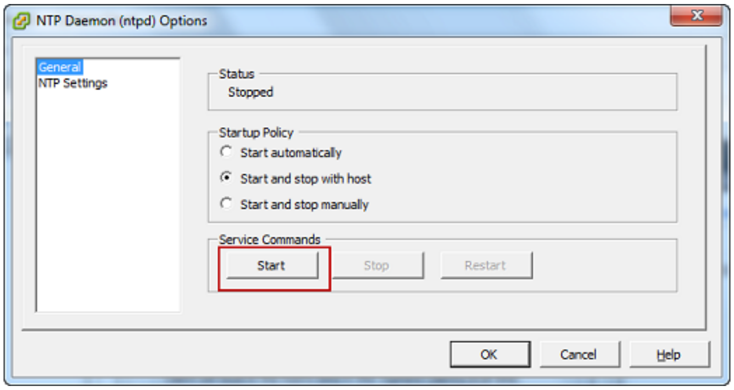

Status
Stopped

Startup Policy
○ Start automatically
● Start and stop with host
○ Start and stop manually

Service Commands

e. Choose **OK** to close the **NTP Daemon (ntpd) Options** dialog box.

f. Choose **OK** to close the **Time Configuration** dialog box.

Provisioning Local Disk Storage for the Gateway VM (VMWare)

Next, you allocate local storage for your deployed gateway VM. After you activate your gateway, you configure this local storage for the gateway's use.

Topics

Deciding the Amount of Local Disk Storage

In this step, you decide the number and size of disks to allocate for your gateway, as follows:

- All gateways require one or more disks for an upload buffer.
- Depending on the storage solution you deploy (see Plan Your AWS Storage Gateway Deployment (p. 8)), the gateway requires the following additional storage:
  - Tape Gateway and gateways created with cached volumes require one or more disks for cache storage.
  - Gateways created with stored volumes require one or more disks for volume storage—one local disk required for each volume you will create on the gateway.

You need at least two disks to begin with. The following table recommends sizes for local disk storage for your deployed gateway.

Disk Size	Gateway-Cached	Gateway-Stored	Tape Gateway
Upload buffer (minimum)	150 GiB	150 GiB	150 GiB
Upload buffer (maximum)	2 TiB	2 TiB	2 TiB
Cache storage (minimum)	1.1 x upload buffer	—	1.1 x upload buffer
Cache storage (maximum)	16 TiB	—	16 TiB
Volume storage (minimum)	—	1 GiB	—
Volume storage (maximum)	32 TiB	16 TiB	—

Before going to the next step, decide the number and size of disks to allocate. You can add more local storage after you set the gateway up, and as your workload demands. If you're not sure what number and size of disks to use, provision two local disks with the minimum recommended size for your gateway type.

If you plan to deploy your gateway in production, you should consider your real workload in determining disk sizes. For information about disk size guidelines, see Adding and Removing Upload Buffer (p. 144) and Adding Cache Storage (p. 146).

For more information about how gateways use local storage, see How AWS Storage Gateway Works (Architecture) (p. 3).

In the next step, you allocate the local disk storage to the gateway VM you deployed.

Provisioning Local Disk Storage

You can add disks from direct attached storage (DAS) or from a storage area network (SAN). This section tells how to add a virtual disk from a DAS. For instructions on attaching iSCSI volumes from an existing SAN for this step, see To add a new iSCSI target (p. 209).

Repeat the following procedure for each local disk you planed in the preceding step.

To allocate a local disk from DAS

1. Start the VMware vSphere client, and then connect to your host.
2. In the client, open the context (right-click) menu for your gateway VM, and then choose **Edit Settings**.

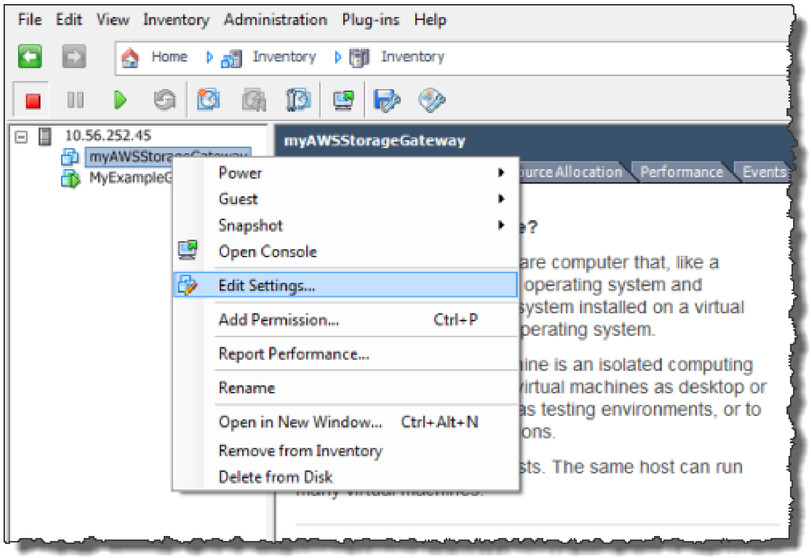

3. Choose the **Hardware** tab of the **Virtual Machine Properties** dialog box, and then choose **Add** to add a device.

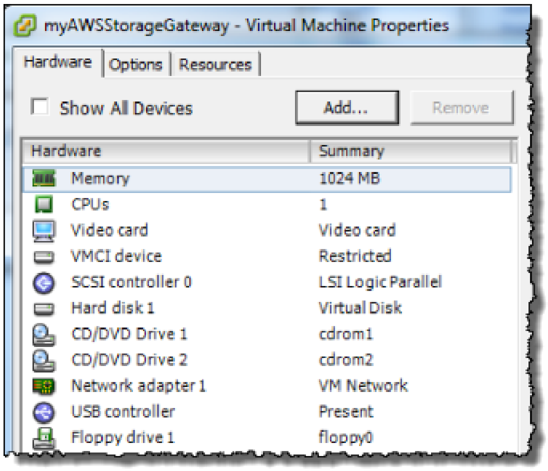

Hardware	Summary
Memory	1024 MB
CPUs	1
Video card	Video card
VMCI device	Restricted
SCSI controller 0	LSI Logic Parallel
Hard disk 1	Virtual Disk
CD/DVD Drive 1	cdrom1
CD/DVD Drive 2	cdrom2
Network adapter 1	VM Network
USB controller	Present
Floppy drive 1	floppy0

4. Follow the steps in the **Add Hardware** wizard to add a disk:

   a. In the **Device Type** pane, choose **Hard Disk** to add a disk, and then choose **Next**.

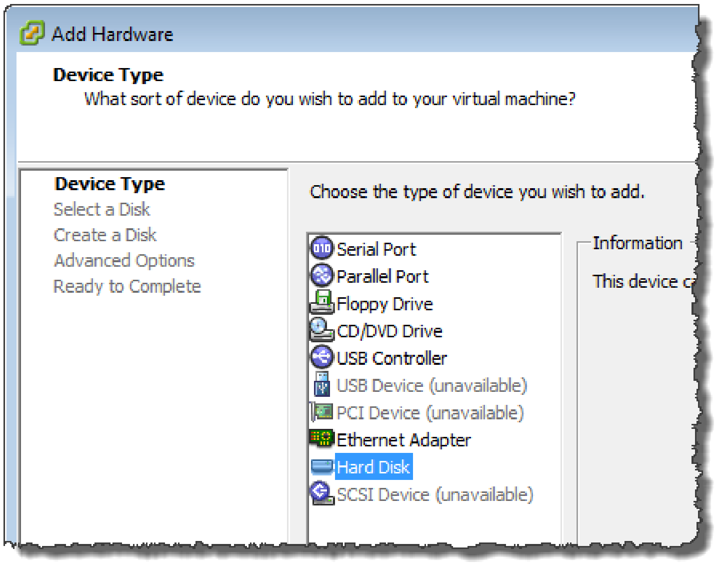

b. In the **Select a Disk** pane, select **Create a new virtual disk**, and then choose **Next**.

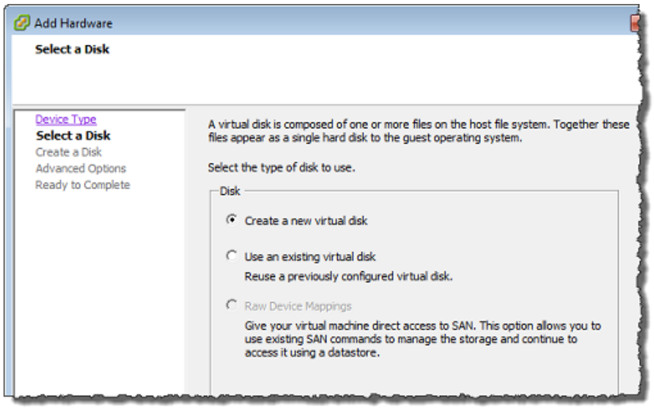

A virtual disk is composed of one or more files on the host file system. Together these files appear as a single hard disk to the guest operating system.

Select the type of disk to use.

Disk

● Create a new virtual disk

○ Use an existing virtual disk

Reuse a previously configured virtual disk.

○ Raw Device Mappings

Give your virtual machine direct access to SAN. This option allows you to use existing SAN commands to manage the storage and continue to access it using a datastore.

c. In the **Create a Disk** pane, specify the following.

- Capacity: The disk size you previously decided.

  Note
  The disk sizes used in these examples are not suitable for real-world workloads. If your setup is a trial setup, you can use small disk sizes.

- Disk provisioning: We recommend thick provisioning,
- Location: We strongly recommend that you choose the **Specify a datastore or datastore cluster** option and select a data store different from the one used for the gateway VM. In addition, we strongly recommend you dedicate a data store for the disks that you will use for upload buffer and another data store for other disks (see Deciding the Amount of Local Disk Storage (p. 218)). For this example, we store with the VM.

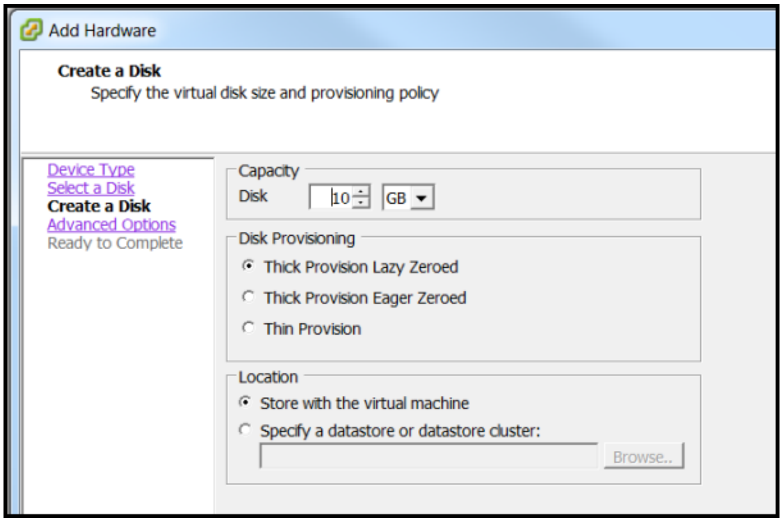

Add Hardware

Create a Disk

Specify the virtual disk size and provisioning policy

Capacity

Disk: 10 GB

Disk Provisioning

● Thick Provision Lazy Zeroed

○ Thick Provision Eager Zeroed

○ Thin Provision

Location

● Store with the virtual machine

○ Specify a datastore or datastore cluster:

d. Choose **Next**, and in the **Advanced Options** pane, accept the default values, and then choose **Next**.

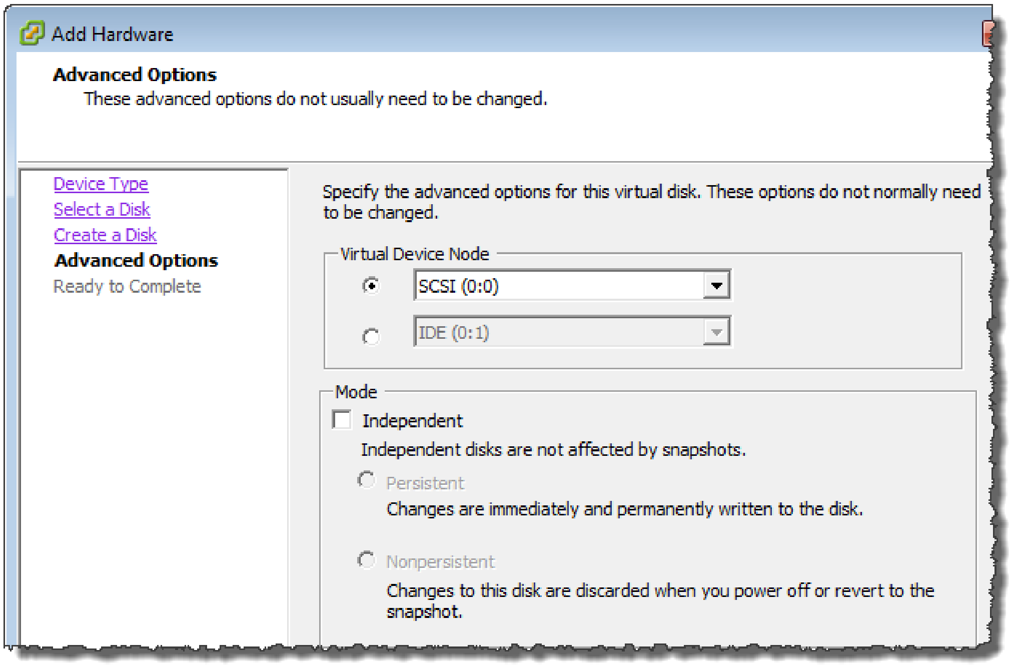

Add Hardware

Advanced Options

These advanced options do not usually need to be changed.

Specify the advanced options for this virtual disk. These options do not normally need to be changed.

Virtual Device Node

● SCSI (0:0)

○ IDE (0:1)

Mode

□ Independent

Independent disks are not affected by snapshots.

○ Persistent

Changes are immediately and permanently written to the disk.

○ Nonpersistent

Changes to this disk are discarded when you power off or revert to the snapshot.

e. In the **Ready to Complete** pane, accept the default values, and then choose **Finish**.

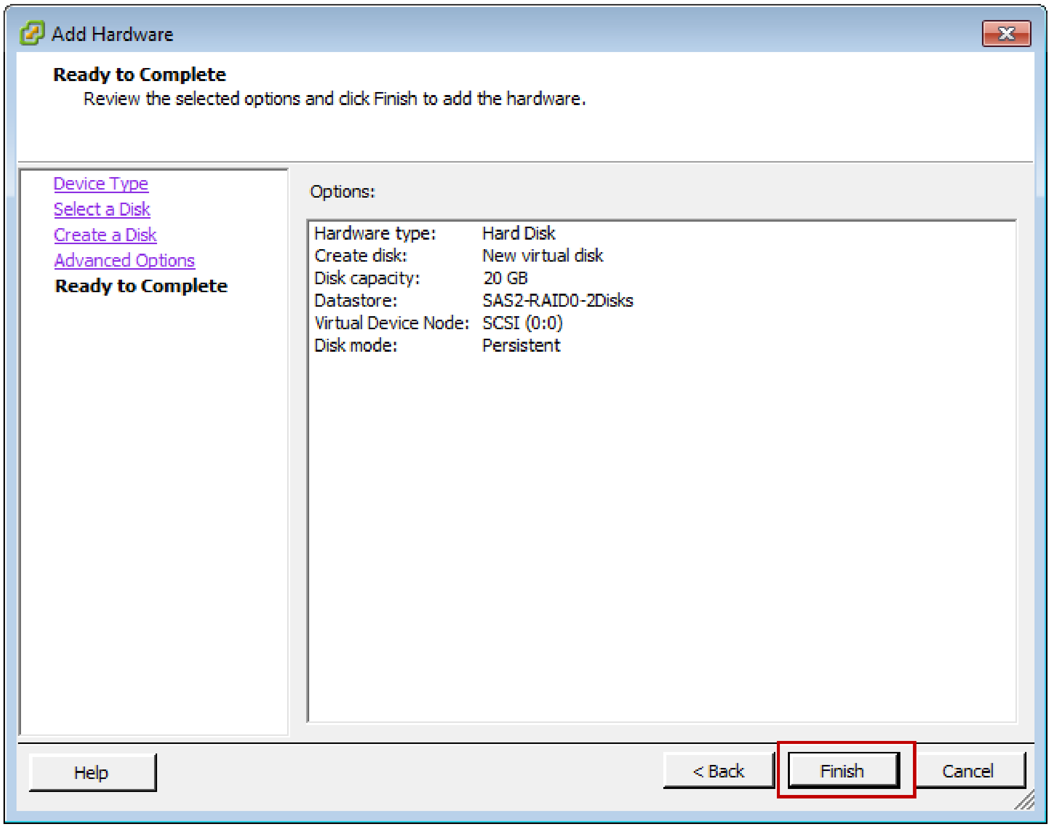

Hardware type:	Hard Disk
Create disk:	New virtual disk
Disk capacity:	20 GB
Datastore:	SAS2-RAID0-2Disks
Virtual Device Node:	SCSI (0:0)
Disk mode:	Persistent

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

5. Repeat steps 1 through 4 for each disk you planned to allocate.

6. Verify your gateway VM has the disks you added by following these steps:

   a. In the client, open the context (right-click) menu for your gateway VM and choose **Edit Settings**.

   b. In the **Hardware** tab of the **Virtual Machine Properties** dialog box, verify the disks in the hardware list.

      For example, the following screenshot shows two disks added to the gateway. The disks you add appear in the AWS Storage Gateway console as SCSI (0:0), SCSI (0:1), and so on.

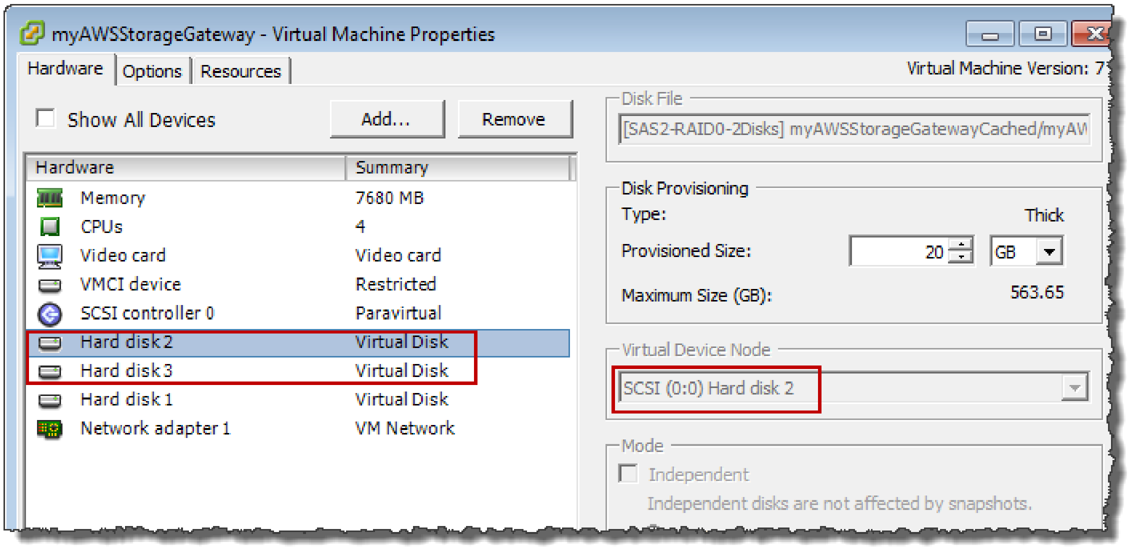

myAWSStorageGateway - Virtual Machine Properties

Hardware | Options | Resources — Virtual Machine Version: 7

□ Show All Devices — Add... — Remove

Hardware	Summary
Memory	7680 MB
CPUs	4
Video card	Video card
VMCI device	Restricted
SCSI controller 0	Paravirtual
Hard disk 2	Virtual Disk
Hard disk 3	Virtual Disk
Hard disk 1	Virtual Disk
Network adapter 1	VM Network

Disk File: [SAS2-RAID0-2Disks] myAWSStorageGatewayCached/myAW

Disk Provisioning
Type: Thick
Provisioned Size: 20 GB
Maximum Size (GB): 563.65

Virtual Device Node: SCSI (0:0) Hard disk 2

Mode
□ Independent
Independent disks are not affected by snapshots.

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**.

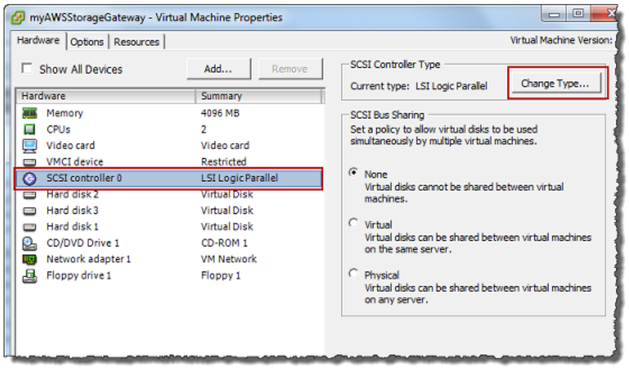

myAWSStorageGateway - Virtual Machine Properties

Hardware | Options | Resources — Virtual Machine Version:

□ Show All Devices — Add... — Remove

Hardware	Summary
Memory	4096 MB
CPUs	2
Video card	Video card
VMCI device	Restricted
SCSI controller 0	LSI Logic Parallel
Hard disk 2	Virtual Disk
Hard disk 3	Virtual Disk
Hard disk 1	Virtual Disk
CD/DVD Drive 1	CD-ROM 1
Network adapter 1	VM Network
Floppy drive 1	Floppy 1

SCSI Controller Type
Current type: LSI Logic Parallel — Change Type...

SCSI Bus Sharing
Set a policy to allow virtual disks to be used simultaneously by multiple virtual machines.

● None
Virtual disks cannot be shared between virtual machines.

○ Virtual
Virtual disks can be shared between virtual machines on the same server.

○ Physical
Virtual disks can be shared between virtual machines on any server.

3. In the **Change SCSI Controller Type** dialog box, select the **VMware Paravirtual** SCSI controller type, and then choose **OK**.

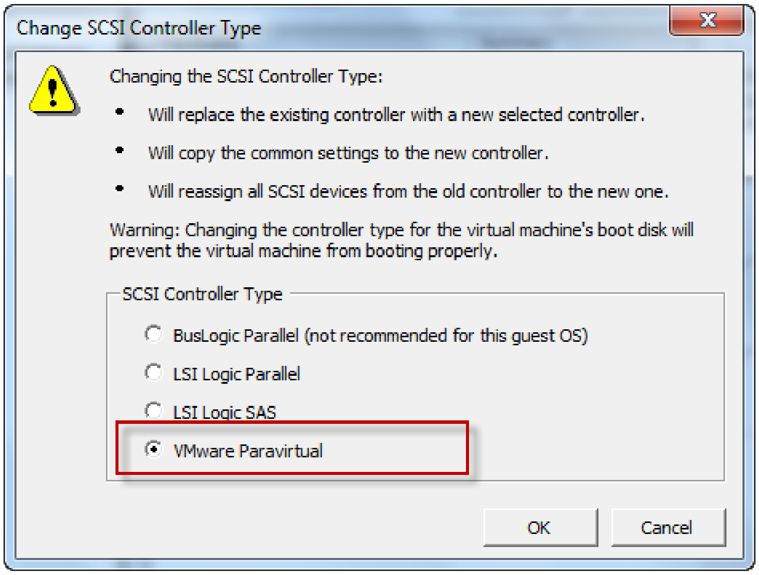

Change SCSI Controller Type

Changing the SCSI Controller Type:

- Will replace the existing controller with a new selected controller.
- Will copy the common settings to the new controller.
- Will reassign all SCSI devices from the old controller to the new one.

Warning: Changing the controller type for the virtual machine's boot disk will prevent the virtual machine from booting properly.

SCSI Controller Type

○ BusLogic Parallel (not recommended for this guest OS)
○ LSI Logic Parallel
○ LSI Logic SAS
● VMware Paravirtual

Now that you have deployed your gateway, synchronized the time and provisioned local disks for your gateway, you begin the activation process for your gateway.

Next Step

Activating Your Gateway (p. 55)

Using AWS 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 AWS Storage Gateway with VMware HA, we recommend doing the following things:

- Deploy the VMware ESX `.ova` downloadable package that contains the AWS 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 (p. 280). For more information on customizing Linux clients' timeout settings, see Customizing Your Linux iSCSI Settings (p. 45).
- 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.

Components in Your Hyper-V Environment for AWS Storage Gateway

You use Microsoft Hyper-V to create an on-premises virtual machine that hosts a storage gateway. You use the Hyper-V Manager to interact with a Hyper-V server and create your virtual machines. A gateway virtual machine definition—or *template*—that contains all the files and data for creating a new gateway is available from the AWS Storage Gateway console. The template is distributed as a single `.zip` file that you import into the Hyper-V server. This section discusses the components of the Microsoft Hyper-V environment that you need to know to use AWS Storage Gateway.

The following table describes the subset of Hyper-V components that you typically work with when using AWS Storage Gateway.

Component	Description
Microsoft Hyper-V	The Microsoft virtualization platform for managing a virtual computing infrastructure including the client and server.
Hyper-V hypervisor OS	The Hyper-V server operating system (OS) that hosts the gateway virtual machine. You interact with the OS through the Microsoft Hyper-V Manager graphical user interface. To provision a storage gateway, you need to only access the host during the activation of the gateway. For all other management and maintenance-related functions, you use the AWS Management Console.
Hyper-V Manager	The Hyper-V client software that you use on your computer to access and manage your Hyper-V environment. You manage your virtual machine (that contains the gateway) using the client.
Virtual machine	The software implementation of a computer that contains the components of AWS Storage Gateway. The virtual machine (VM) runs on the Microsoft Hyper-V platform.
Import files (packaging of VM)	The AWS Storage Gateway appliance is distributed as a compressed directory containing the following: • `Snapshots` folder, which will be empty for AWS Storage Gateway. • `Virtual Hard Disks` folder, which contains one virtual hard disk file called `AWS-Storage-Gateway.vhd`. **Note** AWS Storage Gateway supports the .vhdx file type, which enables you create larger virtual hard disks than other file types. • `Virtual Machines` folder, which contains an exported configuration files *`GUID`*`.exp`, where `GUID` is the virtual machine ID.

Component	Description
	• `config.xml`, which contains configuration information used for importing. **Note** This `config.xml` file is only provided in the Hyper-V 2008 R2 gateway appliance. You deploy AWS Storage Gateway to Hyper-V by first uncompressing the directory and then importing the uncompressed folder using the Hyper-V Manager.

Configuring a Hyper-V Host for AWS Storage Gateway

For basic information on how you can set up, configure, and troubleshoot your Microsoft Hyper-V 2008 R2 virtualization host, see the following:

- How to set up and configure the host:
  - Installing Microsoft Hyper-V (p. 227)
  - Connecting to Microsoft Hyper-V Host (p. 229)
  - Configuring Virtual Network Settings (p. 230)
  - Adding a Virtual Disk Backed by a Hard Disk (p. 232)
- How to troubleshoot certain basic setup issues: Troubleshooting Your Microsoft Hyper-V Setup (p. 233)

Setting Up and Configuring a Microsoft Hyper-V Host

The AWS Storage Gateway service includes an on-premises software appliance that communicates with the AWS Cloud storage infrastructure. The appliance is packaged as a virtual machine that you can deploy on a host running Microsoft Hyper-V virtualization software. For more information on the Microsoft Hyper-V software, see Microsoft Server Virtualization on the Microsoft website. For requirements that your Hyper-V environment must meet to run AWS Storage Gateway, see Requirements (p. 11).

Installing Microsoft Hyper-V

Following you can find a procedure for installing Microsoft Hyper-V. If you already have a Microsoft Hyper-V virtualization environment or your environment will be set up by an administrator familiar with the platform, then you do not need to understand these steps in detail.

Refer to the Hyper-V Getting Started Guide on the *Microsoft TechNet* site for more information about the installation of Hyper-V.

To install the Microsoft Hyper-V hypervisor OS on your host

1. Insert the Microsoft Hyper-V disk in the disk drive.
2. Restart the computer.

   Depending on your computer's BIOS settings, the computer might automatically boot off your disk. If not, check the relevant settings to boot the computer from the hypervisor disk.

3. Follow the instructions on the monitor to install the Hyper-V hypervisor OS.

   This installation wipes any existing content on the disk and installs the hypervisor.

   After a successful Hyper-V hypervisor host installation, you will be prompted to create an Administrator account password. After creating this account, the monitor displays a **Server Configuration** menu where you will do further configuration of the host.

4. In the **Server Configuration** menu, configure the host. We recommend the following:

To	Do This
Configure remote management.	Choose option 4 and then enable the following: • Option 1, **Allow MMC Remote Management** • Option 2, **Enable Windows PowerShell** • Option 3, **Allow Server Manager Remote Management**
Find the network address of the host.	Choose option 8 and follow the prompts. Note the IP address for use later.
Set the date and time.	Choose option 9 and follow the prompts.
(Optional) Change the computer name.	Choose option 2 and follow the prompts. Because this change requires a reboot, you might want to make this configuration change last.
(Optional) Add local administrators.	Choose option 3 and follow the prompts.
(Optional) Enable remote desktop.	Choose option 7 and follow the prompts.

   The following example shows a Server Configuration menu.

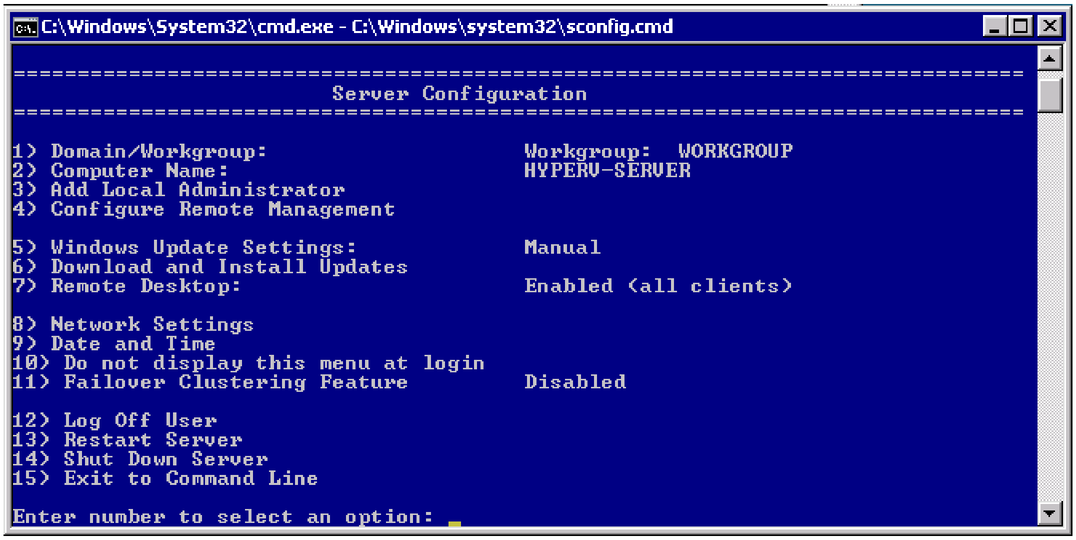
C:\Windows\System32\cmd.exe - C:\Windows\system32\sconfig.cmd

```
===============================================================================
                         Server Configuration
===============================================================================

1) Domain/Workgroup:                    Workgroup:  WORKGROUP
2) Computer Name:                       HYPERV-SERVER
3) Add Local Administrator
4) Configure Remote Management

5) Windows Update Settings:             Manual
6) Download and Install Updates
7) Remote Desktop:                      Enabled (all clients)

8) Network Settings
9) Date and Time
10) Do not display this menu at login
11) Failover Clustering Feature         Disabled

12) Log Off User
13) Restart Server
14) Shut Down Server
15) Exit to Command Line

Enter number to select an option: _
```

5. (Optional) You might need to put the IP address of the hypervisor host in your hosts file of client computers that connect to the hypervisor host.

   For example, in Windows 7 and 8, the hosts file can be found at this location:

```
%SystemRoot%\system32\drivers\etc\hosts
```

Connecting to Microsoft Hyper-V Host

The Hyper-V Manager runs on your client computer and connects to the hypervisor host. You use the Microsoft Hyper-V Manager to import, configure, and start the AWS Storage Gateway VM.

To connect to the hypervisor host

1. Start the Microsoft Hyper-V Manager (virtmgmt.msc).

   Note
   The Hyper-V Manager is a feature that you enable for your client computer. For more information about enabling it, see Install and Configure Hyper-V Tools for Remote Administration on the Microsoft Windows Server website.

   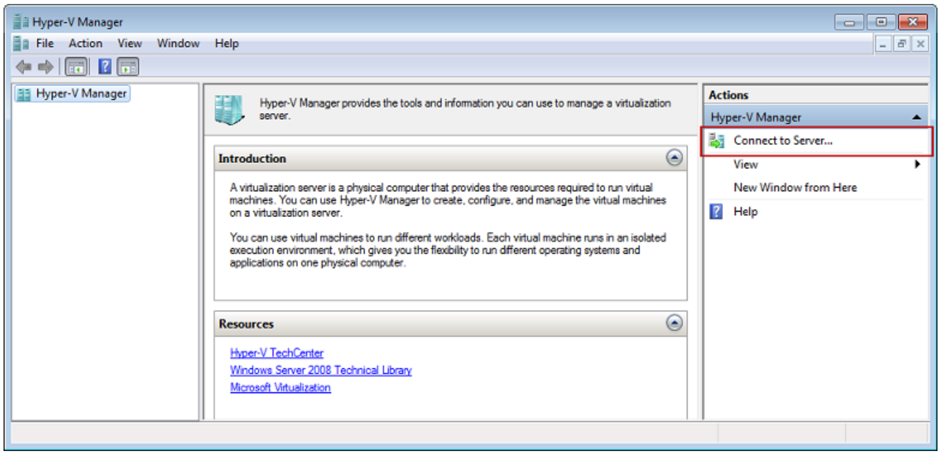

   Hyper-V Manager provides the tools and information you can use to manage a virtualization server.

   A virtualization server is a physical computer that provides the resources required to run virtual machines. You can use Hyper-V Manager to create, configure, and manage the virtual machines on a virtualization server.

   You can use virtual machines to run different workloads. Each virtual machine runs in an isolated execution environment, which gives you the flexibility to run different operating systems and applications on one physical computer.

   Hyper-V TechCenter
   Windows Server 2008 Technical Library
   Microsoft Virtualization

2. In the **Actions** pane, choose **Connect to Server**.
3. In the **Select Computer** dialog box, choose **Another computer**, either type the IP address of the hypervisor host or the host name, and then choose **OK**.

   Note
   To connect to a hypervisor host using the host name, you might need to make an entry in your `hosts` file so that the host name can be mapped to the correct IP address.

   Note
   If you have not been added to the local administrators for the hypervisor host, you might be prompted for credentials.

   The following example shows Hyper-V Manager connected to a hypervisor host called `HYPERV-SERVER` with one gateway VM.

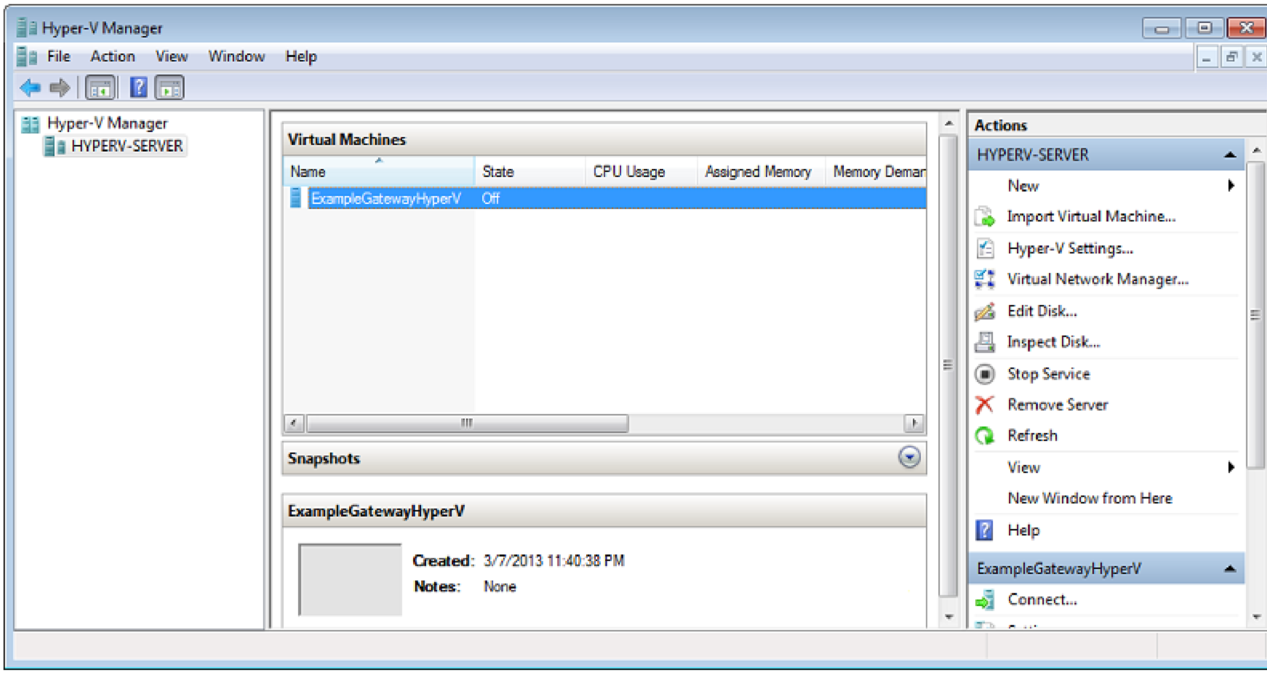

Name	State	CPU Usage	Assigned Memory	Memory Demar
ExampleGatewayHyperV	Off			

Created: 3/7/2013 11:40:38 PM
Notes: None

Configuring Virtual Network Settings

After you install and configure a Microsoft Hyper-V host, we recommend that you set up virtual networks by creating a new virtual network and associating it with a network interface of the host. Later, when you configure your gateway VM, you must associate it with one or more virtual networks so that the VM has connectivity.

To configure virtual network settings for the your Hyper-V host

1. Start the Microsoft Hyper-V Manager (virtmgmt.msc).
2. In the hypervisor host list (left pane), select your hypervisor.
3. In the **Actions** menu, under the hypervisor host name (for example, `HYPERV-SERVER`), choose **Virtual Network Manager**.

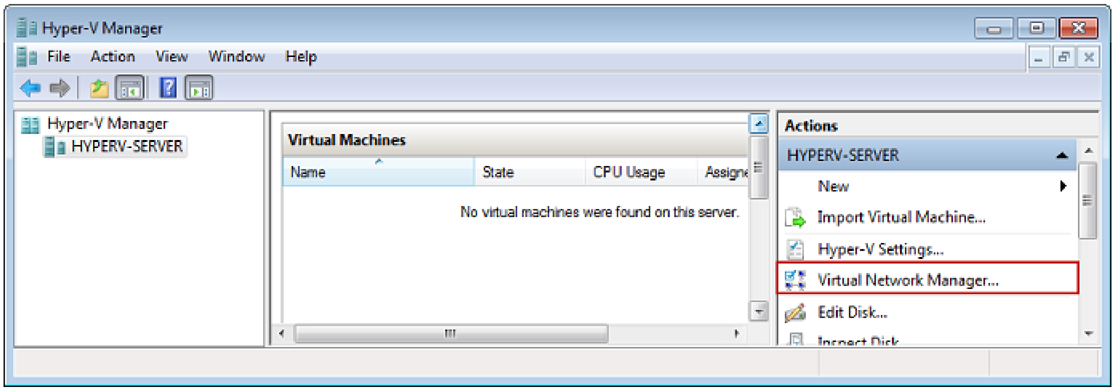

Name	State	CPU Usage	Assigne

No virtual machines were found on this server.

4. In the **Virtual Network Manager** dialog box, choose **New virtual network**.

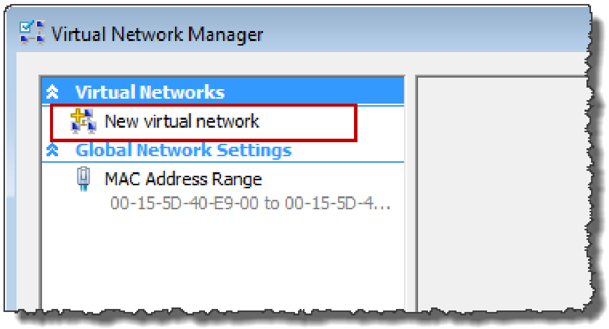

5. Choose **External** as the virtual network type, and then choose **Add**.

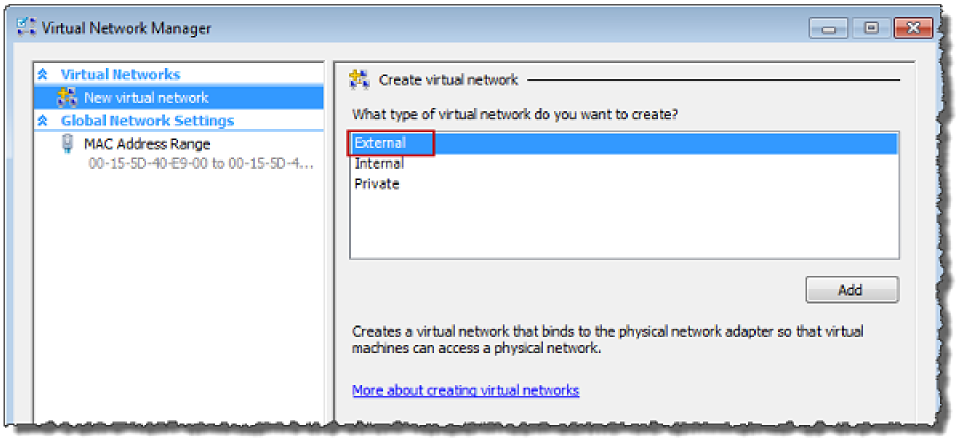

What type of virtual network do you want to create?

- External (selected)
- Internal
- Private

Creates a virtual network that binds to the physical network adapter so that virtual machines can access a physical network.

6. Type a name for the network, and then choose **OK**.

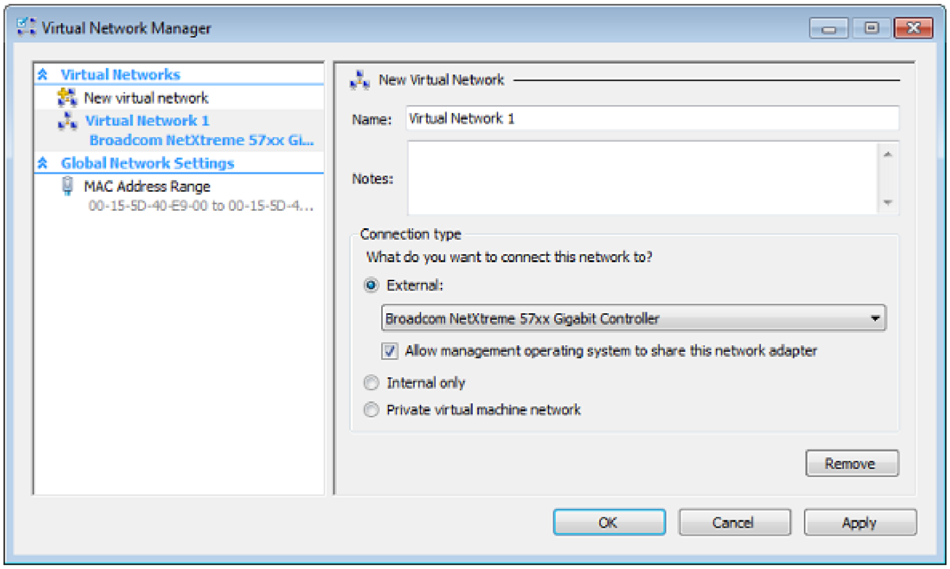

Name: Virtual Network 1

Notes:

Connection type

What do you want to connect this network to?

● External:

Broadcom NetXtreme 57xx Gigabit Controller

☑ Allow management operating system to share this network adapter

○ Internal only

○ Private virtual machine network

When you configure your gateway virtual machine, you can use this virtual network.

Adding a Virtual Disk Backed by a Hard Disk

In a preceding section, you provisioned a host with Hyper-V hypervisor. The hypervisor is aware of host computer configuration, such as available processors, memory, and local hard disks. The host provides these resources to AWS Storage Gateway. You can optionally configure this host by adding more storage, such as additional direct-attached disks or SAN disks. In this section, we show you how to add a virtual disk backed by a direct-attached disk.

To add a virtual disk backed by a physical hard disk

1. Start the Microsoft Hyper-V Manager (virtmgmt.msc).
2. Choose the VM.
3. For the **Actions** list for the VM, choose **Settings**.
4. For **Hardware**, choose **SCSI Controller**.
5. Choose **Hard Drive** in the **SCSI Controller** pane, and then choose **Add**.

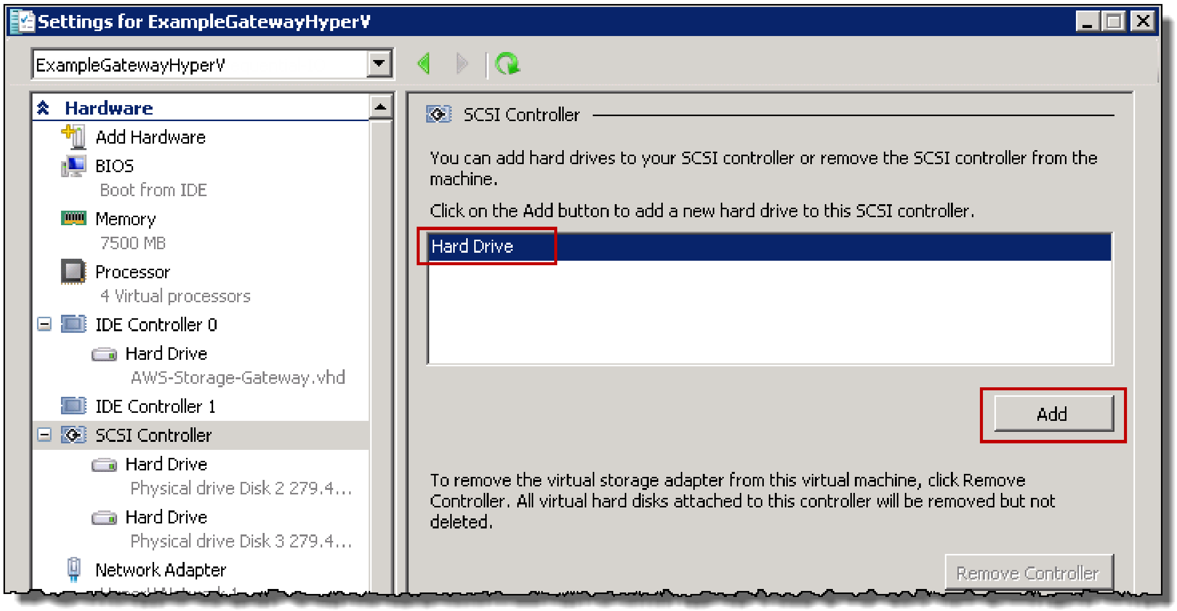

Settings for ExampleGatewayHyperV

ExampleGatewayHyperV

Hardware
- Add Hardware
- BIOS: Boot from IDE
- Memory: 7500 MB
- Processor: 4 Virtual processors
- IDE Controller 0
  - Hard Drive: AWS-Storage-Gateway.vhd
- IDE Controller 1
- SCSI Controller
  - Hard Drive: Physical drive Disk 2 279.4...
  - Hard Drive: Physical drive Disk 3 279.4...
- Network Adapter

SCSI Controller

You can add hard drives to your SCSI controller or remove the SCSI controller from the machine.

Click on the Add button to add a new hard drive to this SCSI controller.

To remove the virtual storage adapter from this virtual machine, click Remove Controller. All virtual hard disks attached to this controller will be removed but not deleted.

6. In the **Hard Drive** pane, choose **Physical hard disk**.

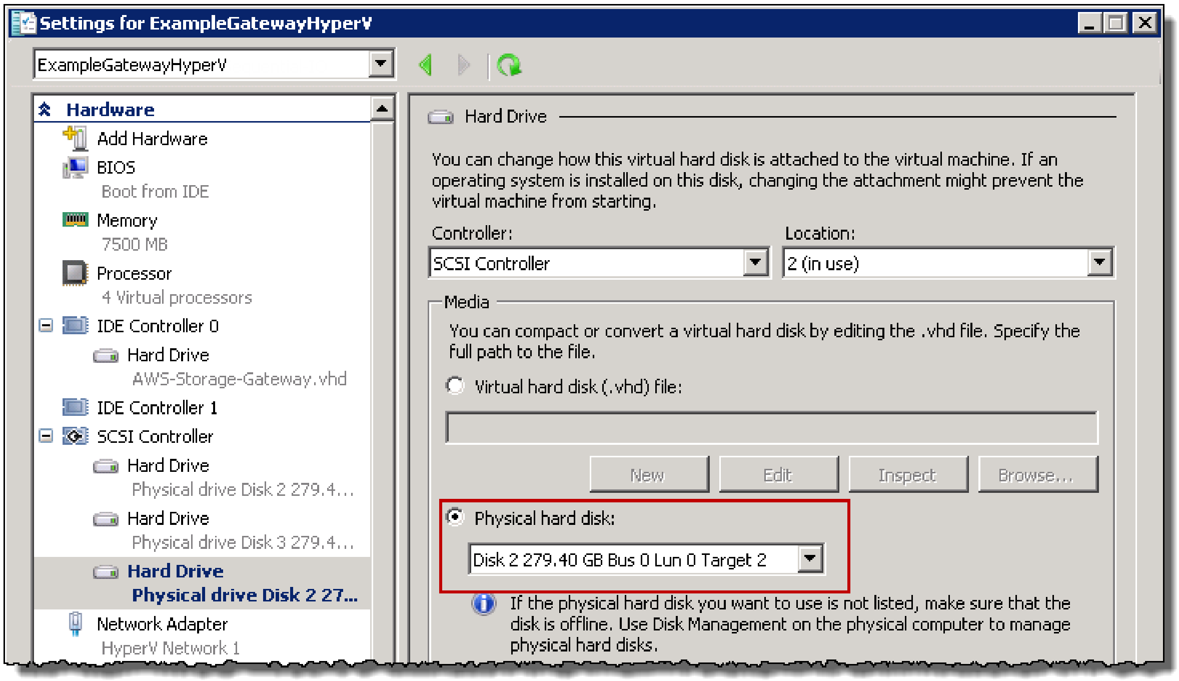

Settings for ExampleGatewayHyperV

ExampleGatewayHyperV

Hardware
- Add Hardware
- BIOS
  Boot from IDE
- Memory
  7500 MB
- Processor
  4 Virtual processors
- IDE Controller 0
  - Hard Drive
    AWS-Storage-Gateway.vhd
- IDE Controller 1
- SCSI Controller
  - Hard Drive
    Physical drive Disk 2 279.4...
  - Hard Drive
    Physical drive Disk 3 279.4...
  - **Hard Drive**
    Physical drive Disk 2 27...
- Network Adapter
  HyperV Network 1

Hard Drive

You can change how this virtual hard disk is attached to the virtual machine. If an operating system is installed on this disk, changing the attachment might prevent the virtual machine from starting.

Controller: SCSI Controller

Location: 2 (in use)

Media

You can compact or convert a virtual hard disk by editing the .vhd file. Specify the full path to the file.

○ Virtual hard disk (.vhd) file:

New | Edit | Inspect | Browse...

● Physical hard disk:

Disk 2 279.40 GB Bus 0 Lun 0 Target 2

If the physical hard disk you want to use is not listed, make sure that the disk is offline. Use Disk Management on the physical computer to manage physical hard disks.

7. Choose **OK**.

Troubleshooting Your Microsoft Hyper-V Setup

The following table lists typical issues that you might encounter when deploying AWS Storage Gateway on the Microsoft Hyper-V platform.

Issue	Action to Take
You try to import a gateway and receive the error message: "Import failed. Unable to find virtual machine import file under location ...". 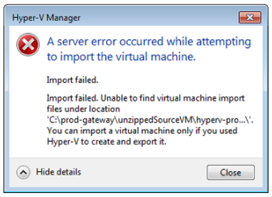 Hyper-V Manager A server error occurred while attempting to import the virtual machine. Import failed. Import failed. Unable to find virtual machine import files under location 'C:\prod-gateway\unzippedSourceVM\hyperv-pro...\'. You can import a virtual machine only if you used Hyper-V to create and export it. Hide details Close	This error can occur for the following reasons: • 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: 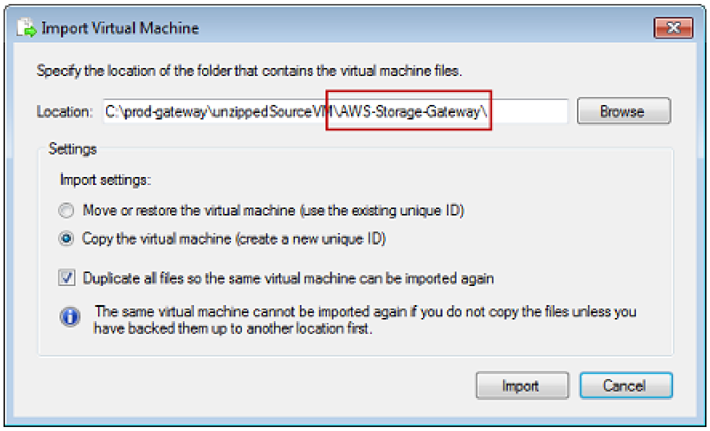 Import Virtual Machine Specify the location of the folder that contains the virtual machine files. Location: C:\prod-gateway\unzippedSourceVM\AWS-Storage-Gateway\ Browse Settings Import settings: ○ Move or restore the virtual machine (use the existing unique ID) ● Copy the virtual machine (create a new unique ID) ☑ Duplicate all files so the same virtual machine can be imported again The same virtual machine cannot be imported again if you do not copy the files unless you have backed them up to another location first. Import Cancel

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. 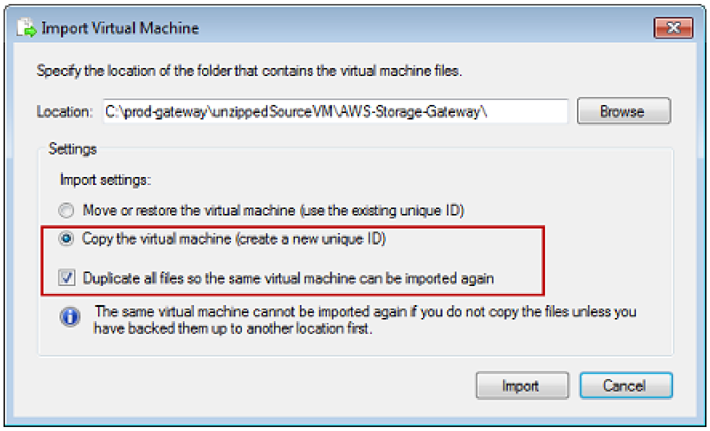 Import Virtual Machine Specify the location of the folder that contains the virtual machine files. Location: C:\prod-gateway\unzippedSourceVM\AWS-Storage-Gateway\ Browse Settings Import settings: ○ Move or restore the virtual machine (use the existing unique ID) ● Copy the virtual machine (create a new unique ID) ☑ Duplicate all files so the same virtual machine can be imported again The same virtual machine cannot be imported again if you do not copy the files unless you have backed them up to another location first. Import Cancel For more information about deploying a gateway, see Download and Deploy the AWS Storage Gateway VM on Your Host (p. 236).
You try to import a gateway and receive the error message: "Import failed. Import task failed to copy file." 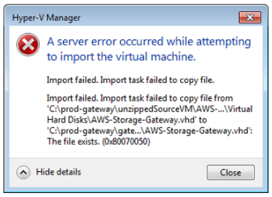 Hyper-V Manager A server error occurred while attempting to import the virtual machine. Import failed. Import task failed to copy file. Import failed. Import task failed to copy file from 'C:\prod-gateway\unzippedSourceVM\AWS-...\Virtual Hard Disks\AWS-Storage-Gateway.vhd' to 'C:\prod-gateway\gate...\AWS-Storage-Gateway.vhd': The file exists. (0x80070050) Hide details Close	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. 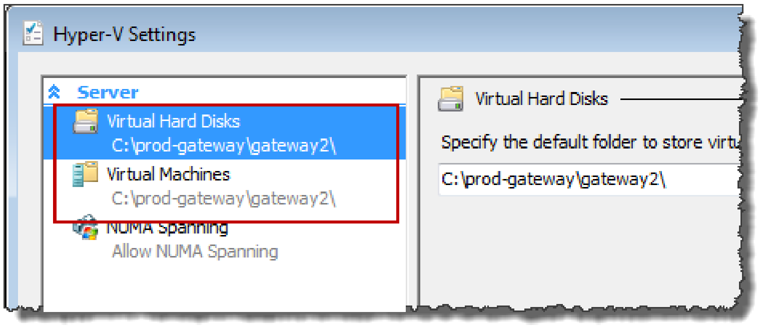 Hyper-V Settings Server Virtual Hard Disks C:\prod-gateway\gateway2\ Virtual Machines C:\prod-gateway\gateway2\ NUMA Spanning Allow NUMA Spanning Virtual Hard Disks Specify the default folder to store virtu C:\prod-gateway\gateway2\ For more information about deploying a gateway, see To import the VM (p. 239).

Issue	Action to Take
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." 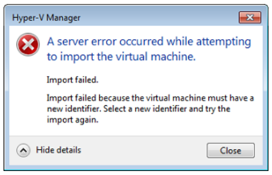 A server error occurred while attempting to import the virtual machine. Import failed. Import failed because the virtual machine must have a new identifier. Select a new identifier and try the import again. 	When you import the gateway make sure you select the **Copy the virtual machine** option and 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. 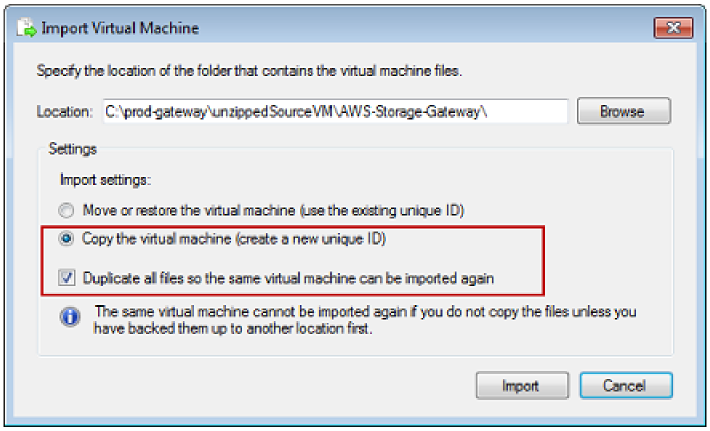 Specify the location of the folder that contains the virtual machine files. Location: C:\prod-gateway\unzippedSourceVM\AWS-Storage-Gateway\ Settings Import settings: ○ Move or restore the virtual machine (use the existing unique ID) ● Copy the virtual machine (create a new unique ID) ☑ Duplicate all files so the same virtual machine can be imported again The same virtual machine cannot be imported again if you do not copy the files unless you have backed them up to another location first. For more information about importing a gateway, see To import the VM (p. 239).
You try to start a gateway VM and receive an error message "The child partition processor setting is incompatible with parent partition." 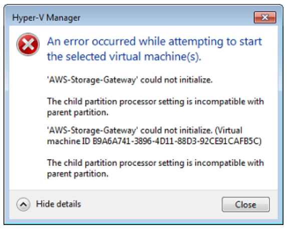 An error occurred while attempting to start the selected virtual machine(s). 'AWS-Storage-Gateway' could not initialize. The child partition processor setting is incompatible with parent partition. 'AWS-Storage-Gateway' could not initialize. (Virtual machine ID B9A6A741-3896-4D11-88D3-92CE91CAFB5C) The child partition processor setting is incompatible with parent partition. 	This error is likely caused by a CPU discrepancy between the required CPUs for the gateway and the available CPUs on the host. Ensure that the VM CPU count is supported by the underlying hypervisor. For more information about the requirements for AWS Storage Gateway, see Requirements (p. 11).
You try to start a gateway VM and receive an error message "Failed to create partition: Insufficient resources exist to complete the requested service." 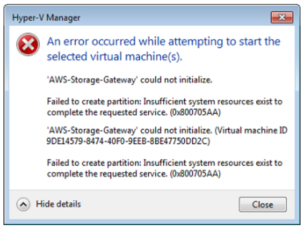 An error occurred while attempting to start the selected virtual machine(s). 'AWS-Storage-Gateway' could not initialize. Failed to create partition: Insufficient system resources exist to complete the requested service. (0x800705AA) 'AWS-Storage-Gateway' could not initialize. (Virtual machine ID 9DE14579-8474-40F0-9EEB-8BE47750DD2C) Failed to create partition: Insufficient system resources exist to complete the requested service. (0x800705AA) 	This error is likely caused by a RAM discrepancy between the required RAM for the gateway and the available RAM on the host. For more information about the requirements for AWS Storage Gateway, see Requirements (p. 11).

Issue	Action to Take
Your snapshots and gateway software updates are occurring at slightly different times than expected.	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. 167).
You need to put the unzipped Microsoft Hyper-V AWS Storage gateway files on the host file system.	Access the host as you do a typical Microsoft Windows server. For example, if the hypervisor host is name `hyperv-server`, then you can use the following UNC path `\\hyperv-server\c$`, which assumes that the name `hyperv-server` can be resolved or is defined in your local hosts file.
You are prompted for credentials when connecting to hypervisor. 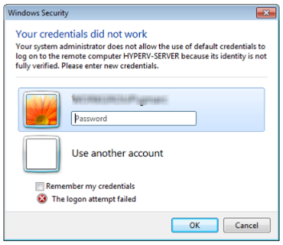 Windows Security Your credentials did not work Your system administrator does not allow the use of default credentials to log on to the remote computer HYPERV-SERVER because its identity is not fully verified. Please enter new credentials. [redacted] Password Use another account ☐ Remember my credentials ☒ The logon attempt failed 	Add your user credentials as a local administrator for the hypervisor host by using the Sconfig.cmd tool. For more information, see Setting Up and Configuring a Microsoft Hyper-V Host (p. 227).

Deploying a AWS Storage Gateway VM on a Microsoft Hyper-V Host

You can choose to run AWS Storage Gateway on-premises either as a VMware ESXi Hypervisor or Microsoft Hyper-V Hypervisor virtual machine (VM) appliance. Following, you can find out how to deploy your gateway on a Microsoft Hyper-V Hypervisor VM appliance. AWS Storage Gateway supports Microsoft Hyper-V Hypervisor versions 2008 R2, 2012, and 2012 R2. For information about supported hosts, see Supported Hypervisors and Host Requirements (p. 16)

Deploying your gateway on Hyper-V Hypervisor is a three-step process. When you complete all the tasks in this section, you can begin the activation process for your gateway.

Topics

Download and Deploy the AWS Storage Gateway VM on Your Host

The AWS Storage Gateway virtual machine (VM) is available as a Hyper-V downloadable package. Following, you can find out how to provision an on-premises Hyper-V host, download the AWS Storage Gateway VM, and deploy a gateway.

Note

For gateways deployed on-premises, restoring a snapshot of a gateway VM that is taken from your hypervisor is not supported.

Topics

- Deploy the AWS Storage Gateway VM to Your Host (p. 237)

Deploy the AWS Storage Gateway VM to Your Host

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

In this exercise, we use Microsoft Hyper-V 2008 R2 Manager to show you how to deploy the VM. Depending on the version of Hyper-V Manager you are using, the screenshots in the exercise might look slightly different from what you see in your UI. However, the procedure is similar for all versions of Microsoft Hyper-V that AWS Storage Gateway supports. For information about the Hyper-V versions that AWS Storage Gateway supports, see Requirements (p. 11).

To connect to the hypervisor host

1. Start the Microsoft Hyper-V Manager on your Windows client, and then in the **Actions** pane, choose **Connect to Server**.

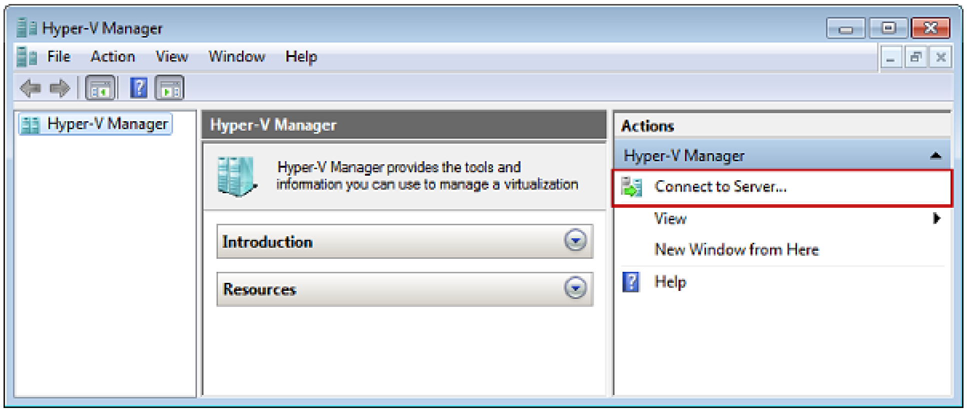

Hyper-V Manager provides the tools and information you can use to manage a virtualization

2. In the **Select Computer** dialog box, choose **Another Computer**, type the IP address or host name, and then choose **OK**.

   Note
   In this exercise, we use `hyperv-server` as a host. Your host name will be different. If you don't find your host name in the **Select Computer** dialog box, you might need to add an entry to your hosts file so that Hyper-V Manager can resolve the server name.

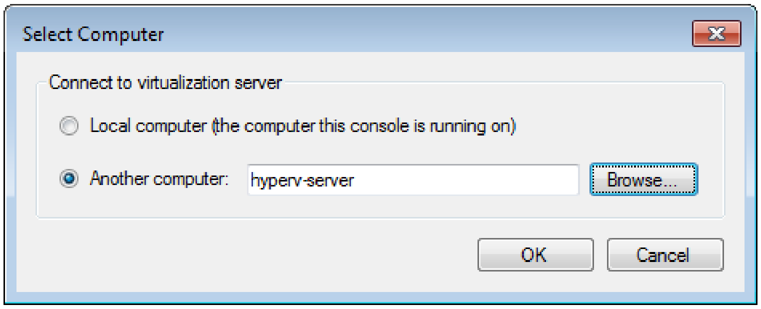

Select Computer

Connect to virtualization server

○ Local computer (the computer this console is running on)

● Another computer: hyperv-server

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

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

To specify a location for the virtual hard disks and VM

1. Create locations on the hypervisor host for the gateway virtual hard disks and VM.

   Note
   For Microsoft Hyper-V 2012, you can find this step in the **Choose Folders to Store Virtual Hard Disks** dialog box.

   a. Navigate to the hypervisor drive.

      For example, assuming the name of the host in this exercise and the C drive as the correct drive for the host, on the **Start** menu you can type `\\hyperv-server\c$`.

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

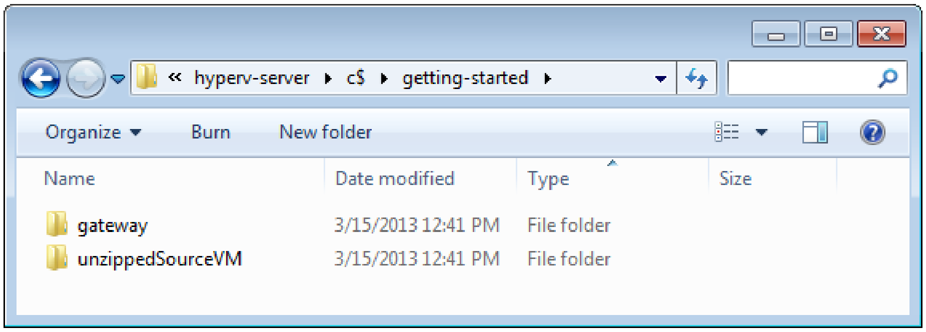

Name	Date modified	Type	Size
gateway	3/15/2013 12:41 PM	File folder	
unzippedSourceVM	3/15/2013 12:41 PM	File folder	

2. Configure the Hyper-V Manager to point to the `gateway` folder you created. The running VM stores its configuration in this folder.

   a. On the **Actions** menu, choose **Hyper-V Settings**.

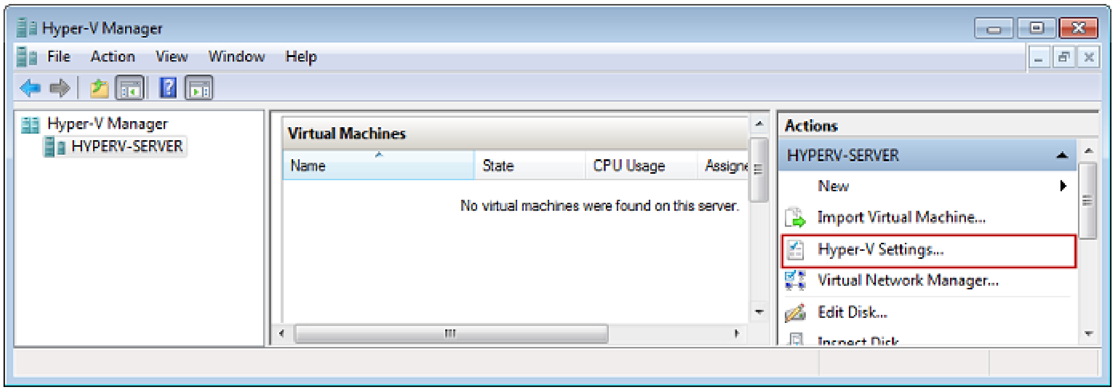

No virtual machines were found on this server.

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

      i. On the left pane, under **Server**, choose **Virtual Hard Disks**.

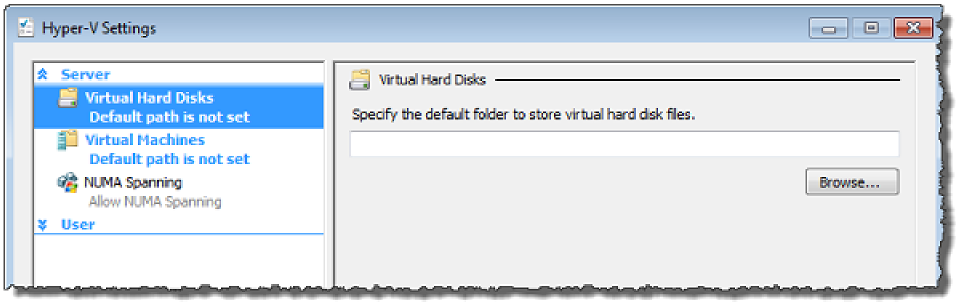

ii. Browse to find the `gateway` folder you created earlier. Note that you are browsing on the hypervisor (host) server.

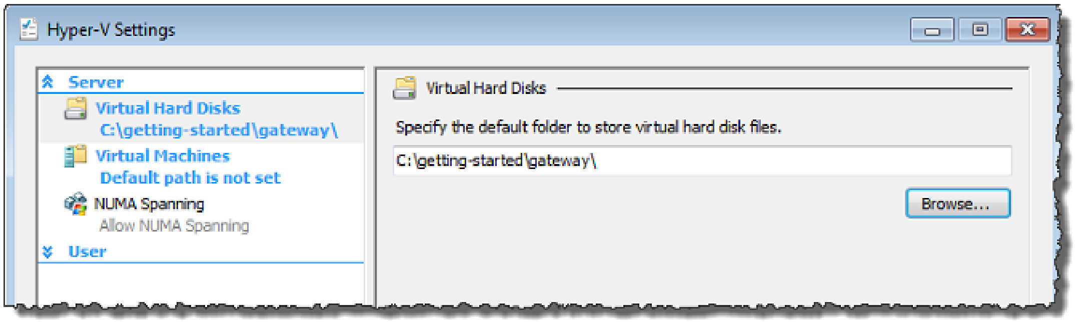

Virtual Hard Disks: C:\getting-started\gateway\

iii. On the left pane, choose **Virtual Machines** under **Server**, and then browse to set the location to the same gateway folder and choose **OK**.

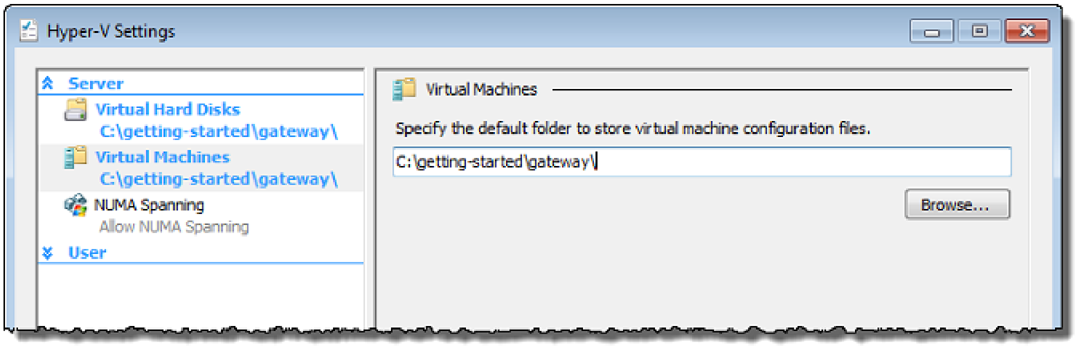

Virtual Hard Disks: C:\getting-started\gateway\
Virtual Machines: C:\getting-started\gateway\

To import the VM

1. Copy the unzipped source VM files to the folder you created on the host computer. In this exercise, the path is `\\hyperv-server\c$\getting-started\unzippedSourceVM\AWS-Storage-Gateway`.
2. Import the AWS Storage Gateway VM to the host.

   a. In the Hyper-V Manager, select the host **`hyperv-server`** in the left pane, which shows the console tree.

   b. On the **Actions** menu, choose **Import Virtual Machine**.

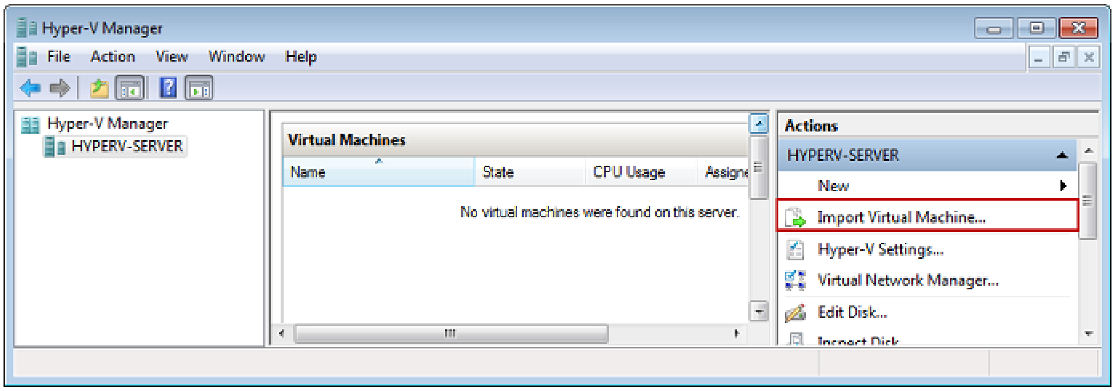

c. In the **Import Virtual Machine** dialog box, do the following.

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

Note
For Microsoft Hyper-V 2012, this step can be found in the **Choose Folders for Virtual Machine Files** dialog box (the next dialog box).

Caution
You must point to the correct folder for the import to succeed. The correct folder to select (`AWS-Storage-Gateway`) contains three other folders (`Snapshots`, `Virtual Hard Disks`, `Virtual Machines`) and one file (`config.xml`). Depending on how you unzip the gateway source files, you might end up with an extra folder level. For help troubleshooting imports, see Troubleshooting Your Microsoft Hyper-V Setup (p. 233).

ii. Choose **Copy the virtual machine (create a new unique ID)**.

Note
For Microsoft Hyper-V 2012, this step can be found in the **Choose Import Type** tab.

iii. Check **Duplicate all files so the same virtual machine can be imported again**.

iv. Choose **Import**.

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

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

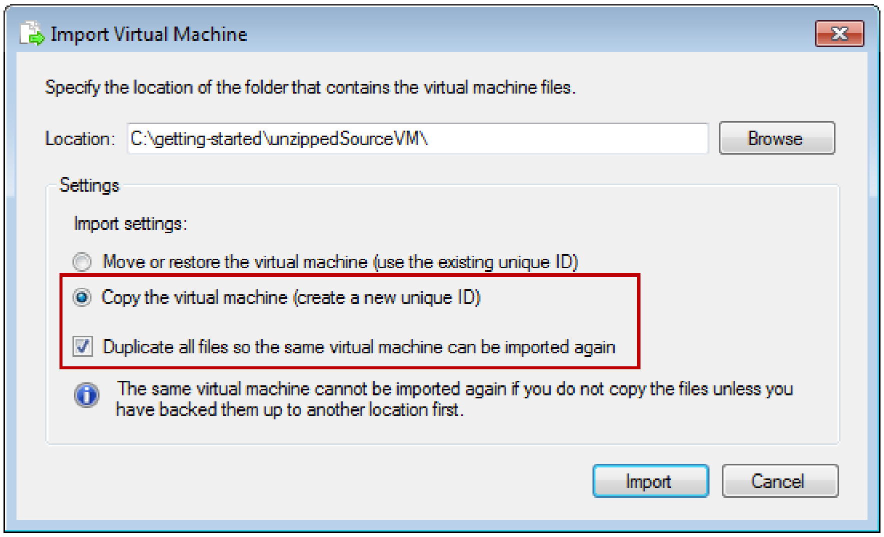

Import Virtual Machine

Specify the location of the folder that contains the virtual machine files.

Location: C:\getting-started\unzippedSourceVM\

Settings

Import settings:

○ Move or restore the virtual machine (use the existing unique ID)

● Copy the virtual machine (create a new unique ID)

☑ Duplicate all files so the same virtual machine can be imported again

The same virtual machine cannot be imported again if you do not copy the files unless you have backed them up to another location first.

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

3. Rename the virtual machine to avoid confusion with other virtual machines that you might import to the host.

   a. Open the context (right-click) menu for the virtual machine name, and then choose **Rename**.

   b. Provide a new name for the virtual machine.

      In this exercise, we will use `ExampleGatewayHyperV`.

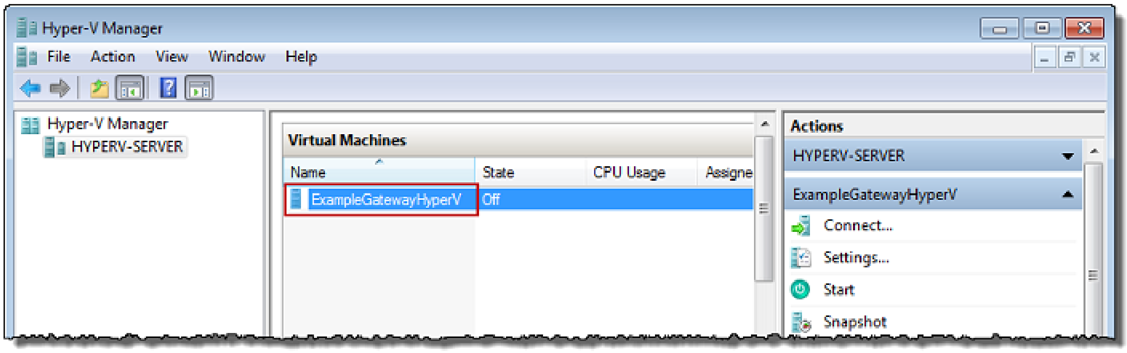

Virtual Machines

Name	State	CPU Usage	Assigne
ExampleGatewayHyperV	Off		

4. Confirm that **Time synchronization** for the VM is selected in **Integration Services**.

   a. For **Virtual Machines**, select the virtual machine **`ExampleGatewayHyperV`**.

   b. On the **Actions** menu, choose **Settings**.

   c. In the **Settings** dialog box, select **Integration services** under **Management** and confirm that **Time synchronization** is checked.

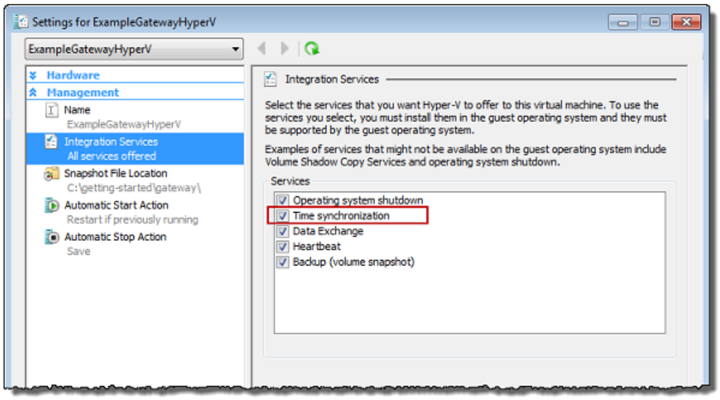

Settings for ExampleGatewayHyperV

ExampleGatewayHyperV

Integration Services

Select the services that you want Hyper-V to offer to this virtual machine. To use the services you select, you must install them in the guest operating system and they must be supported by the guest operating system.

Examples of services that might not be available on the guest operating system include Volume Shadow Copy Services and operating system shutdown.

Services

- ☑ Operating system shutdown
- ☑ Time synchronization
- ☑ Data Exchange
- ☑ Heartbeat
- ☑ Backup (volume snapshot)

5. Configure the host time if you have not already done so, and then choose **OK**.

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

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

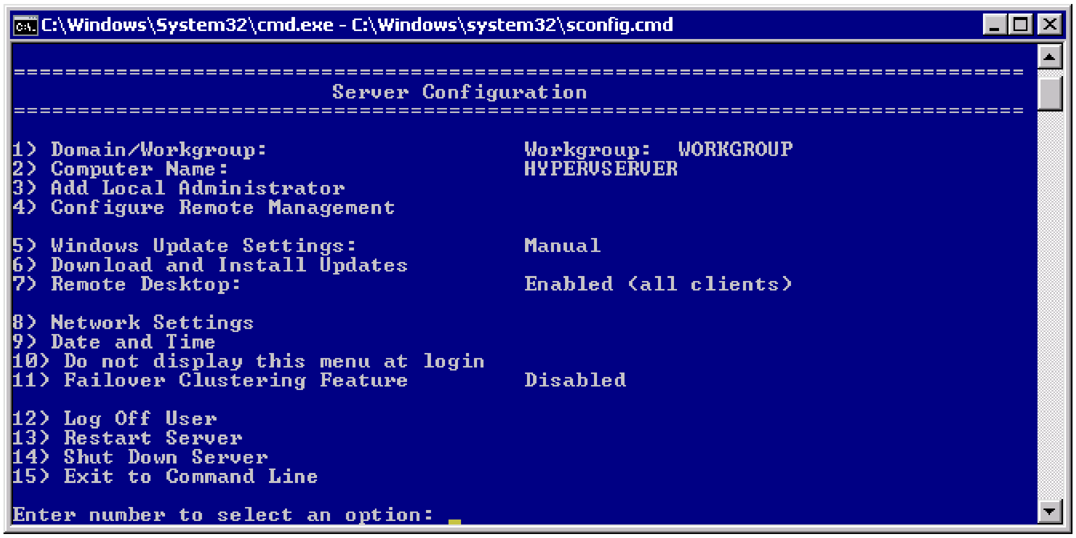

C:\Windows\System32\cmd.exe - C:\Windows\system32\sconfig.cmd

```
===============================================================================
                         Server Configuration
===============================================================================

1) Domain/Workgroup:                        Workgroup:  WORKGROUP
2) Computer Name:                           HYPERVSERVER
3) Add Local Administrator
4) Configure Remote Management

5) Windows Update Settings:                 Manual
6) Download and Install Updates
7) Remote Desktop:                          Enabled (all clients)

8) Network Settings
9) Date and Time
10) Do not display this menu at login
11) Failover Clustering Feature             Disabled

12) Log Off User
13) Restart Server
14) Shut Down Server
15) Exit to Command Line

Enter number to select an option: _
```

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

      The **Date and Time** control panel appears.

   c. Configure the time, and then choose **OK**.

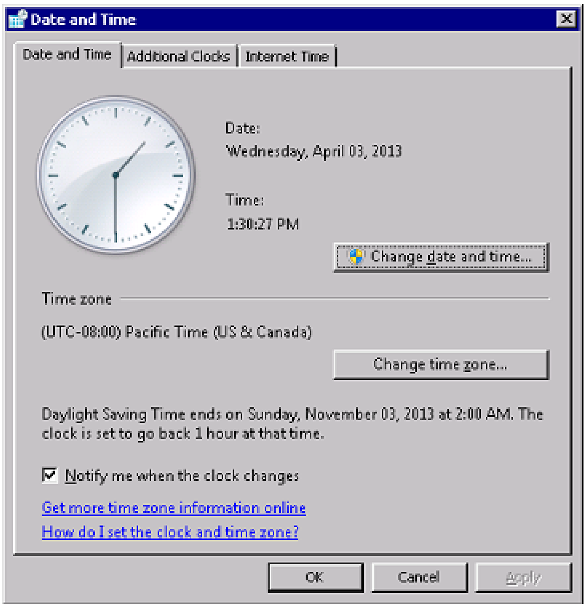

Date and Time

Date and Time | Additional Clocks | Internet Time

Date:
Wednesday, April 03, 2013

Time:
1:30:27 PM

Change date and time...

Time zone

(UTC-08:00) Pacific Time (US & Canada)

Change time zone...

Daylight Saving Time ends on Sunday, November 03, 2013 at 2:00 AM. The clock is set to go back 1 hour at that time.

☑ Notify me when the clock changes

Get more time zone information online
How do I set the clock and time zone?

OK | Cancel | Apply

Now you can log off the remote computer.

To configure a virtual network and use it for the VM

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

   Note
   In this exercise, we assume that the host doesn't have virtual network settings configured. If you already have a virtual network configured, see step 2.

   a. On the **Actions** menu, under the hypervisor host name (for example, `hyperv-server`), choose **Virtual Network Manager**.

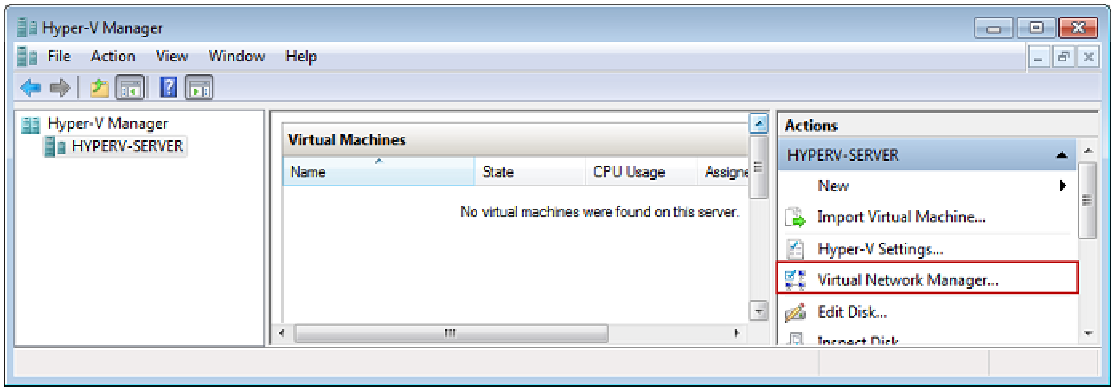

No virtual machines were found on this server.

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

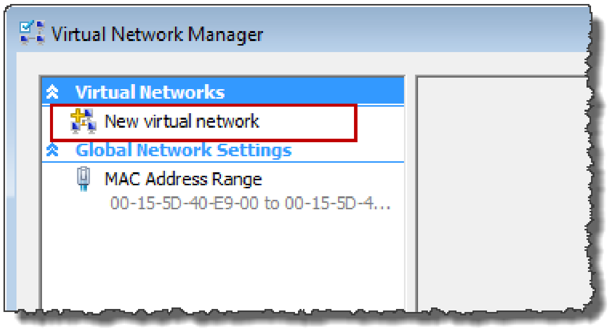

c. Select **External** as the virtual virtual network type, and then choose **Add**.

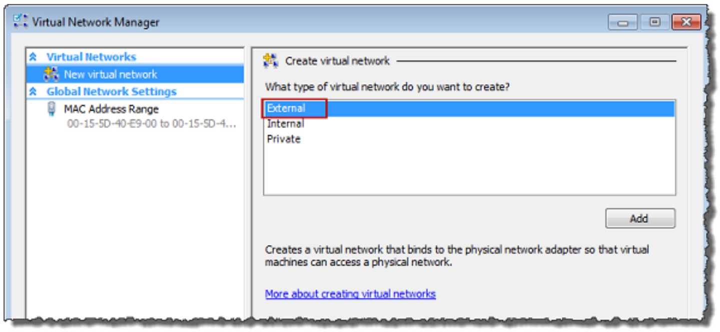

Creates a virtual network that binds to the physical network adapter so that virtual machines can access a physical network.

d. Provide a name for the network, and then choose **OK**.

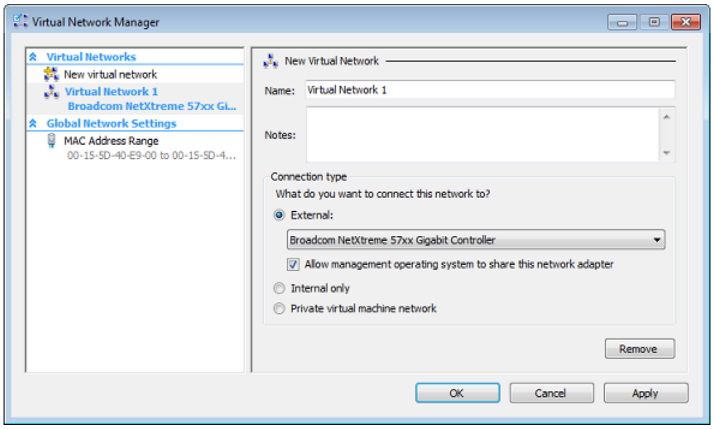

Name: Virtual Network 1

Notes:

● External:

Broadcom NetXtreme 57xx Gigabit Controller

☑ Allow management operating system to share this network adapter

○ Internal only

○ Private virtual machine network

2. Configure the virtual machine to use the virtual network.

a. For **Virtual Machines**, select the virtual machine **`ExampleGatewayHyperV`**.

b. On the **Actions** pane, choose **Settings**.

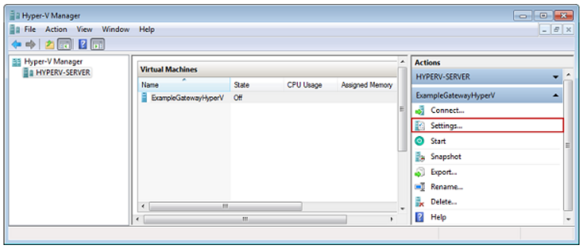

Name	State	CPU Usage	Assigned Memory
ExampleGatewayHyperV	Off		

c. In the **Settings** window, choose **Network Adapter**. The **Network Adapter** should have a status of **Not connected**.

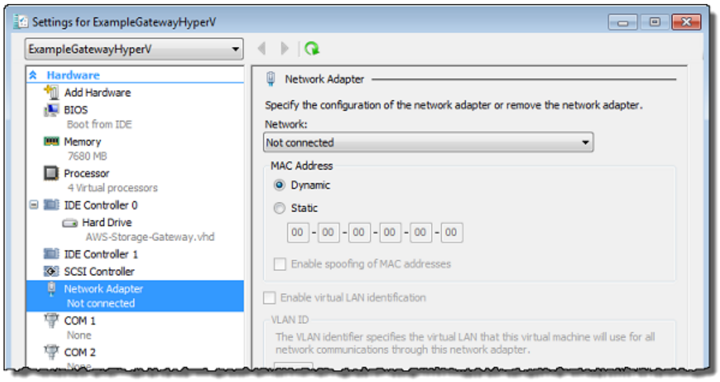

Settings for ExampleGatewayHyperV

ExampleGatewayHyperV

Network Adapter

Specify the configuration of the network adapter or remove the network adapter.

Network: Not connected

MAC Address
● Dynamic
○ Static
00 - 00 - 00 - 00 - 00 - 00
□ Enable spoofing of MAC addresses

□ Enable virtual LAN identification

VLAN ID
The VLAN identifier specifies the virtual LAN that this virtual machine will use for all network communications through this network adapter.

d. In the **Network** box at right, select a network, and then choose **OK**. In the following screenshot, **Virtual Network 1** is selected.

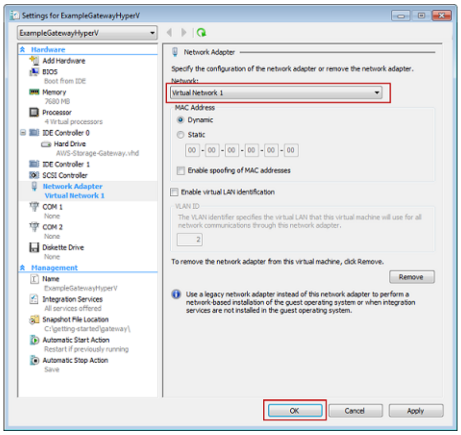

Settings for ExampleGatewayHyperV

ExampleGatewayHyperV

Network Adapter

Specify the configuration of the network adapter or remove the network adapter.

Network: Virtual Network 1

MAC Address
● Dynamic
○ Static
00 - 00 - 00 - 00 - 00 - 00
□ Enable spoofing of MAC addresses

□ Enable virtual LAN identification

VLAN ID
The VLAN identifier specifies the virtual LAN that this virtual machine will use for all network communications through this network adapter.
2

To remove the network adapter from this virtual machine, click Remove.

Use a legacy network adapter instead of this network adapter to perform a network-based installation of the guest operating system or when integration services are not installed in the guest operating system.

Next Step

Provision Local Storage for the AWS Storage Gateway VM (Hyper-V) (p. 245)

Provision Local Storage for the AWS Storage Gateway VM (Hyper-V)

Next, you allocate local storage for your deployed gateway VM. After you activate your gateway, you configure this local storage for the gateway's use.

Topics

Decide the Amount of Local Disk Storage

In this step, you decide the number and size of disks to allocate for your gateway, as follows:

- All gateways require one or more disks for an upload buffer.

- Depending on the storage solution you deploy (see Plan Your AWS Storage Gateway Deployment (p. 8)), the gateway requires the following additional storage:
  - Tape Gateways and gateways created with cached volumes require one or more disks for cache storage.
  - Stored volumes require one or more disks for volume storage—one local disk for each volume created on the gateway.

You will need at least two disks to begin with. The following table recommends sizes for the local disk storage you will allocate in this step.

Disk Size	Gateway-Cached	Gateway-Stored	Tape Gateway
Upload buffer (minimum)	150 GiB	150 GiB	150 GiB
Upload buffer (maximum)	2 TiB	2 TiB	2 TiB
Cache storage (minimum)	1.1 x upload buffer	—	1.1 x upload buffer
Cache storage (maximum)	16 TiB	—	16 TiB
Volume storage (minimum)	—	1GiB	—
Volume storage (maximum)	32 TiB	16 TiB	—

Before going to the next step, decide the number and size of disks to allocate. You can add more local storage after you set the gateway up, as your workload demands. If you're not sure what number and size of disks to use, provision two local disks with minimum recommended size.

If you plan to deploy your gateway in production, you should consider your real workload in determining disk sizes. For information about disk size guidelines, see Adding and Removing Upload Buffer (p. 144) and Adding Cache Storage (p. 146).

For more information about how gateways use local storage, see How AWS Storage Gateway Works (Architecture) (p. 3).

In the next step, you allocate the local disk storage to the gateway VM you deployed.

Provision Local Disk Storage

You can add disks from direct attached storage (DAS) or from a storage area network (SAN). This section tells how to add a virtual disk from a DAS. For instructions on attaching iSCSI volumes from an existing SAN for this step, see To add a new iSCSI target (p. 209).

Repeat the following procedure for each local disk you planned in the preceding step.

To allocate a local disk from DAS

1. Start the Microsoft Hyper-V Manager, and then connect to the hypervisor.
2. For **Virtual Machines**, select the virtual machine **`ExampleGatewayHyperV`**.
3. On the **Actions** pane of your gateway, select **Settings**.

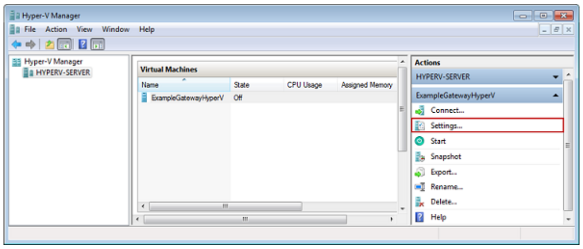

4. In the **Settings** window, select **SCSI Controller**, and then choose **Add**.

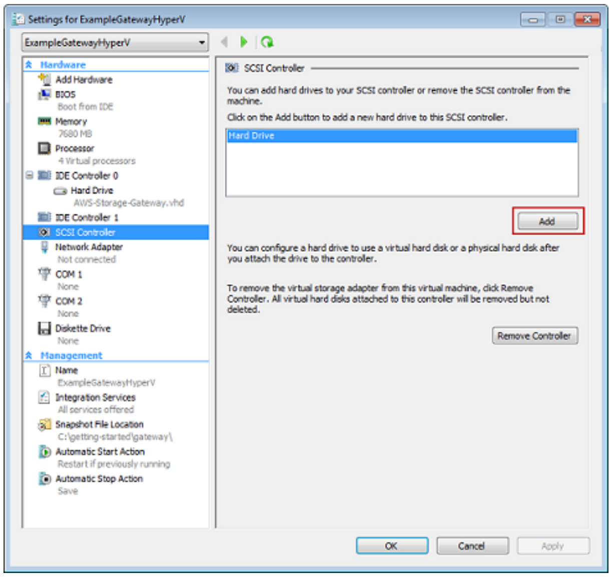

You can add hard drives to your SCSI controller or remove the SCSI controller from the machine.

Click on the Add button to add a new hard drive to this SCSI controller.

You can configure a hard drive to use a virtual hard disk or a physical hard disk after you attach the drive to the controller.

To remove the virtual storage adapter from this virtual machine, click Remove Controller. All virtual hard disks attached to this controller will be removed but not deleted.

5. On the **Hard Drive** pane, under **Media**, choose **New**.

   Note

   In this example, notice that **virtual hard disk (.vhd) file** is selected. AWS Storage Gateway supports .vhdx file type. This file type enables you to create larger virtual disks than other file types. If you create a .vhdx type virtual disk, make sure that the size of the virtual disks you create does not exceed the recommended disk size for your gateway. For more information, see Decide the Amount of Local Disk Storage (p. 245).

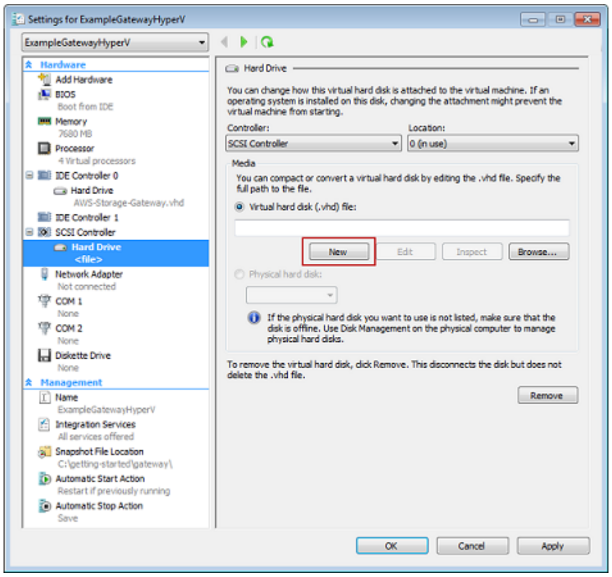

You can change how this virtual hard disk is attached to the virtual machine. If an operating system is installed on this disk, changing the attachment might prevent the virtual machine from starting.

Controller: SCSI Controller

Location: 0 (in use)

You can compact or convert a virtual hard disk by editing the .vhd file. Specify the full path to the file.

● Virtual hard disk (.vhd) file:

○ Physical hard disk:

If the physical hard disk you want to use is not listed, make sure that the disk is offline. Use Disk Management on the physical computer to manage physical hard disks.

To remove the virtual hard disk, click Remove. This disconnects the disk but does not delete the .vhd file.

6. In the **New Virtual Hard Disk Wizard**, create a new virtual hard disk.

   a. On the **Before You Begin** page, choose **Next**.

      Note
      If you are using the Microsoft Hyper-V 2012 Hypervisor, you will be prompted to choose a disk format (**VHD** or **VHDX**).

   b. On the **Choose Disk Type** page, choose **Fixed size**, and then choose **Next**.

      When you use fixed-size provisioning, the disk storage is allocated immediately, resulting in better performance. If you don't use fixed-size provisioning, the storage is allocated on demand. On-demand allocation can affect the functioning of AWS Storage Gateway. For AWS Storage Gateway to function properly, the VM disks must be stored in fixed-size provisioned format.

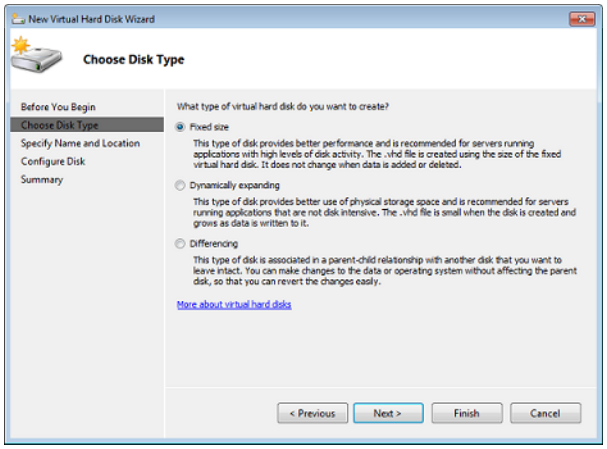

What type of virtual hard disk do you want to create?

● Fixed size

This type of disk provides better performance and is recommended for servers running applications with high levels of disk activity. The .vhd file is created using the size of the fixed virtual hard disk. It does not change when data is added or deleted.

○ Dynamically expanding

This type of disk provides better use of physical storage space and is recommended for servers running applications that are not disk intensive. The .vhd file is small when the disk is created and grows as data is written to it.

○ Differencing

This type of disk is associated in a parent-child relationship with another disk that you want to leave intact. You can make changes to the data or operating system without affecting the parent disk, so that you can revert the changes easily.

   c. On the **Specify Name and Location** page, specify a name and location for the virtual hard disk.

      i. Specify `UploadBuffer.vhd` as the name.

      ii. Specify the location as `c$\getting-started\gateway`.

Note

In this example setup, you store the virtual disk with the virtual machine. For real-world workloads, we strongly recommend that you not provision local disks using the same underlying physical storage disk. Depending on your hosting environment, performance, and portability requirements, consider selecting a different physical disk in this step.

iii. Choose **Next**.

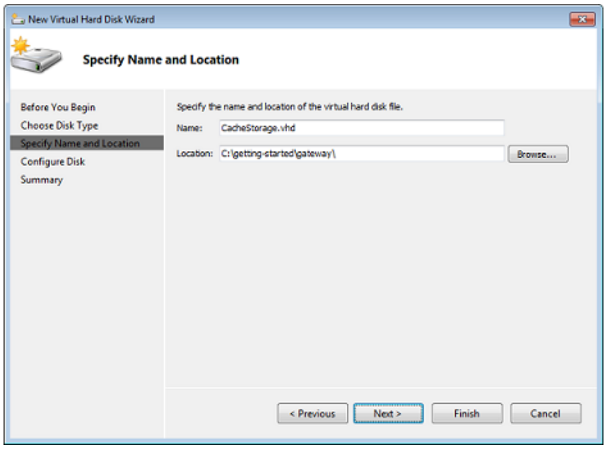

Specify the name and location of the virtual hard disk file.

Name: CacheStorage.vhd

Location: C:\getting-started\gateway\

d. In the **Configure Disk** page, specify the disk size you previously decided, and choose **Finish**.

Note

The disk sizes used in these examples in this documentation are not suitable for real-world workloads. If your setup is a trial setup, you can use small disk sizes.

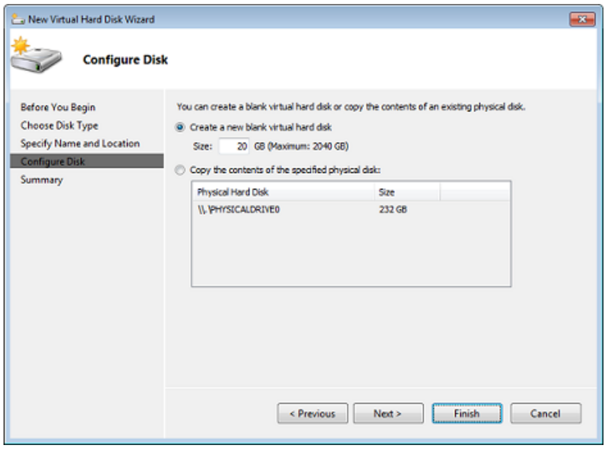

You can create a blank virtual hard disk or copy the contents of an existing physical disk.

● Create a new blank virtual hard disk

Size: 20 GB (Maximum: 2040 GB)

○ Copy the contents of the specified physical disk:

Physical Hard Disk	Size
\\.\PHYSICALDRIVE0	232 GB

e. After the virtual disk is created, verify that **Hard Drive** shows up under **SCSI controller**.

f. Choose **SCSI Controller** to prepare to add another hard drive, and then choose **OK**.

Warning

When you add an additional hard drive, you must first choose **SCSI Controller** and then follow the steps in the following procedure. Choosing **New** when viewing the details of an existing hard drive replaces the existing drive.

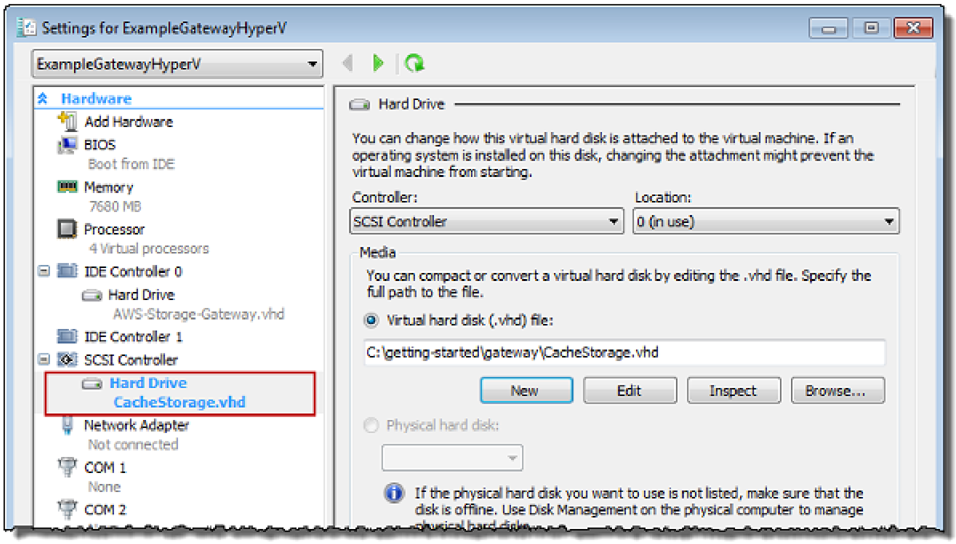

You can change how this virtual hard disk is attached to the virtual machine. If an operating system is installed on this disk, changing the attachment might prevent the virtual machine from starting.

Controller: SCSI Controller

Location: 0 (in use)

You can compact or convert a virtual hard disk by editing the .vhd file. Specify the full path to the file.

● Virtual hard disk (.vhd) file: C:\getting-started\gateway\CacheStorage.vhd

If the physical hard disk you want to use is not listed, make sure that the disk is offline. Use Disk Management on the physical computer to manage physical hard disks.

7. Repeat steps 1 through 6 for each disk to allocate.
8. Verify your gateway VM has the disks you added.

   a. Start the Microsoft Hyper-V Manager and connect to the hypervisor if it is not already started.
   b. For **Virtual Machines**, select the virtual machine **`ExampleGatewayHyperV`**.
   c. On the **Actions** pane, choose **Settings**, and then on the **Hardware** pane select **SCSI Controller** and verify there are two disks under the SCSI controller.

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

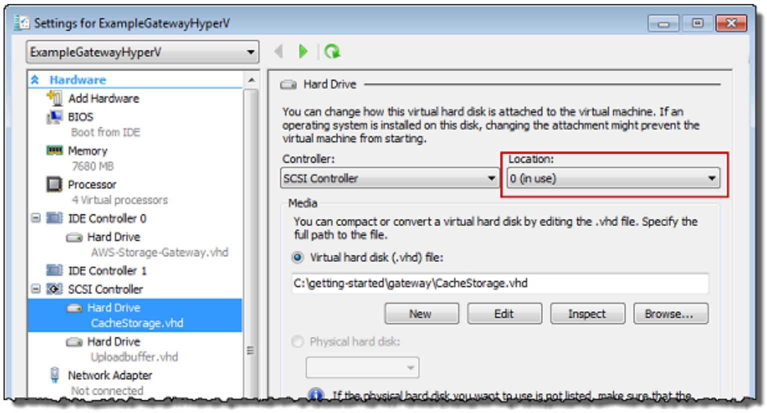

You can change how this virtual hard disk is attached to the virtual machine. If an operating system is installed on this disk, changing the attachment might prevent the virtual machine from starting.

Controller: SCSI Controller

Location: 0 (in use)

You can compact or convert a virtual hard disk by editing the .vhd file. Specify the full path to the file.

● Virtual hard disk (.vhd) file: C:\getting-started\gateway\CacheStorage.vhd

If the physical hard disk you want to use is not listed, make sure that the

Next Step

Activating Your Gateway (p. 55)

Volume Gateway

Topics

- Adding and Removing Disks for Your Gateway (p. 251)
- Working with Snapshots (p. 253)
- Working With Volumes (p. 265)
- Adding and Removing Amazon EBS Volumes for Your Gateway Hosted on Amazon EC2 (p. 271)

Adding and Removing Disks for Your Gateway

You can add or remove underlying disks from your gateway as described following. For example, you might add disks to your gateway to use as an upload buffer or cache storage if you need additional upload buffer space or cache storage. You can also remove the underlying disks from your gateway. For example, you might want to remove a failed disk from your gateway.

> **Important**
> Do not remove a disk allocated for cache storage.

For information about how to add a disk to a gateway hosted on VMware ESXi, see Provisioning Local Disk Storage for the Gateway VM (VMWare) (p. 218).

For information about how to add a disk to a gateway hosted on Microsoft Hyper-V, see Provision Local Storage for the AWS Storage Gateway VM (Hyper-V) (p. 245).

Topics

- Remove a Disk from a Gateway Hosted on VMware ESXi (p. 251)
- Remove a Disk from Gateway Hosted on Microsoft Hyper-V (p. 252)

Remove 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.

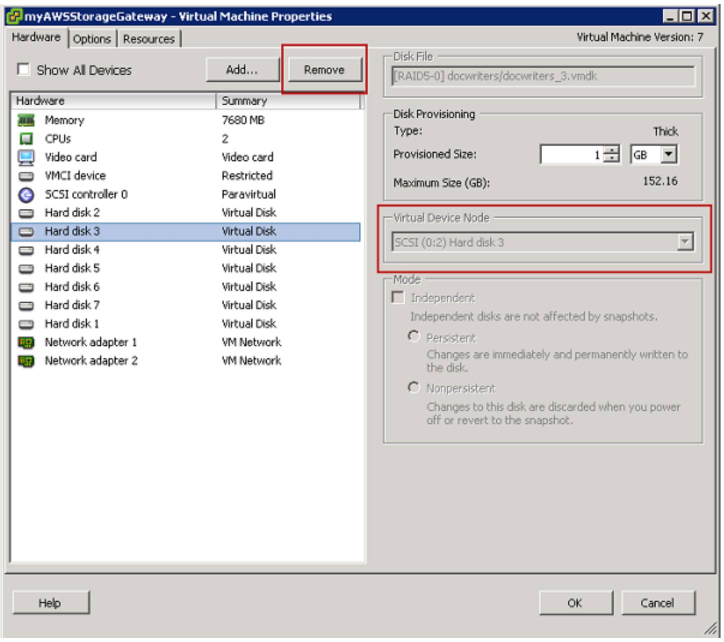
myAWSStorageGateway - Virtual Machine Properties
Hardware
Options
Resources
Virtual Machine Version: 7
Add...
Remove
Help
OK
Cancel

☐ Show All Devices

Hardware	Summary
Memory	7680 MB
CPUs	2
Video card	Video card
VMCI device	Restricted
SCSI controller 0	Paravirtual
Hard disk 2	Virtual Disk
Hard disk 3	Virtual Disk
Hard disk 4	Virtual Disk
Hard disk 5	Virtual Disk
Hard disk 6	Virtual Disk
Hard disk 7	Virtual Disk
Hard disk 1	Virtual Disk
Network adapter 1	VM Network
Network adapter 2	VM Network

Disk File
[RAID5-0] docwriters/docwriters_3.vmdk

Disk Provisioning
Type: Thick
Provisioned Size: 1 GB
Maximum Size (GB): 152.16

Virtual Device Node
SCSI (0:2) Hard disk 3

Mode
☐ Independent
Independent disks are not affected by snapshots.
○ Persistent
Changes are immediately and permanently written to the disk.
○ Nonpersistent
Changes to this disk are discarded when you power off or revert to the snapshot.

3. Choose an option in the **Removal Options** panel, and then choose **OK** to complete the process of removing the disk.

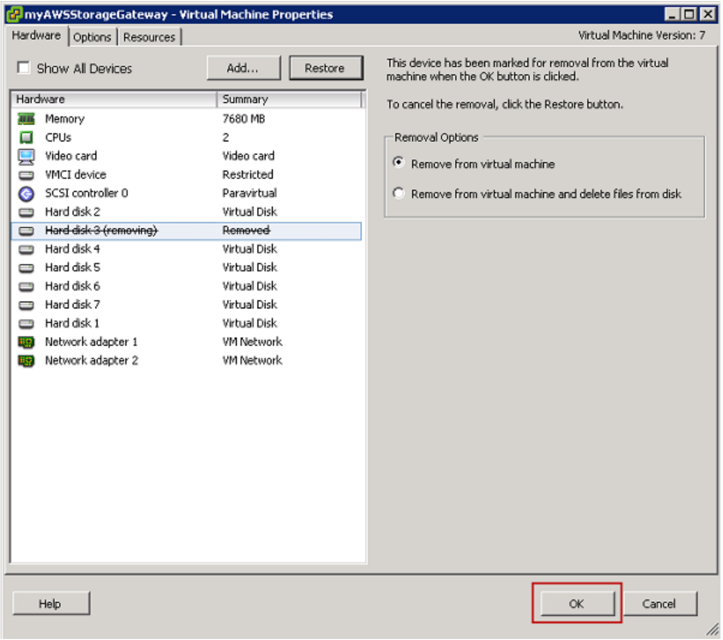
myAWSStorageGateway - Virtual Machine Properties
Hardware
Options
Resources
Virtual Machine Version: 7
Add...
Restore
Help
OK
Cancel

☐ Show All Devices

Hardware	Summary
Memory	7680 MB
CPUs	2
Video card	Video card
VMCI device	Restricted
SCSI controller 0	Paravirtual
Hard disk 2	Virtual Disk
~~Hard disk 3 (removing)~~	~~Removed~~
Hard disk 4	Virtual Disk
Hard disk 5	Virtual Disk
Hard disk 6	Virtual Disk
Hard disk 7	Virtual Disk
Hard disk 1	Virtual Disk
Network adapter 1	VM Network
Network adapter 2	VM Network

This device has been marked for removal from the virtual machine when the OK button is clicked.

To cancel the removal, click the Restore button.

Removal Options
● Remove from virtual machine
○ Remove from virtual machine and delete files from disk

Remove a Disk from 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.

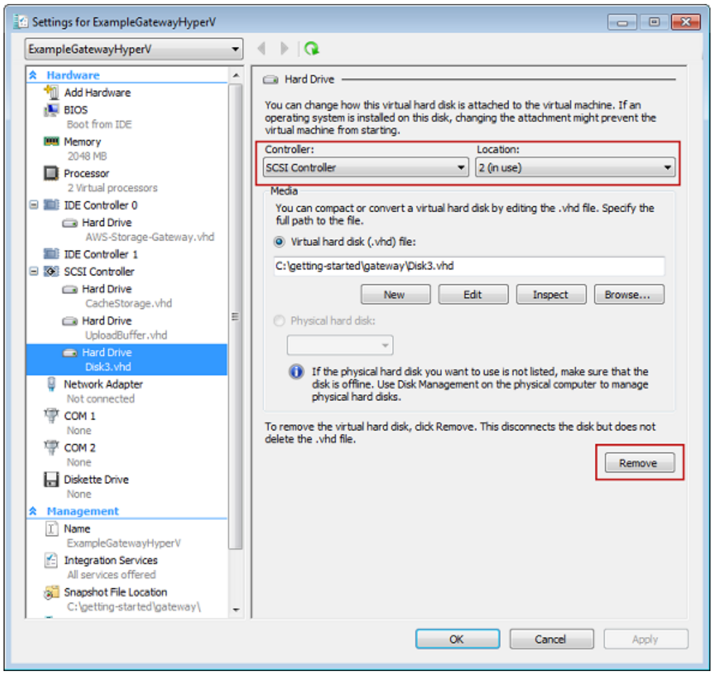

Settings for ExampleGatewayHyperV

Hard Drive

You can change how this virtual hard disk is attached to the virtual machine. If an operating system is installed on this disk, changing the attachment might prevent the virtual machine from starting.

Controller: SCSI Controller

Location: 2 (in use)

Media

You can compact or convert a virtual hard disk by editing the .vhd file. Specify the full path to the file.

● Virtual hard disk (.vhd) file: C:\getting-started\gateway\Disk3.vhd

○ Physical hard disk:

If the physical hard disk you want to use is not listed, make sure that the disk is offline. Use Disk Management on the physical computer to manage physical hard disks.

To remove the virtual hard disk, click Remove. This disconnects the disk but does not delete the .vhd file.

3. Choose **OK** to apply the change.

Working with Snapshots

Topics

- Common Snapshot Tasks (p. 254)
- Snapshot Data Consistency (p. 255)
- Finding and Restoring Snapshots (p. 255)

Using AWS Storage Gateway, you can back up point-in-time snapshots of your data to Amazon Simple Storage Service (Amazon S3) for durable recovery. You can use the snapshot backups later on-premises or in Amazon Elastic Compute Cloud (Amazon EC2), and you can take snapshots on a one-time or scheduled basis. Following, you can find information about the most common tasks that you can perform with snapshots, including creating a snapshot and restoring the snapshot to a volume that can then be mounted as an iSCSI device. When you restore a snapshot to an Amazon Elastic Block Store (Amazon EBS) volume, the volume can then be attached to an Amazon EC2 instance.

AWS Storage Gateway continually and asynchronously uploads data to AWS to keep your local data synchronized with a copy stored in AWS. A benefit of this is that when snapshots are initiated, some or all the data has already been uploaded and snapshots complete quickly. Furthermore, snapshots are incremental—that is, the gateway uploads only the blocks of your volume that have changed since the last snapshot. For example, if you have 100 GiB of data and only 5 GiB data changed since the last snapshot, then the gateway uploads only the 5 GiB of changed data. You can delete any snapshot. AWS Storage Gateway removes only the snapshot data that is not needed by other snapshots, letting you restore a volume from any of the active snapshots.

How snapshots can be effectively used in your AWS Storage Gateway setup depends on the type of gateway you set up—that is, gateway-cached or gateway-stored. For more information on gateway-cached and gateway-stored architecture, see How AWS Storage Gateway Works (Architecture) (p. 3).

- For cached volumes, your volume data is already stored in Amazon S3. You can use snapshots to preserve older versions of your data.
- For stored volumes, your volume data is stored on-premises. You can use snapshots for durable, off-site backups in Amazon S3.

You can restore a snapshot either to a gateway volume or to an Amazon EBS volume. To restore a snapshot to a gateway storage volume, first set up and activate a local gateway. Creating a Tape Gateway (p. 48) walks you through the steps of setting up a gateway. When you restore a snapshot to an Amazon EBS volume, you can then attach the EBS volume to an Amazon EC2 instance. For a complete list of charges and specific prices for Amazon EC2, see the Amazon EC2 Pricing page.

Common Snapshot Tasks

Because snapshots are key to using the AWS Storage Gateway service, you should understand what each snapshot operation does and why it does that. Each task is covered in detail following.

To perform snapshot tasks, you should have one or more gateways that have been running for enough time that there are snapshots to work with. You can work with snapshots using the AWS Storage Gateway console, one of the AWS Software Development Kits (SDKs), or the AWS Storage Gateway REST API; for more information on using the AWS Storage Gateway REST API, see Operations in AWS Storage Gateway (p. 379). Following, we primarily show how to work with the console to perform gateway tasks.

Snapshot Action	Common Scenarios
Finding	You might want to find a snapshot to see if it is complete, what time it was taken, what the size of the snapshot is, or the name of the volume the snapshot was taken from. For more information, see Finding a Snapshot (p. 255).
Scheduling	When you first set up a stored volume, a default snapshot schedule of once a day is set. You can change the frequency and timing of the snapshot schedule to fit your application needs. For more information, see Editing a Snapshot Schedule (p. 104). **Note** For stored volumes, you cannot delete the default snapshot schedule.
Creating	Snapshots for stored volumes are automatically created on a schedule by default, and you can change the schedule of the snapshots as needed. You can also take an instantaneous snapshot at any time for both stored and cached volumes. For more information, see Creating a One-Time Snapshot (p. 104). For information about troubleshooting snapshot issues, see Troubleshooting Snapshot Issues (p. 200).

Snapshot Action	Common Scenarios
Restoring	You can restore a snapshot to a new local volume, or you can use the snapshot to create an Amazon EBS volume and attach that to an Amazon EC2 instance. For more information, see Restoring a Snapshot to a Storage Volume (p. 262) and Restoring a Snapshot to an Amazon EBS Volume (p. 265).
Deleting	If you don't need a snapshot anymore, you can delete it. Because snapshots are incremental backups, if you delete a snapshot only the data that is not needed in other snapshots is deleted. For more information, see Deleting a Snapshot (p. 105).

Snapshot Data Consistency

Snapshots provide a point-in-time view of data that has been written to your AWS Storage Gateway volumes. However, snapshots only capture data that has been written to your storage volumes, which can exclude data that has been buffered by your client application or operating system. Your application and operating system will eventually flush this buffered data to your storage volumes. If you need to guarantee that your application data is flushed to disk prior to taking a snapshot, you should consult your specific application's documentation to understand if and how your application buffers data and how to flush this data.

If you need to guarantee that your operating system and file system have flushed their buffered data to disk prior to taking a snapshot, you can do this by taking your storage volume offline before taking a snapshot. Doing this forces your operating system to flush its data to disk. After the snapshot is complete, you can bring the volume back online. In Microsoft Windows, use Disk Management (`diskmgmt.msc`) to choose the storage volume and take it offline or bring it back online. To script this process in Windows, you can use a command line tool such as Diskpart.exe. In Linux, you can use the `mount` and `umount` commands.

Finding and Restoring Snapshots

Following, you can find information about how to find and restore snapshots

Topics

- Finding a Snapshot (p. 255)
- Restoring a Snapshot (p. 261)

Finding a Snapshot

When you want to restore a snapshot to a new volume—for example, in a disaster recovery scenario—or to restore a previous version of your application data, you need to find a snapshot associated with the volume in question. To locate a snapshot, you need to know the snapshot ID. You can find a snapshot ID in several ways, including by using the Amazon EBS console or Amazon EC2 console, or programmatically by using one of the AWS Software Development Kits (SDKs).

Topics

- Finding Snapshots by Using the AWS SDK for Java (p. 256)
- Finding Snapshots by Using the AWS SDK for .NET (p. 258)
- Finding Snapshots by Using the AWS Tools for Windows PowerShell (p. 260)

If you list snapshots using the Amazon EBS console, the list includes all snapshots generated from your gateway and also snapshots that you might have generated from Amazon EBS volumes. If you list snapshots on the Amazon EC2 console, you can see more snapshot properties to help you find your snapshot. The Amazon EC2 console also has a search filtering capability. Both console experiences are described following.

In some scenarios, you might need to search using several snapshot properties at once—for example, status, start date, and description. In this case, you can use a programmatic approach. For examples, see Finding Snapshots by Using the AWS SDK for Java (p. 256), Finding Snapshots by Using the AWS SDK for .NET (p. 258), and Finding Snapshots by Using the AWS Tools for Windows PowerShell (p. 260).

When you find your snapshot, you can view its details, including the date and time the snapshot was started and the storage volume on your gateway that was the source for the snapshot.

To find a snapshot for a volume using the Amazon EBS console

1. On the AWS Storage Gateway console, choose **Volumes** in the navigation pane.
2. Choose the volume that the snapshot was created from, and then choose the link in the **Snapshot** column. You are redirected to the Amazon EBS console where you can search for your snapshot.

   The link shows the number of snapshots created from the volume.

Finding Snapshots by Using the AWS SDK for Java

You can use a programmatic approach to quickly find snapshots and filter the results returned using snapshot properties such as snapshot status, description, and the date the snapshot was initiated. The following example demonstrates how to find snapshots using the AWS SDK for Java using several snapshot properties. To use the example code, you should be familiar with running a Java console application. For more information, see Getting Started in the *AWS SDK for Java Developer Guide*.

Example : Finding Snapshots by Using the AWS SDK for Java

The following Java code example finds snapshots for a specified volume of a gateway using several properties of the snapshot to filter the results returned. It uses the AWS SDK for Java and the Amazon EC2 API. The Amazon EC2 API includes operations for working with snapshots.

```
import com.amazonaws.services.ec2.model.DescribeSnapshotsResult;
import com.amazonaws.services.ec2.model.Filter;
import com.amazonaws.services.ec2.model.Snapshot;

public class FindingSnapshotsExample {

    static AmazonEC2Client ec2Client;
    // A full volume id or partial fragment with "*".
    static String volumeID = "vol-424*";
    // Snapshot status to filter on: "completed", "pending", "error".
    static String status = "completed";
    // The number of days before which to not return snapshot results.
    static int daysBack = 10;
    // Service end point. Should be same region as volume/gateway.
    public static String serviceURLEC2 = "https://ec2.us-
east-1.amazonaws.com";

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

        ec2Client = new AmazonEC2Client(new PropertiesCredentials(

 FindingSnapshotsExample.class.getResourceAsStream("AwsCredentials.properties")));
        ec2Client.setEndpoint(serviceURLEC2);

        FindingSnapshotsForAVolume();

    }

    private static void FindingSnapshotsForAVolume() {

        try {
            Filter[] filters = new Filter[2];
            filters[0] = new Filter().withName("volume-
id").withValues(volumeID);
            filters[1] = new Filter().withName("status").withValues(status);

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

            List<Snapshot> snapshots = describeSnapshotResult.getSnapshots();
            System.out.println("volume-id = " + volumeID);
            for (Snapshot s : snapshots) {
                if (CompareDates(daysBack, s.getStartTime())) {
                    StringBuilder sb = new StringBuilder();
                    sb.append(s.getSnapshotId() + ", " + s.getStartTime() +
 ", " + s.getDescription());
                    System.out.println(sb.toString());
                }
            }
        } catch (AmazonClientException ace) {
            System.err.println(ace.getMessage());
        }
    }
    public static boolean CompareDates(int daysBack, Date snapshotDate) {
        Date today = new Date();
        Calendar cal = new GregorianCalendar();
        cal.setTime(today);
        cal.add(Calendar.DAY_OF_MONTH, -daysBack);
        Date cutoffDate = cal.getTime();
        return (snapshotDate.compareTo(cutoffDate) > 0) ? true : false;
    }

}
```

Finding Snapshots by Using the AWS SDK for .NET

You can use a programmatic approach to quickly find snapshots and filter the results returned using snapshot properties such as snapshot status, description, and the date the snapshot was initiated. The following example demonstrates how to find snapshots using the AWS SDK for .NET using several snapshot properties. To use the example code, you should be familiar with running a .NET console application. For more information, see Getting Started in the *AWS SDK for .NET Developer Guide*.

The following C# code example finds snapshots for a specified volume of a gateway using several properties of the snapshot to filter the results returned. It uses the AWS SDK for .NET and the Amazon EC2 API. The Amazon EC2 API includes operations for working with snapshots.

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

```
        static String status = "completed";
        // The number of days before which to not return snapshot results.
        static int daysBack = 4;
        // Service endpoint. Should be same region as volume/gateway.
        static String serviceURLEC2 = "https://ec2.us-east-1.amazonaws.com";


        public static void Main(string[] args)
        {
            //Create a ec2 client
            ec2Config = new AmazonEC2Config();
            ec2Config.ServiceURL = serviceURLEC2;
            ec2Client = new AmazonEC2Client(ec2Config);

            FindingSnapshotsForAVolume();

            Console.WriteLine("\nTo continue, press Enter.");
            Console.Read();
        }

        private static void FindingSnapshotsForAVolume()
        {
            try
            {
                Filter[] filters = new Filter[2];
                filters[0] = new Filter().WithName("volume-
id").WithValue(volumeID);
                filters[1] = new
 Filter().WithName("status").WithValue(status);

                DescribeSnapshotsRequest describeSnapshotsRequest =
                    new DescribeSnapshotsRequest().WithFilter(filters);
                DescribeSnapshotsResponse describeSnapshotsResponse =
                    ec2Client.DescribeSnapshots(describeSnapshotsRequest);

                List<Snapshot> snapshots =
 describeSnapshotsResponse.DescribeSnapshotsResult.Snapshot;
                Console.WriteLine("volume-id = " + volumeID);
                foreach (Snapshot s in snapshots)
                {
                    if (CompareDates(daysBack, s.StartTime))
                    {
                        StringBuilder sb = new StringBuilder();
                        sb.Append(s.SnapshotId + ", " + s.StartTime + ", " +
 s.Description);
                        Console.WriteLine(sb.ToString());
                    }
                }
            }
            catch (AmazonEC2Exception ex)
            {
                Console.WriteLine(ex.Message);
            }

        }

        public static Boolean CompareDates(int daysBack, String d)
        {
            DateTime snapshotDate = DateTime.Parse(d);
            DateTime cutoffDate = DateTime.Now.Add(new TimeSpan(-daysBack, 0,
 0, 0));
            return (DateTime.Compare(snapshotDate, cutoffDate) < 0) ? true :
 false;
        }
    }
}
```

Finding Snapshots by Using the AWS Tools for Windows PowerShell

You can use a programmatic approach to quickly find snapshots and filter the results returned using snapshot properties such as snapshot status, description, and the date the snapshot was initiated. The following example demonstrates how to find snapshots using the AWS Tools for Windows PowerShell using several snapshot properties. To use the example code, you should be familiar with running a PowerShell script. For more information, see Getting Started in the *AWS Tools for Windows PowerShell User Guide*.

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

The following PowerShell script example finds snapshots for a specified volume of a gateway using several properties of the snapshot to filter the results returned. It uses AWS Tools for Windows PowerShell cmdlets for Amazon EC2. The Amazon EC2 cmdlets include operations for working with snapshots.

You need to update the script and provide a full or partial volume ID, a snapshot status, and a number of days to indicate a cutoff date for the snapshots to return.

```
<#
.DESCRIPTION
    Finds snapshots for a given volume and criteria about the snapshot.

.NOTES
    PREREQUISITES:
    1) AWS Tools for PowerShell from http://console.aws.amazon.com/
powershell/
    2) Credentials and region stored in session using Initialize-AWSDefault.
    For more info, see http://docs.aws.amazon.com/powershell/latest/
userguide//specifying-your-aws-credentials.html
.EXAMPLE
    powershell.exe .\SG_FindSnapshots.ps1
#>

# Criteria to use to filter the results returned.
$volumeID = "vol-424*"
$status = "completed"
$daysBack = 4

# Define filters.
$filter1 = New-Object Amazon.EC2.Model.Filter
$filter1.Name = "volume-id"
$filter1.Value.Add($volumeID)

$filter2 = New-Object Amazon.EC2.Model.Filter
$filter2.Name = "status"
$filter2.Value.Add($status)

$snapshots = get-EC2Snapshot -Filter $filter1, $filter2
$count = 0

foreach ($s in $snapshots)
 {
   $d = ([DateTime]::Now).AddDays(-$daysBack)
   if ([DateTime]::Compare($d, $s.StartTime) -gt 0)
   {
        # Meets criteria.
       $count +=1
       $sb = $s.SnapshotId + ", " + $s.StartTime + ", " + $s.Description
       Write-Output($sb)
   }
}
Write-Output ("Found " + $count + " snapshots that matched the criteria.")
```

Restoring a Snapshot

You can restore a snapshot of a volume to a new AWS Storage Gateway volume, or you can use the snapshot to create an Amazon EBS volume and attach this volume to an Amazon EC2 instance.

When you restore the snapshot to a new AWS Storage Gateway volume, you can mount the volume as an iSCSI device to your on-premises application server. You can then access the contents of the snapshot. Overall, the process is similar to creating a new volume.

The use cases for restoring snapshots depend on the type of gateway you set up—for more information on gateway types, see How AWS Storage Gateway Works (Architecture) (p. 3).

- For gateways created with cached volumes, your volume data is already stored in Amazon S3. In this case, you typically use snapshots to preserve older versions of your data. After initiating a snapshot restore to a gateway-cached volume, snapshot data is downloaded to the local cache only upon first access of the data.
- For stored volumes, your volume data is stored on-premises, In this case, you can use snapshots for durable, off-site backups in Amazon S3. For example, if a local disk allocated as a storage volume crashes, you can provision a new local disk and restore a snapshot to it during the volume creation process. For more information on this approach, see Adding a Volume (p. 266).

  After you initiate a snapshot restore to a gateway-stored volume, snapshot data is downloaded in the background. This functionality means that once you create a volume from a snapshot, you don't need to wait for all of the data to transfer from Amazon S3 to your volume before your application can start accessing the volume and all of its data. If your application accesses a piece of data that has yet to be loaded, the gateway immediately downloads the requested data from Amazon S3. The gateway then continues loading the rest of the volume's data in the background.

Topics

Restoring a Snapshot to a Storage Volume

The following procedure applies to both gateway-cached and gateway-stored volumes.

To restore a snapshot to a new volume

1. On the AWS Storage Gateway console, choose **Create Volume**, and provide the requested information. Depending on the type of gateway you created the requested information might be slightly different, as shown in the screenshots following.

   Information for Cached Volumes

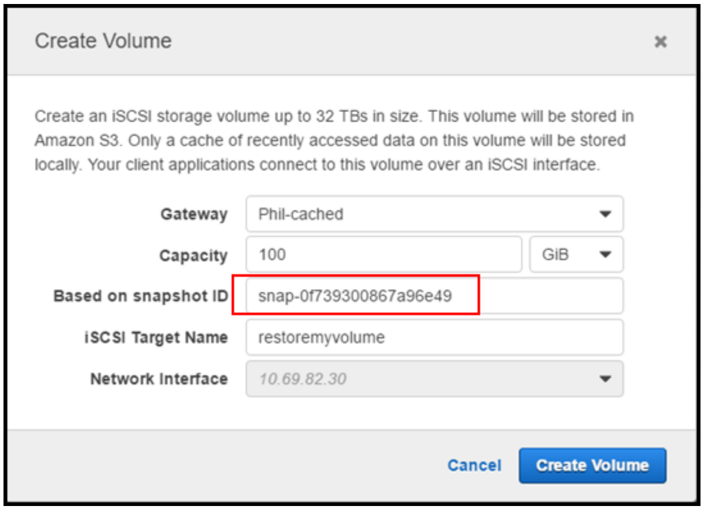

Create Volume

Create an iSCSI storage volume up to 32 TBs in size. This volume will be stored in Amazon S3. Only a cache of recently accessed data on this volume will be stored locally. Your client applications connect to this volume over an iSCSI interface.

Gateway: Phil-cached
Capacity: 100 GiB
Based on snapshot ID: snap-0f739300867a96e49
iSCSI Target Name: restoremyvolume
Network Interface: 10.69.82.30

   Information for Stored Volumes

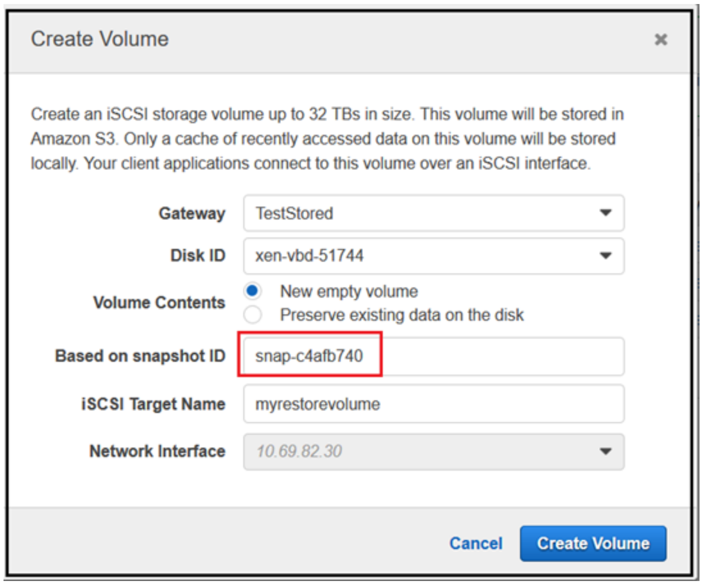

Create an iSCSI storage volume up to 32 TBs in size. This volume will be stored in Amazon S3. Only a cache of recently accessed data on this volume will be stored locally. Your client applications connect to this volume over an iSCSI interface.

Gateway	TestStored
Disk ID	xen-vbd-51744
Volume Contents	● New empty volume ○ Preserve existing data on the disk
Based on snapshot ID	snap-c4afb740
iSCSI Target Name	myrestorevolume
Network Interface	10.69.82.30

2. In the **Create Volume** dialog box, choose a gateway for **Gateway**. The volume is created for the gateway you select.

3. For cached volumes, enter the same **Capacity** value as for the original volume from which you took the snapshot.

   For stored volumes, select the ID for the virtual disk on your VM that you want to use, and then choose **New empty volume**.

4. For **Based on snapshot ID**, specify the snapshot ID

   The gateway downloads your existing snapshot data to the storage volume. For more information, see Restoring a Snapshot to a Storage Volume (p. 262).

5. Type a target name in the **iSCSI Target Name** box.

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

6. Verify that the **Network Interface** setting is the IP address of your gateway, and then choose **Create Volume**. The **Configure CHAP Authentication** dialog box appears.

7. Choose **Create Volume**.

   Doing this creates a storage volume based on the snapshot you specified. The volume details appear in the AWS Storage Gateway console.

8. Connect to the new volume target:

   a. On the **Start** menu of your Windows client computer, choose **Run**, type **`iscsicpl.exe`**, and then choose **OK** to run the iSCSI initiator program.

   b. In the **iSCSI Initiator Properties** dialog box, choose the **Targets** tab. If the new target does not appear in the **Discovered Targets** pane, choose **Refresh**.

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

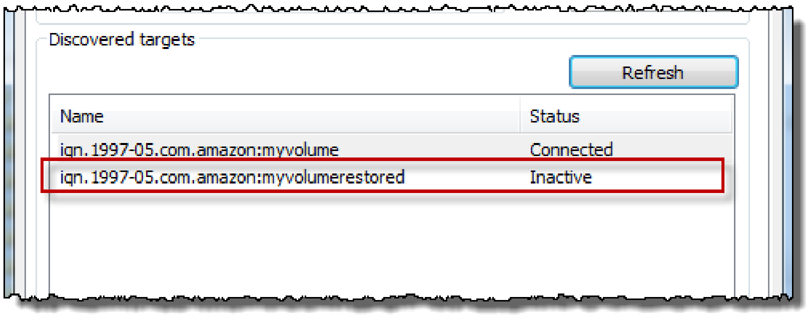

Discovered targets

Name	Status
iqn.1997-05.com.amazon:myvolume	Connected
iqn.1997-05.com.amazon:myvolumerestored	Inactive

c. Select the new target, and choose **Connect**.

d. In the **Connect to Target** dialog box, choose **OK**.

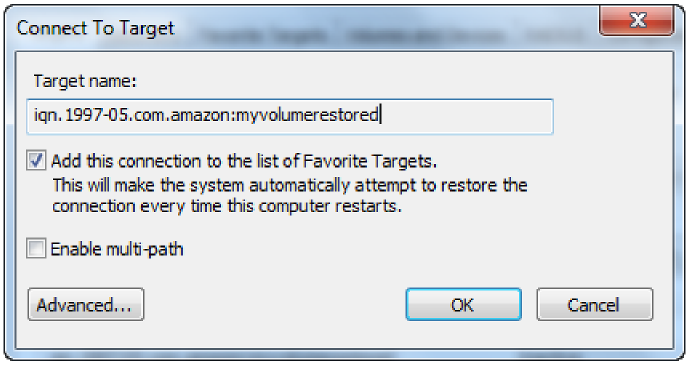

Connect To Target

Target name:
iqn.1997-05.com.amazon:myvolumerestored

☑ Add this connection to the list of Favorite Targets.
This will make the system automatically attempt to restore the connection every time this computer restarts.

☐ Enable multi-path

9. Bring the restored volume online:

a. If the **Disk Management** console is not already open, on the **Start** menu of your Windows client computer, choose **Run**, type `diskmgmt.msc`, and then choose **OK** to open the console.

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

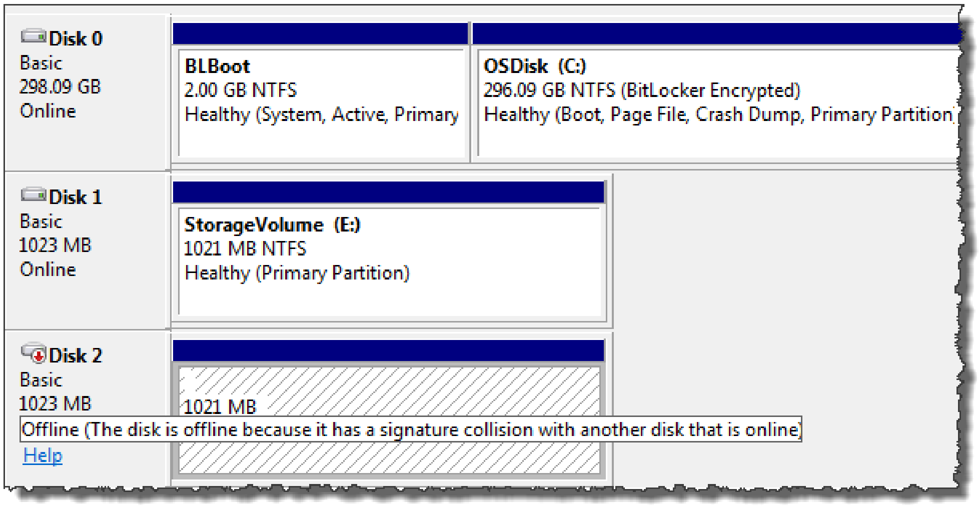

Disk 0
Basic
298.09 GB
Online

BLBoot
2.00 GB NTFS
Healthy (System, Active, Primary

OSDisk (C:)
296.09 GB NTFS (BitLocker Encrypted)
Healthy (Boot, Page File, Crash Dump, Primary Partition

Disk 1
Basic
1023 MB
Online

StorageVolume (E:)
1021 MB NTFS
Healthy (Primary Partition)

Disk 2
Basic
1023 MB
Offline (The disk is offline because it has a signature collision with another disk that is online)
Help

1021 MB

b. Open the context (right-click) menu for the restored volume, and choose **Online**. Doing this brings the volume online and assigns it a different drive letter.

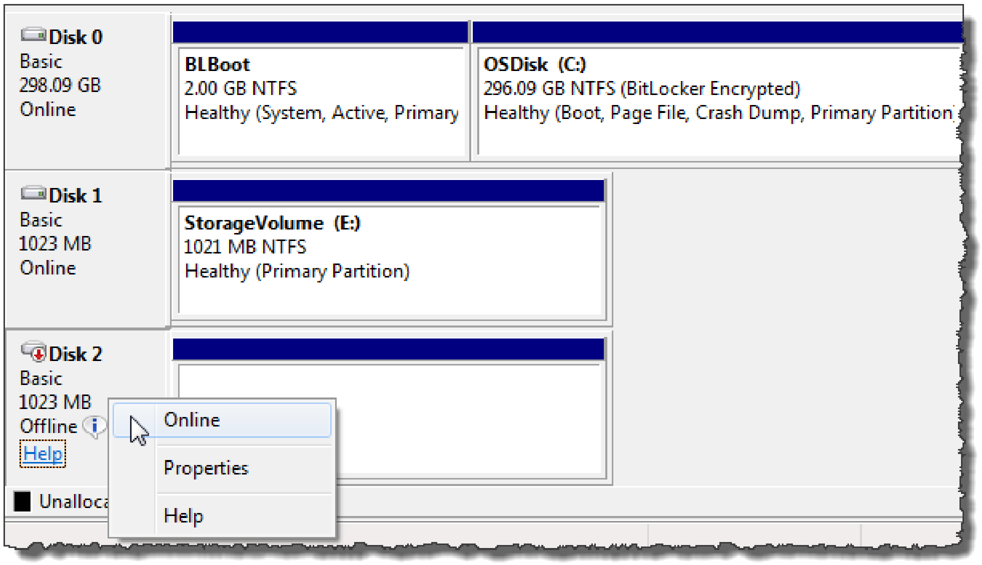

10. Open the restored volume, and verify that the data you saved earlier is there.

Restoring a Snapshot to an Amazon EBS Volume

Your snapshots of local storage volumes taken by AWS Storage Gateway are stored in Amazon S3 as Amazon EBS snapshots. For snapshots up to 16 TiB in size, you can restore snapshots of your local storage volumes to an Amazon EBS volume. You can then attach the Amazon EBS volume to an Amazon EC2 instance. By doing, you can easily migrate data from your on-premises applications to your applications running on Amazon EC2, in case you need to use the Amazon EC2 compute capacity for disaster recovery or data processing. To see detailed pricing for Amazon EC2 and Amazon EBS, see the Amazon EC2 Pricing page.

To restore a snapshot to an Amazon EBS volume

1. Create an Amazon EBS volume.

   - Follow the instructions in Creating an Amazon EBS Volume in the *Amazon Elastic Compute Cloud User Guide*.

     The volume size that you specify must be greater than or equal to the size of the snapshot. To specify the snapshot to use, on the **EBS Volumes** pane of the Amazon EC2 console, choose **Create Volume**, and then choose the snapshot's ID in the list. Alternatively, you can use the Amazon EC2 API to create your Amazon EBS volumes.

     > **Note**
     > By default, AWS limits the maximum size of Amazon EBS volumes you can create per AWS account. For information about default Amazon EBS volume limits, see Amazon EBS Limits. in the *Amazon EC2 User Guide*. For information about how to increase these limits, see Amazon EC2 Service Limits.

2. Attach the Amazon EBS volume to an Amazon EC2 instance. For more information, see Attaching the Volume to an Instance in the *Amazon Elastic Compute Cloud User Guide*.

Working With Volumes

Topics

Managing Volumes

In this section, you can find information about managing existing volumes, including adding new volumes, removing existing volumes, and viewing the status of a volume.

Working With Volumes for Gateway-Stored Volumes

Stored volumes are volumes that are exposed as iSCSI targets on which you can store your application data. The volumes are created on the local virtual disks that you added to your gateway virtual machine (VM)— provides information on how to add and remove volumes for your gateway-stored setup.

Topics

Resizing the underlying disk of a volume is not supported. To change the size of an underlying disk, delete the volume that is using the disk, resize the disk, and then create a new volume from the resized disk. When you recreate the storage volume, be sure to preserve the data on the disk. For steps describing how to remove a volume, see Deleting a Volume (p. 103) or To remove the underlying local disk (Microsoft Hyper-V) (p. 270). For steps describing how to add a volume and preserve existing data, see To create a volume using the console (p. 266).

Adding a Volume

As your application needs grow, you might need to add more volumes to your gateway. As you add more volumes, you must consider the size of your upload buffer you allocated to the gateway. The gateway must have sufficient buffer space. For more information, see Adding and Removing Upload Buffer (p. 144).

You can add volumes using the AWS Storage Gateway console or the AWS Storage Gateway API. For information on using the API to add storage volumes, see CreateStorediSCSIVolume. The following procedure demonstrates using the console and assumes that you already have a deployed and activated gateway. Furthermore, the procedure assumes that there is at least one locally provisioned disk of the gateway that is not used and can be allocated as a storage volume. For information on how to provision a local disk for application storage, see Provisioning Local Disk Storage for the Gateway VM (VMWare) (p. 218).

To create a volume using the console

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Volumes**, and then choose the **Create Volume** button to open the **Create Volume** dialog box.

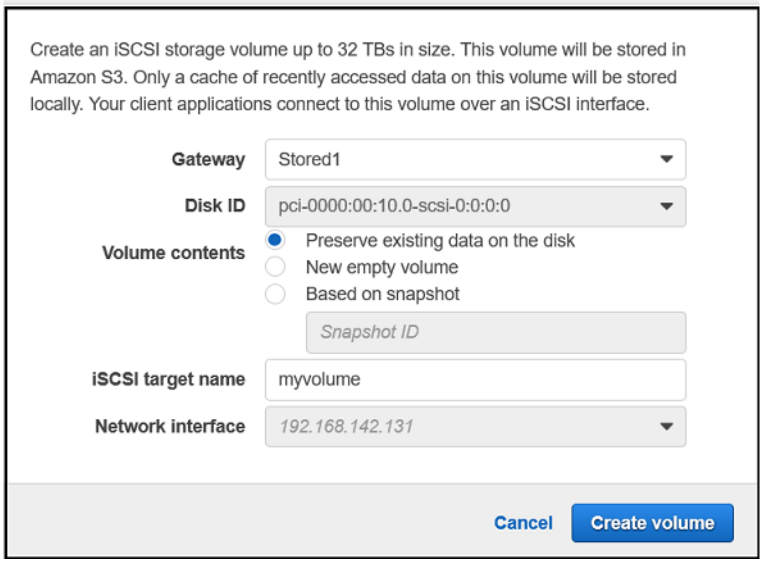

Create an iSCSI storage volume up to 32 TBs in size. This volume will be stored in Amazon S3. Only a cache of recently accessed data on this volume will be stored locally. Your client applications connect to this volume over an iSCSI interface.

Gateway: Stored1

Disk ID: pci-0000:00:10.0-scsi-0:0:0:0

Volume contents:
● Preserve existing data on the disk
○ New empty volume
○ Based on snapshot

Snapshot ID

iSCSI target name: myvolume

Network interface: 192.168.142.131

3. For **Gateway**, choose the gateway you want to create the volume for and configure the volume:
4. In the **Create Volume** dialog box, specify the following information:

   a. For **Disk ID**, choose the local virtual disk that you provisioned for the gateway.

      For information about provisioning disks, see Provisioning Local Disk Storage for the Gateway VM (VMWare) (p. 218).

   b. For **Volume Contents**, choose **New empty volume**.

   c. If you want to preserve data on the disk, select **Preserve existing data on the disk**.

      If you preserve data, AWS Storage Gateway bootstraps your volume upon creation, uploading your volume's existing data to AWS. This process might take some time to complete.

      When you create a volume from an existing snapshot and you don't want to preserve any existing data on the disk, make sure **Preserve existing data on the disk** is clear.

   d. Leave **Based on snapshot ID** blank.

      If you are creating a volume from a snapshot, for **Based on Snapshot ID**, type the snapshot ID.

      You can specify the ID of an existing AWS Storage Gateway or Amazon Elastic Block Store (Amazon EBS) snapshot that you previously created. This approach is useful if you want to restore a snapshot of another storage volume. In this case, the gateway creates the storage volume and downloads your existing snapshot data to the volume. However, you don't need to wait for all of the data to transfer from Amazon S3 to your volume before your application can start accessing the volume and all of its data. For more information about snapshots, see Working with Snapshots (p. 253).

      Important
      If the Amazon EBS volume you create the snapshot from contains data from a public Amazon Machine Image (AMI), such as an AMI from AWS Marketplace, you will not have the permission to create the volume.
      Storage volumes cannot be created from snapshots that are encrypted or not owned by you but shared with your account.
      The size of the storage volume you create must be equal to or greater than the size of the snapshot you create it from. For information on how to add a disk to your gateway VM that can be used as a gateway-stored volume, see Provisioning Local Disk Storage for the Gateway VM (VMWare) (p. 218). Once you add such

a volume, you can access the contents of the volume from your on-premises applications. For more information, see Connecting to Volumes on Your Volume Gateway (p. 279).

By default, AWS limits the maximum size of Amazon EBS volumes you can create per AWS account. For information about default Amazon EBS volume limits, see Amazon EBS Limits in the *Amazon EC2 User Guide*. For information about how to increase these limits, see Amazon EC2 Service Limits.

e. For **iSCSI Target Name**, type a name.

   The target name can contain lowercase letters, numbers, periods (.), and hyphens (-). This target name appears as the **iSCSI Target Node** name on the **Targets** tab of the iSCSI Microsoft Initiator UI after discovery. For example, the name `target1` appears as `iqn.1997-05.com.amazon:target1`. Ensure that the target name is globally unique within your storage area network (SAN).

f. If you've configured your local gateway host with multiple network interface cards (NICs), for **Host IP**, specify the IP address for the NIC to use for this storage volume.

g. Check that **Port** shows 3260. The **Port** box shows the port to which to map an iSCSI target. AWS Storage Gateway supports only port 3260.

h. Choose **Create Volume**.

   Doing this creates a storage volume and makes your disk available as an iSCSI target for your applications to connect to and store data on.

   Choosing **Create Volume** also creates a snapshot schedule for your new volume. By default, AWS Storage Gateway takes snapshots once a day. You can change both the time the snapshot occurs each day and the frequency (every 1, 2, 4, 8, 12, or 24 hours). For more information, see Editing a Snapshot Schedule (p. 104).

   Note
   Snapshots are incremental, compressed backups. For a given storage volume, the gateway saves only the blocks that have changed since the last snapshot. This approach minimizes the amount of storage that is used for your backups.

Removing a Volume

You might need to remove a volume as your application needs change—for example, if you migrate your application to use a larger volume and you want to reclaim the underlying local disk space of the old volume. To reclaim the local disk space, you need to remove the local disk from the VM.

Before removing the volume, make sure that there are no applications currently writing to the volume. Also, make sure that there are no snapshots in progress for the volume. You can check the snapshot schedule of storage volumes on the **Snapshot Schedules** tab of the console. For more information, see Editing a Snapshot Schedule (p. 104).

You can remove volumes using the AWS Storage Gateway console or the AWS Storage Gateway API. For information on using the API to remove volumes, see DeleteVolume. The following procedure demonstrates using the console and either the vSphere client for a gateway deployed on the VMware ESXi platform or the Microsoft Hyper-V Manager for a gateway deployed on the Microsoft Hyper-V platform.

To remove a storage volume on the AWS Storage Gateway console

Note
Perform the following procedure with the gateway running.

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

2. On the **Volumes** tab, choose the volume.
3. If you plan to remove the disk from the VM that backs the volume, choose the **Details** tab, and note the value for **Local Disk**.

   This value is the disk's virtual device node value, which you use in the hypervisor client to help ensure that you remove the correct disk.
4. Choose **Delete Volume**.
5. If you also want to remove the underlying local disk, do one of the following:

For a Gateway Hosted In	Do This
VMware ESXi	Follow the steps in Deleting a Volume (p. 103).
Microsoft Hyper-V	Follow the steps in To remove the underlying local disk (Microsoft Hyper-V) (p. 270).

To remove the underlying local disk (VMware ESXi)

1. In the vSphere client, open the context (right-click) menu for your gateway VM, and then choose **Edit Settings**.
2. On the **Hardware** tab of the **Virtual Machine Properties** dialog box, choose the disk to remove, 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 in step 2 of the preceding procedure. Doing this helps ensure that you remove the correct disk. The first SCSI controller displayed in the Microsoft Hyper-V Manager is controller 0.

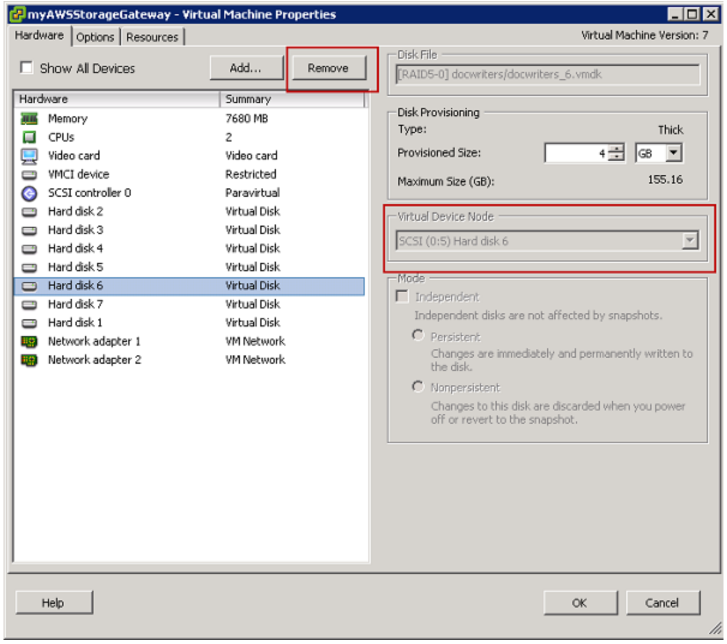

Virtual Machine Version: 7

□ Show All Devices

Hardware	Summary
Memory	7680 MB
CPUs	2
Video card	Video card
VMCI device	Restricted
SCSI controller 0	Paravirtual
Hard disk 2	Virtual Disk
Hard disk 3	Virtual Disk
Hard disk 4	Virtual Disk
Hard disk 5	Virtual Disk
Hard disk 6	Virtual Disk
Hard disk 7	Virtual Disk
Hard disk 1	Virtual Disk
Network adapter 1	VM Network
Network adapter 2	VM Network

Disk File
[RAID5-0] docwriters/docwriters_6.vmdk

Disk Provisioning
Type: Thick
Provisioned Size: 4 GB
Maximum Size (GB): 155.16

Virtual Device Node
SCSI (0:5) Hard disk 6

Mode
□ Independent
Independent disks are not affected by snapshots.
○ Persistent
Changes are immediately and permanently written to the disk.
○ Nonpersistent
Changes to this disk are discarded when you power off or revert to the snapshot.

3. Choose the appropriate option in the **Removal Options** panel, and then choose **OK** to complete the process of removing the disk.

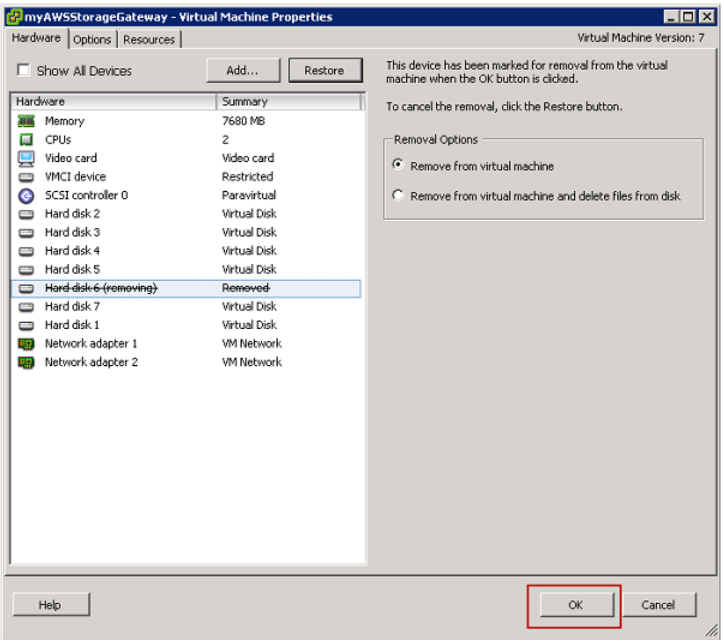
myAWSStorageGateway - Virtual Machine Properties
Hardware
Options
Resources
Virtual Machine Version: 7
□ Show All Devices
Add...
Restore
Help
OK
Cancel

Hardware	Summary
Memory	7680 MB
CPUs	2
Video card	Video card
VMCI device	Restricted
SCSI controller 0	Paravirtual
Hard disk 2	Virtual Disk
Hard disk 3	Virtual Disk
Hard disk 4	Virtual Disk
Hard disk 5	Virtual Disk
~~Hard disk 6 (removing)~~	~~Removed~~
Hard disk 7	Virtual Disk
Hard disk 1	Virtual Disk
Network adapter 1	VM Network
Network adapter 2	VM Network

This device has been marked for removal from the virtual machine when the OK button is clicked.

To cancel the removal, click the Restore button.

Removal Options

● Remove from virtual machine

○ Remove from virtual machine and delete files from disk

To remove the underlying local disk (Microsoft Hyper-V)

1. In the Microsoft Hyper-V Manager, open the context (right-click) menu for your gateway VM, and then choose **Settings**.
2. For **Hardware** in the **Settings** dialog box, choose the disk to remove, and then choose **Remove**.

   The disks you add to a gateway are under the **SCSI Controller** entry in the **Hardware** list.

   Verify that the **Controller** and **Location** values are the same as the value that you noted from a previous step. Doing this helps ensure that you remove the correct disk.

   The first SCSI controller displayed in the Microsoft Hyper-V Manager is controller 0.

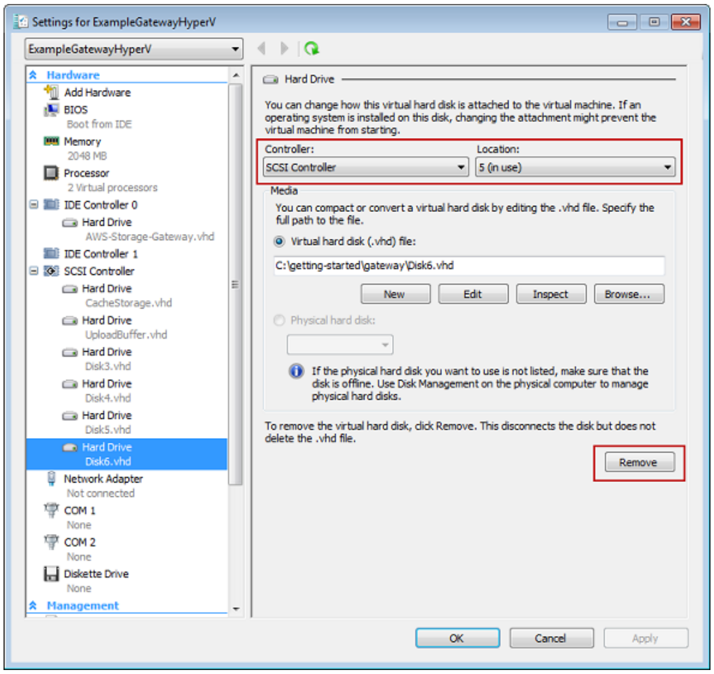

Settings for ExampleGatewayHyperV

Hard Drive

You can change how this virtual hard disk is attached to the virtual machine. If an operating system is installed on this disk, changing the attachment might prevent the virtual machine from starting.

Controller: SCSI Controller

Location: 5 (in use)

Media

You can compact or convert a virtual hard disk by editing the .vhd file. Specify the full path to the file.

● Virtual hard disk (.vhd) file: C:\getting-started\gateway\Disk6.vhd

○ Physical hard disk:

If the physical hard disk you want to use is not listed, make sure that the disk is offline. Use Disk Management on the physical computer to manage physical hard disks.

To remove the virtual hard disk, click Remove. This disconnects the disk but does not delete the .vhd file.

3. Choose **OK** to apply the change.

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 for Linux Instances*.

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 Sizing the Upload Buffer (p. 145) and Adding and Removing Upload Buffer (p. 144).

There are limits to 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 limits. For more information, see AWS Storage Gateway Limits (p. 296).

To create an Amazon EBS volume, attach it, 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 for Linux Instances*.
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 for Linux Instances*.
3. Configure the Amazon EBS volume you added as either an upload buffer or cache storage. For instructions, see Managing Local Disks for Your AWS Storage Gateway (p. 142).

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. If you remove an Amazon EBS volume that is allocated as cache storage from a gateway, virtual tapes on the gateway will have the IRRECOVERABLE status, and you risk data loss. For more information on the IRRECOVERABLE status, see Tape Status Information in a VTL (p. 118).

1. Shut down the gateway by following the approach described in the Starting and Stopping Your Gateway (p. 141) 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 for Linux Instances*.
3. Delete the Amazon EBS volume. For instructions, see Deleting an Amazon EBS Volume in the *Amazon EC2 User Guide for Linux Instances*.
4. Start the gateway by following the approach described in the Starting and Stopping Your Gateway (p. 141) section.

Tape Gateway

Topics

Working with VTL Devices

Your tape gateway setup provides the following SCSI devices, which you select when activating your gateway.

Topics

For medium changers, AWS Storage Gateway works with the following:

- Tape Gateway—This device is provided with the gateway.
- STK-L700—This device emulation is provided with the gateway.

> **Important**
> The type of medium changer you select depends on the backup software you plan to use. The following table shows which medium changer to select for your backup software. This list includes third-party backup software that has been tested and found to be compatible with tape gateway.

Backup Software	Medium Changer Type
Backup Exec 2012	`STK-L700`
Backup Exec 2014	`AWS-Gateway-VTL`
Backup Exec 15	`AWS-Gateway-VTL`

Backup Software	Medium Changer Type
Backup Exec 16	`AWS-Gateway-VTL`
Dell NetVault Backup 10.0	`STK-L700`
EMC NetWorker 8.x	`STK-L700`
HPE Data Protector 9.x	`AWS-Gateway-VTL`
Microsoft System Center 2012 R2 Data Protection Manager	`STK-L700`
Symantec NetBackup Version 7.x	`AWS-Gateway-VTL`
Veeam Backup & Replication V7	`STK-L700`
Veeam Backup & Replication V8	`STK-L700`
Veeam Backup & Replication V9 Update 2 or later	`AWS-Gateway-VTL`

> **Note**
> You must select the medium changer that is recommended for your backup software. Other medium changers might not function properly. You can select a different medium changer after the gateway is activated. For more information, see Selecting a Medium Changer After Gateway Activation (p. 273).

For tape drives, AWS Storage Gateway works with the following:

- IBM-ULT3680-TD5—This device emulation is provided with the gateway.

The following screenshot shows the activation page for a tape gateway.

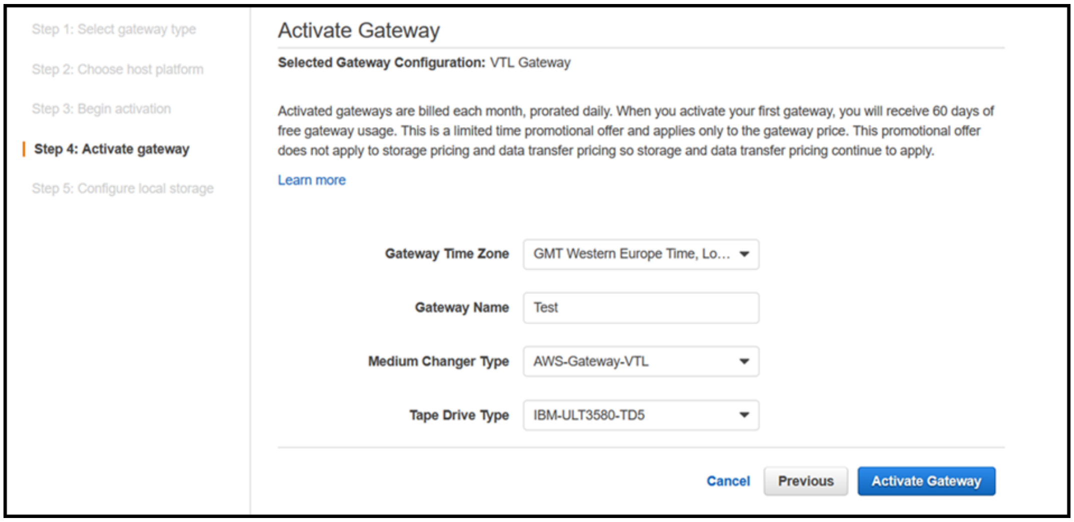
Step 1: Select gateway type
Step 2: Choose host platform
Step 3: Begin activation
Step 4: Activate gateway
Step 5: Configure local storage

Activate Gateway

Selected Gateway Configuration: VTL Gateway

Activated gateways are billed each month, prorated daily. When you activate your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies only to the gateway price. This promotional offer does not apply to storage pricing and data transfer pricing so storage and data transfer pricing continue to apply.

Learn more

Gateway Time Zone GMT Western Europe Time, Lo...

Gateway Name Test

Medium Changer Type AWS-Gateway-VTL

Tape Drive Type IBM-ULT3580-TD5

Selecting a Medium Changer After Gateway Activation

When you activate your tape gateway, you select a medium changer type for the gateway. You can choose to select a different medium changer type after the gateway is activated.

> **Important**
>
> If your tape gateway uses the Symantec Backup Exec 2014 or NetBackup 7.x backup software, you must select the Tape Gateway device type. For more information on how to change the medium changer after gateway activation for these applications, see Best Practices for using Symantec Backup products (NetBackup, Backup Exec) with the Amazon Web Services (AWS) Storage Tape Gateway in *Symantec Support*.

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 AWS Storage Gateway console, choose **Gateways** from the navigation pane, and then choose the gateway whose medium changer you want to change.
6. Choose the **VTL Devices** tab, select the medium changer you want to change, and then choose the **Change Media Changer** button,

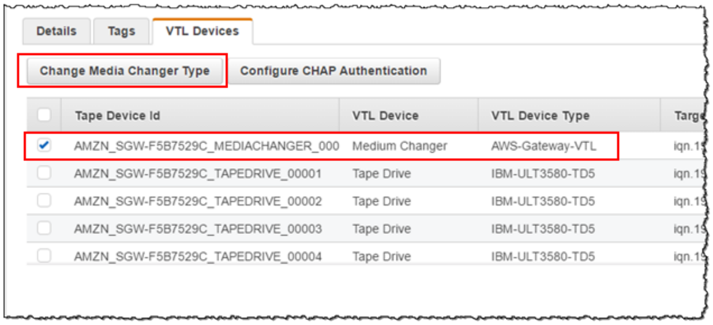

□	Tape Device Id	VTL Device	VTL Device Type	Targe
☑	AMZN_SGW-F5B7529C_MEDIACHANGER_000	Medium Changer	AWS-Gateway-VTL	iqn.1
□	AMZN_SGW-F5B7529C_TAPEDRIVE_00001	Tape Drive	IBM-ULT3580-TD5	iqn.1
□	AMZN_SGW-F5B7529C_TAPEDRIVE_00002	Tape Drive	IBM-ULT3580-TD5	iqn.1
□	AMZN_SGW-F5B7529C_TAPEDRIVE_00003	Tape Drive	IBM-ULT3580-TD5	iqn.1
□	AMZN_SGW-F5B7529C_TAPEDRIVE_00004	Tape Drive	IBM-ULT3580-TD5	iqn.1

7. 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

Depending on the backup software you use on your Windows server, you might need to update the driver for your medium changer.

1. Open Device Manager on your Windows server, and expand the **Medium Changer devices** tree.
2. 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.

3. 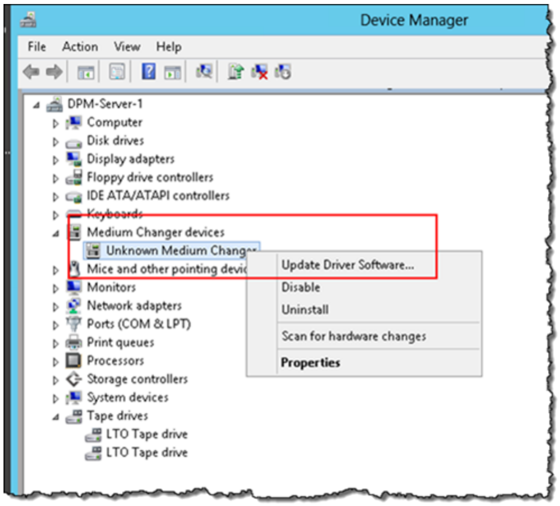

   In the **How do you want to search for driver software?** section, choose **Browse my computer for driver software**.

4. 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 V7, Veeam Backup & Replication V8, and Microsoft System Center 2012 R2 Data Protection Manager backup software. This Sony driver has been tested with these types of backup software.

5. 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**.

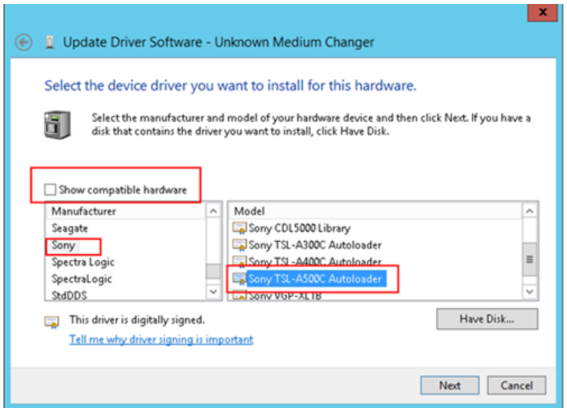

Update Driver Software - Unknown Medium Changer

Select the device driver you want to install for this hardware.

Select the manufacturer and model of your hardware device and then click Next. If you have a disk that contains the driver you want to install, click Have Disk.

☐ Show compatible hardware

This driver is digitally signed.

Manufacturer
Seagate
Sony
Spectra Logic
SpectraLogic
StdDDS
Model
Sony CDL5000 Library
Sony TSL-A300C Autoloader
Sony TSL-A400C Autoloader
Sony TSL-A500C Autoloader
Sony VGP-XL1B
Have Disk...
Tell me why driver signing is important
Next
Cancel

6. In the warning box that appears, choose **Yes**. If the driver is successfully installed, close the **Update drive software** window.

Working With Tapes

AWS 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 Tape Status Information in a VTL (p. 118).

Topics

The tape library shows tapes in your tape gateway. The library shows the tape barcode, status, and size and the gateway the tape is associated with.

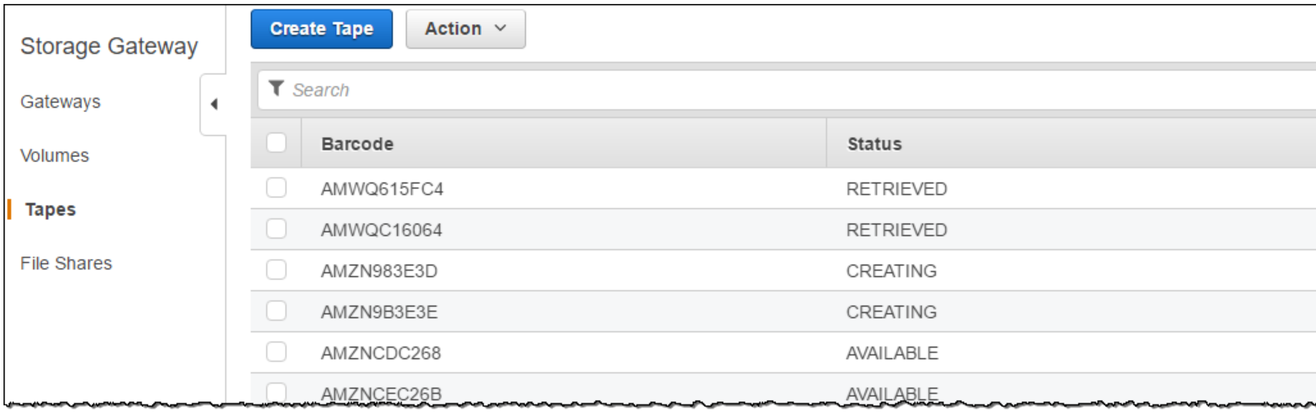

□	Barcode	Status
□	AMWQ615FC4	RETRIEVED
□	AMWQC16064	RETRIEVED
□	AMZN983E3D	CREATING
□	AMZN9B3E3E	CREATING
□	AMZNCDC268	AVAILABLE
□	AMZNCEC26B	AVAILABLE

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 AWS 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, AWS 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 AWS Storage Gateway console.

   If you filter for gateway, tapes that are associated with that gateway appear in the library in the AWS Storage Gateway console.

   > **Note**
   > By default, AWS 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, AWS 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 (p. 277).

> **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 AWS 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 (p. 117).

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 (p. 90).

Canceling Tape Retrieval

After you start tape retrieval from the archive to your tape gateway, and before the retrieval is complete, you might decide that you don’t need to retrieve the tape. You can cancel the retrieval that is in progress. Canceling the retrieval puts the tape back into the archive.

The following procedure shows you how to cancel the retrieval process.

To cancel tape retrieval

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. Choose the **VTL Tape Cartridges** tab, and then choose the tape you want to stop retrieving.

   For information about how to find a tape in the console, see Working With Tapes (p. 275).
3. On the **Details** tab, choose **Stop** in the **Status** row.

   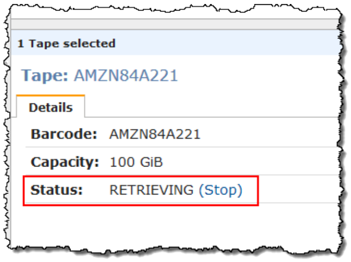

   1 Tape selected

   Tape: AMZN84A221

   Barcode: AMZN84A221
   Capacity: 100 GiB
   Status: RETRIEVING (Stop)
4. In the dialog box that appears, choose **OK**. After the retrieval is canceled, AWS Storage Gateway removes the tape from the list of tapes in your VTL, and changes the tape status from RETRIEVING to ARCHIVED.

   > **Note**
   > You are charged for the portion the data that was retrieved from archive before you canceled the tape retrieval.

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 AWS Storage Gateway console.

To cancel tape archival

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. Choose the **VTL Tape Cartridges** tab, and then choose the tape you want to stop from archiving. The status of the tape must be IN TRANSIT TO VTS.
3. On the **Details** tab, choose **Stop** in the **Status** row.
4. Choose **Cancel Archival**. After archiving is canceled, the status of the tape returns to its original status.

General

Topics

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).

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. AWS Storage Gateway only supports software initiators.

iSCSI target
The server component of the iSCSI network that receives and responds to requests from initiators. Each of your 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 enables 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 is already

installed on Windows Server 2008 R2, Windows 7, Windows Server 2008, and Windows Vista. On these operating systems, you don't need to install the initiator.

Red Hat iSCSI initiator

The iscsi-initiator-utils Resource Package Manager (RPM) package provides you with an iSCSI initiator implemented in software for Red Hat 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.

Topics

Connecting to Volumes on Your Volume Gateway

A volume gateway exposes volumes you have created for the gateway as iSCSI targets. For more information, see Connecting Your Volumes to Your Microsoft Windows Client (p. 34).

Note

To connect to your volume target, your gateway must have an upload buffer configured. If an upload buffer is not configured for your gateway, then the status of your volumes is displayed as UPLOAD BUFFER NOT CONFIGURED. To configure an upload buffer for a gateway in a gateway-stored setup, see To configure upload buffer or cache storage (p. 143). To configure an upload buffer for a gateway in a gateway-cached setup, see To configure upload buffer or cache storage (p. 143).

The following diagram highlights the iSCSI target in the larger picture of the AWS Storage Gateway architecture. For more information, see How AWS Storage Gateway Works (Architecture) (p. 3).

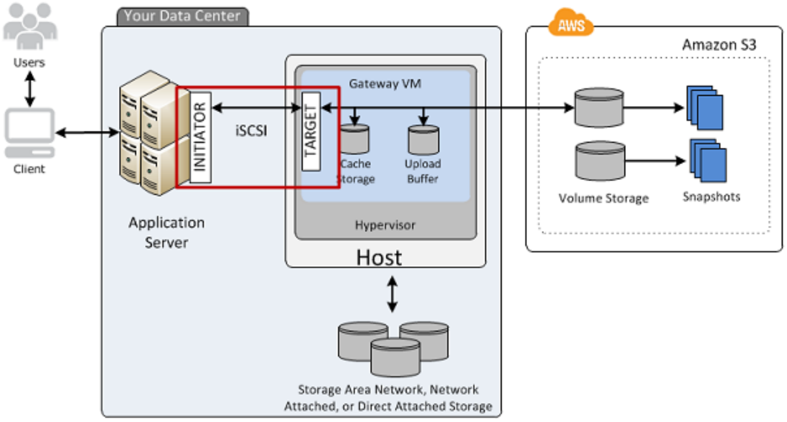

You can connect to your volume from either a Windows or Red Hat Linux client. You can optionally configure CHAP for either client type.

Your gateway exposes your volume as an iSCSI target with a name you specify, prepended by `iqn.1997-05.com.amazon:`. For example, if you specify a target name of `myvolume`, then the iSCSI target you use to connect to the volume is `iqn.1997-05.com.amazon:myvolume`. For more information about how to configure your applications to mount volumes over iSCSI, see Connecting to Volumes on Your Volume Gateway (p. 279).

To	See
Connect to your volume from Windows.	Connecting Your Volumes to Your Microsoft Windows Client (p. 34) in the Getting Started section
Connect to your volume from Red Hat Linux.	Recommended Red Hat Linux iSCSI Settings (p. 44)
Configure CHAP authentication for Windows and Red Hat Linux.	Configuring CHAP Authentication for Your iSCSI Targets (p. 281)

Connecting to VTL Devices on Your Gateway

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 (p. 11).

> **Note**
> You connect only one application to each iSCSI target.

The following diagram highlights the iSCSI target in the larger picture of the AWS Storage Gateway architecture. For more information on AWS Storage Gateway architecture, see Tape Gateway (p. 6).

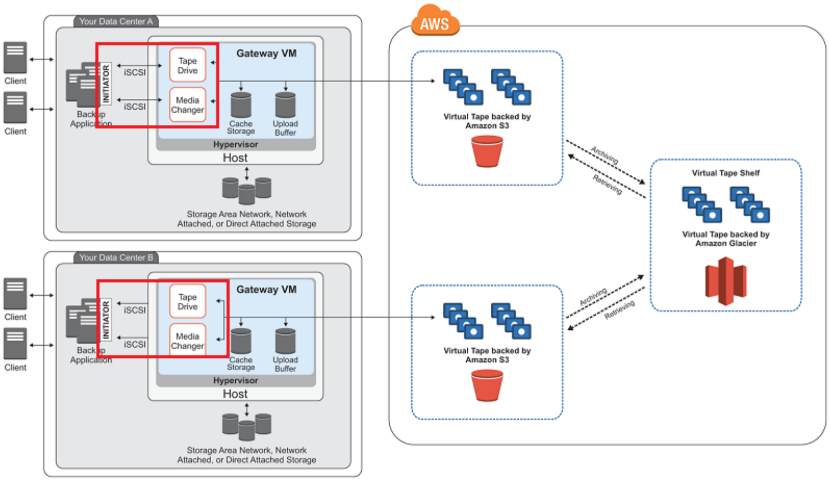

Customizing Your Windows iSCSI Settings

When using a Windows client, you use the Microsoft iSCSI initiator to connect to your gateway volume. For instructions on how to connect to your volumes, see Connecting Your Volumes to Your Microsoft Windows Client (p. 34).

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 Connect Your Tape Gateway Devices to Your Windows Client (p. 59) and Configuring NetBackup Storage Devices (p. 82).

Configuring CHAP Authentication for Your iSCSI Targets

AWS 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.

To set up CHAP, you must configure it both on the AWS 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 AWS Storage Gateway console, as discussed in To configure CHAP for a volume target on the AWS Storage Gateway console (p. 281).
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 (p. 283).
   - To configure mutual CHAP on a Red Hat Linux client, see To configure mutual CHAP on a Red Hat Linux client (p. 288).

To configure CHAP for a volume target on the AWS Storage Gateway console

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

1. On the AWS Storage Gateway console, choose **Volumes** in the navigation pane.
2. On the **Actions** menu, choose **Configure CHAP Authentication**.
3. Provide the requested information in the **Configure CHAP Authentication** dialog box, shown in the screenshot following:

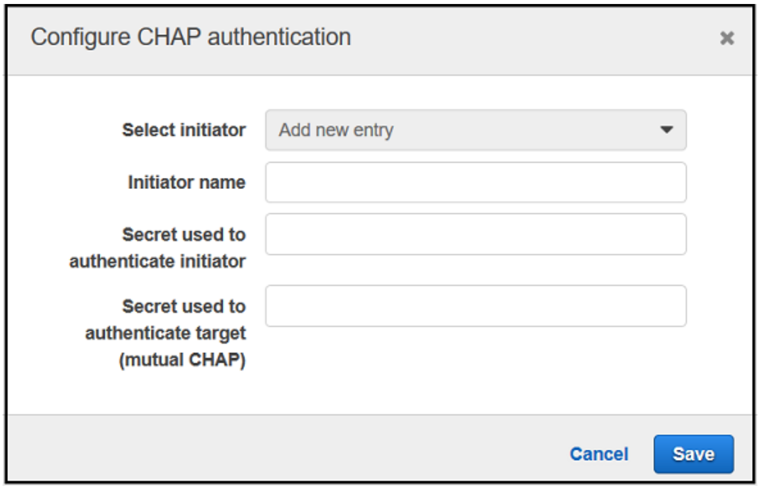
Configure CHAP authentication
×
Select initiator
Add new entry
Initiator name
Secret used to authenticate initiator
Secret used to authenticate target (mutual CHAP)
Cancel
Save

   a. For **Initiator Name**, type the name of your iSCSI initiator.

      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 (p. 283).

> **Note**
> To change an initiator name, you must first disable CHAP, change the initiator name in your iSCSI initiator software, and then enable CHAP with the new name.

b. For **Secret used to Authenticate Initiator**, type 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)**, type the secret requested.

   This secret must be a minimum of 12 characters and a maximum of 16 characters long. This value is the secret key that the target must know to participate in CHAP with the initiator.

   > **Note**
   > The secret used to authenticate the target must be different than the secret to authenticate the initiator.

d. Choose **Save**.

4. Choose the **Details** tab and confirm that **iSCSI CHAP authentication** is set to **true**.

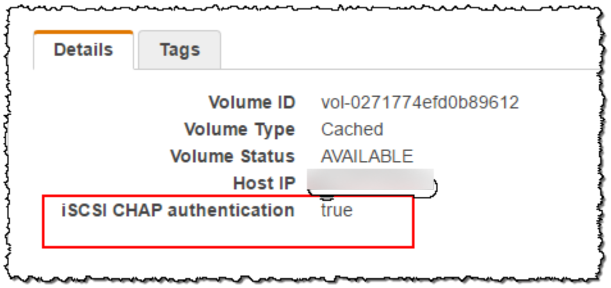

Volume ID	vol-0271774efd0b89612
Volume Type	Cached
Volume Status	AVAILABLE
Host IP	[redacted]
iSCSI CHAP authentication	true

To configure CHAP for a VTL device target on the AWS 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 you want to configure CHAP for.
4. Provide the requested information in the **Configure CHAP Authentication** dialog box, shown in the screenshot following:

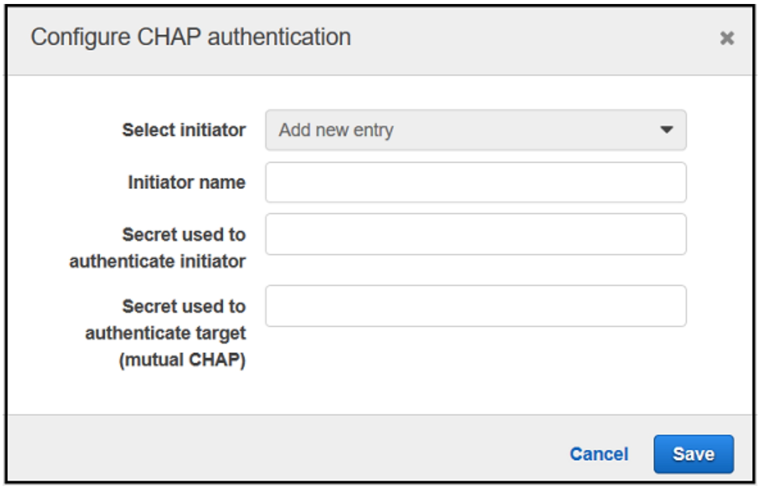
Configure CHAP authentication
×
Initiator name
Secret used to authenticate initiator
Secret used to authenticate target (mutual CHAP)
Cancel
Save

Select initiator: Add new entry

a. For **Initiator Name**, type the name of your iSCSI initiator.

   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 (p. 283).

   > **Note**
   > To change an initiator name, you must first disable CHAP, change the initiator name in your iSCSI initiator software, and then enable CHAP with the new name.

b. For **Secret used to Authenticate Initiator**, type 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)**, type 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**, type `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.

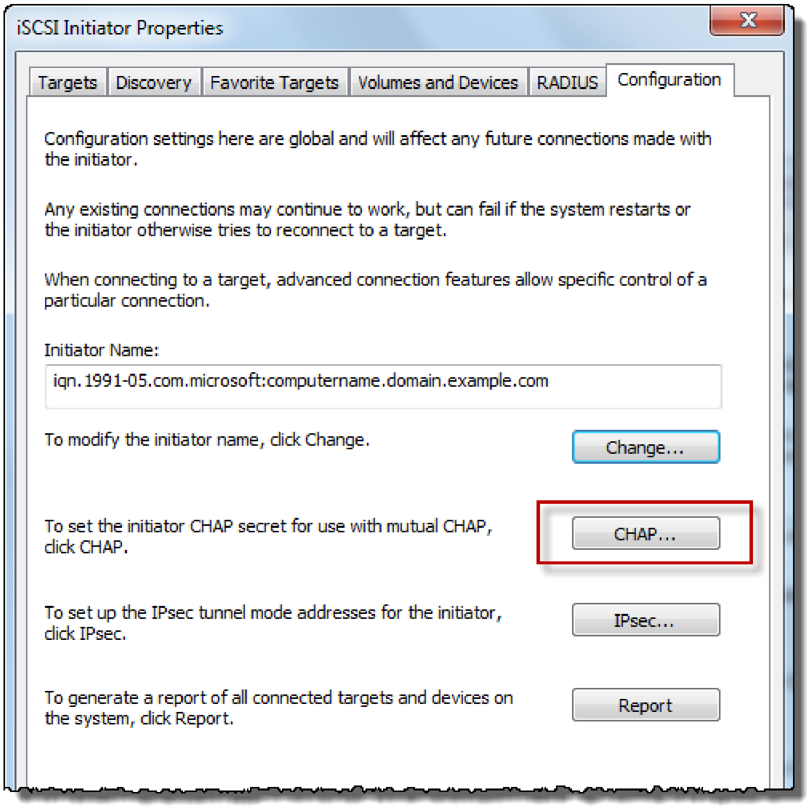

Configuration settings here are global and will affect any future connections made with the initiator.

Any existing connections may continue to work, but can fail if the system restarts or the initiator otherwise tries to reconnect to a target.

When connecting to a target, advanced connection features allow specific control of a particular connection.

Initiator Name:

iqn.1991-05.com.microsoft:computername.domain.example.com

To modify the initiator name, click Change.

To set the initiator CHAP secret for use with mutual CHAP, click CHAP.

To set up the IPsec tunnel mode addresses for the initiator, click IPsec.

To generate a report of all connected targets and devices on the system, click Report.

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 AWS 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, type the mutual CHAP secret value.

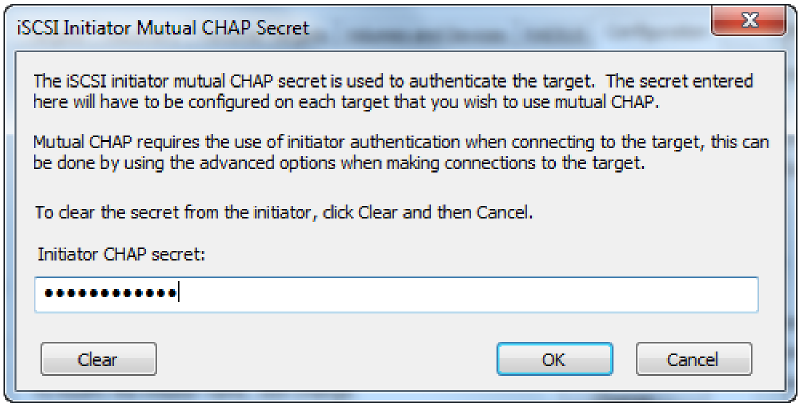

The iSCSI initiator mutual CHAP secret is used to authenticate the target. The secret entered here will have to be configured on each target that you wish to use mutual CHAP.

Mutual CHAP requires the use of initiator authentication when connecting to the target, this can be done by using the advanced options when making connections to the target.

To clear the secret from the initiator, click Clear and then Cancel.

Initiator CHAP secret:

••••••••••••

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 typed 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 (p. 281).

d. If the key that you typed is less than 12 characters or more than 16 characters long, an **Initiator CHAP secret** error dialog box appears.

   Choose **OK**, and then type the key again.

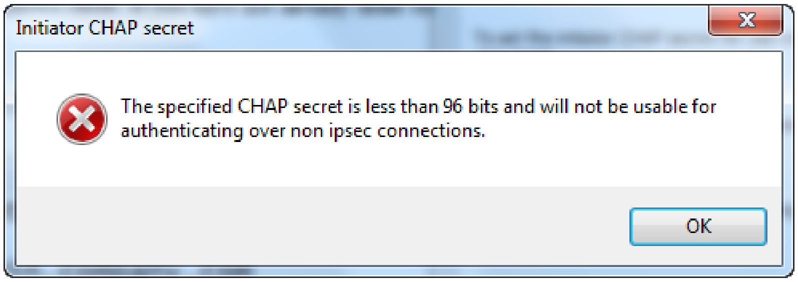

The specified CHAP secret is less than 96 bits and will not be usable for authenticating over non ipsec connections.

3. Configure the target with the initiator's secret to complete the mutual CHAP configuration.

   a. Choose the **Targets** tab.

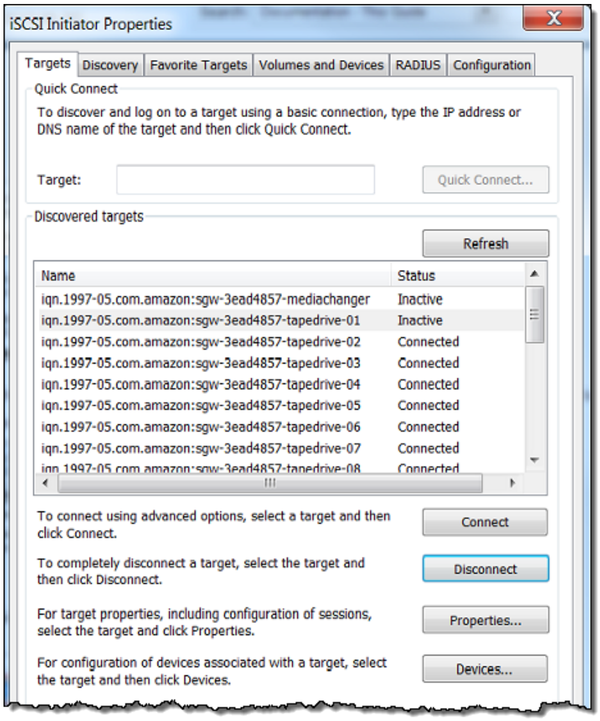

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Name	Status
iqn.1997-05.com.amazon:sgw-3ead4857-mediachanger	Inactive
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-01	Inactive
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-02	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-03	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-04	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-05	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-06	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-07	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-08	Connected

To connect using advanced options, select a target and then click Connect.

To completely disconnect a target, select the target and then click Disconnect.

For target properties, including configuration of sessions, select the target and click Properties.

For configuration of devices associated with a target, select the target and then click Devices.

   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**.

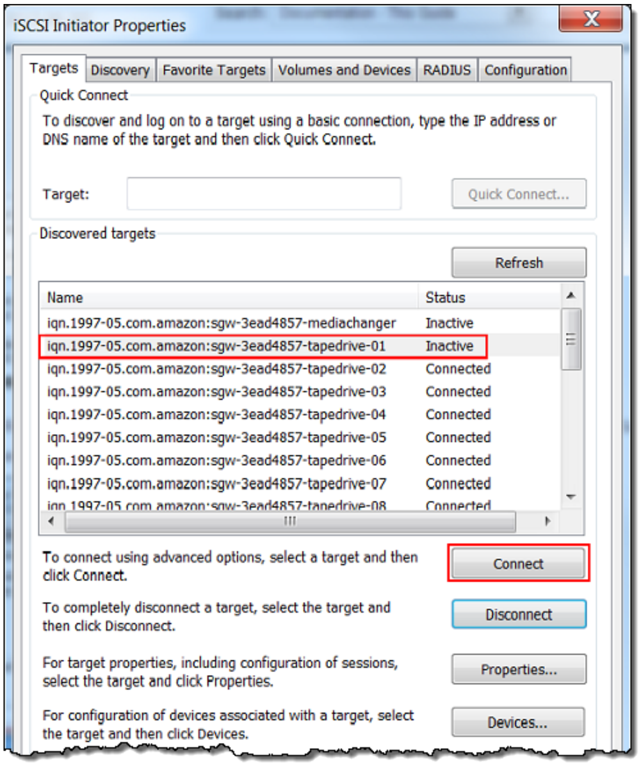

iSCSI Initiator Properties

Quick Connect

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Target:

Discovered targets

Name	Status
iqn.1997-05.com.amazon:sgw-3ead4857-mediachanger	Inactive
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-01	Inactive
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-02	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-03	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-04	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-05	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-06	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-07	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-08	Connected

To connect using advanced options, select a target and then click Connect.

To completely disconnect a target, select the target and then click Disconnect.

For target properties, including configuration of sessions, select the target and click Properties.

For configuration of devices associated with a target, select the target and then click Devices.

d. In the **Connect to Target** dialog box, choose **Advanced**.

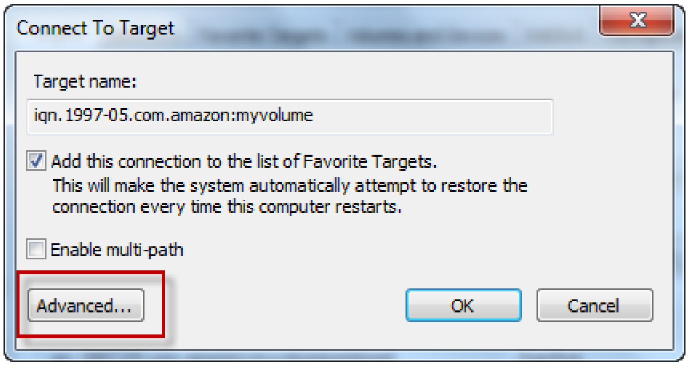

Connect To Target

Target name:

iqn.1997-05.com.amazon:myvolume

☑ Add this connection to the list of Favorite Targets.
This will make the system automatically attempt to restore the connection every time this computer restarts.

☐ Enable multi-path

e. In the **Advanced Settings** dialog box, configure CHAP.

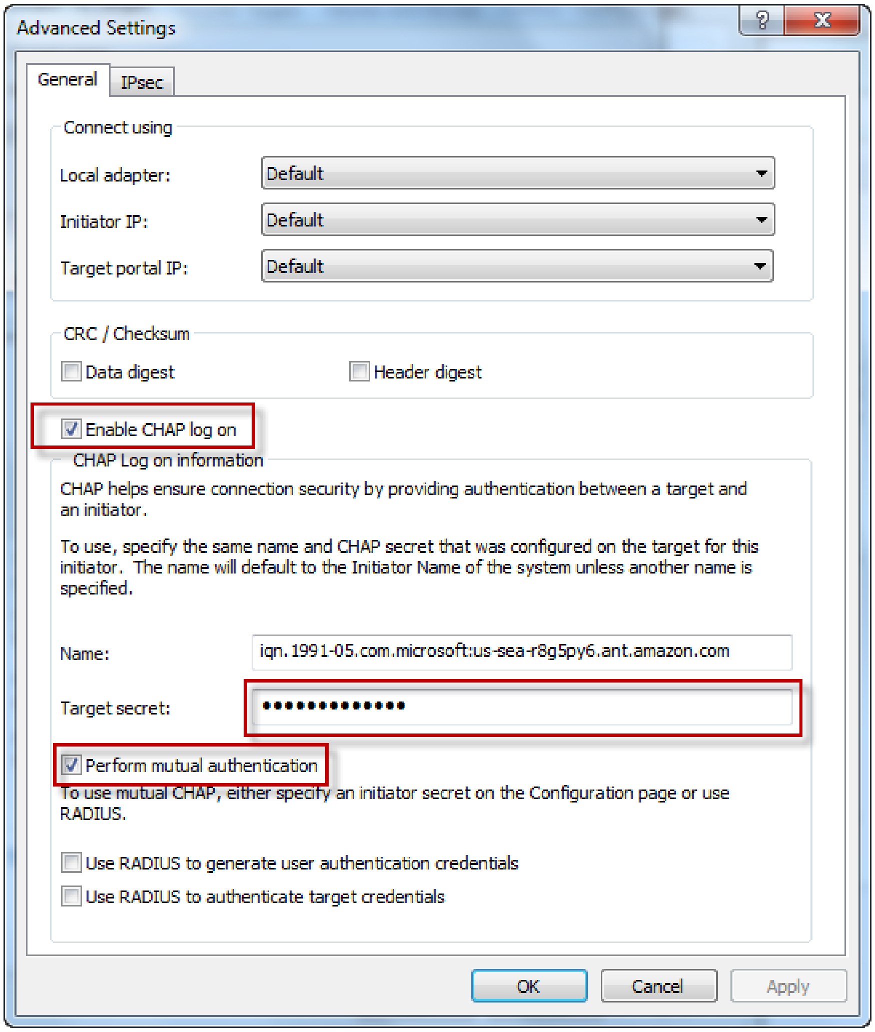

Connect using

Local adapter: Default

Initiator IP: Default

Target portal IP: Default

CRC / Checksum

☐ Data digest ☐ Header digest

☑ Enable CHAP log on

CHAP Log on information

CHAP helps ensure connection security by providing authentication between a target and an initiator.

To use, specify the same name and CHAP secret that was configured on the target for this initiator. The name will default to the Initiator Name of the system unless another name is specified.

Name: iqn.1991-05.com.microsoft:us-sea-r8g5py6.ant.amazon.com

Target secret: ••••••••••••••

☑ Perform mutual authentication

To use mutual CHAP, either specify an initiator secret on the Configuration page or use RADIUS.

☐ Use RADIUS to generate user authentication credentials

☐ Use RADIUS to authenticate target credentials

i. Select **Enable CHAP log on**.

ii. Type the secret that is required to authenticate the initiator. This secret is the same as the secret typed 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 (p. 281).

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**.

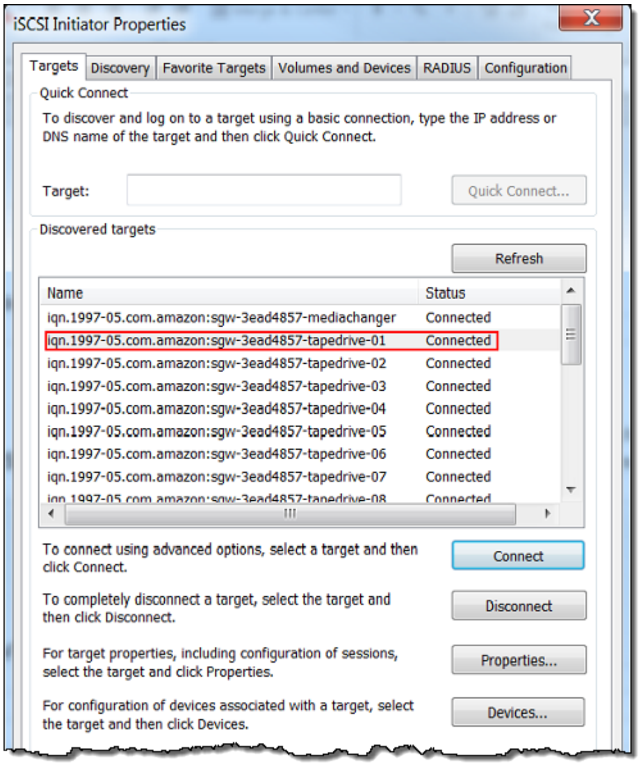

Quick Connect

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Target:

Discovered targets

Name	Status
iqn.1997-05.com.amazon:sgw-3ead4857-mediachanger	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-01	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-02	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-03	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-04	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-05	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-06	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-07	Connected
iqn.1997-05.com.amazon:sgw-3ead4857-tapedrive-08	Connected

To connect using advanced options, select a target and then click Connect.

To completely disconnect a target, select the target and then click Disconnect.

For target properties, including configuration of sessions, select the target and click Properties.

For configuration of devices associated with a target, select the target and then click Devices.

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 AWS 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 Recommended Red Hat Linux iSCSI Settings (p. 44).
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 enable 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
`username`	The initiator name that you found in a previous step in this procedure. The value starts with *iqn*. For example, **`iqn.1994-05.com.redhat:8e89b27b5b8`** is a valid *`username`* value.
`password`	The secret key used to authenticate the initiator (the client you are using) when it communicates with the volume.
`username_in`	The IQN of the target volume. The value starts with *iqn* and ends with the target name. For example, **`iqn.1997-05.com.amazon:myvolume`** is a valid *`username_in`* value.
`password_in`	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 Recommended Red Hat Linux iSCSI Settings (p. 44).

Understanding AWS Storage Gateway Resources and Resource IDs

In AWS 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:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*
Volume ARN	`arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*`/volume/`*`volume-id`*
Tape ARN	`arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:tape/`*`tapebarcode`*
Target ARN (iSCSI target)	`arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*`/target/`*`iSCSItarget`*
VTL Device ARN	`arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*`/device/`*`vtldevice`*

AWS Storage Gateway also supports the use of EC2 instances and EBS volumes and snapshots. These resources are Amazon EC2 resources that are used in AWS 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.

AWS 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.

> **Important**
> IDs for Storage Gateway volumes and Amazon EBS snapshots created from gateway volumes are changing to a longer format. Starting in December 2016, all new volumes and snapshots will be created with a 17-character string. Starting in April 2016, you will be able to use these longer IDs so you can test your systems with the new format. For more information, see Longer EC2 and EBS Resource IDs.
> For example, a volume ARN with the longer volume ID format will look like this:
> `arn:aws:storagegateway:us-west-2:111122223333:gateway/sgw-12A3456B/volume/vol-1122AABBCCDDEEFFG`.

A snapshot ID with the longer ID format will look like this: `snap-78e226633445566ee`. For more information, see Announcement: Heads-up – Longer AWS Storage Gateway volume and snapshot IDs coming in 2016.

Launching and Activating a Gateway's Amazon EC2 AMI in a Nondefault VPC

You can launch and activate your gateway Amazon Machine Image (AMI) into a default Amazon Virtual Private Cloud (Amazon VPC) or nondefault VPC by using the following procedure. To learn more about VPCs, see What is Amazon VPC? in the *Amazon VPC Getting Started Guide.*

To launch and activate an Amazon EC2 gateway AMI into a nondefault VPC

1. Set up the VPC and Internet gateway. Take note of your VPC ID. For information about how to create a VPC, see Getting Started with Amazon VPC in the *Amazon VPC Getting Started Guide.*

   A subnet and an Internet gateway are automatically created for your VPC. The subnet is associated with the Internet gateway. A routing table is also created for your VPC and is associated with the Internet gateway.
2. Create security group rules for your VPC. For more information, see Step 3: Set Up a Security Group for Your VPC in the *Amazon VPC Getting Started Guide.*
3. Create a network access control list (ACL) for your VPC:

   - For inbound traffic, create a rule to allow port 80 (HTTP), and set the source to `0.0.0.0/0`.
   - For inbound traffic, next create a rule to allow ephemeral ports. For more information, see Ephemeral Ports in the *Amazon VPC User Guide.*
   - For outbound traffic, create a rule to allow port 443 (HTTPS) and ports 1024–65535, set the destination to `0.0.0.0/0`, and then associate the ACL with your subnet.

     For more information, see Network ACLs in the *Amazon VPC Getting Started Guide.*
4. Launch your gateway AMI into the VPC.

   Note
   You must assign a public IP address to your instance. For more information, see Assigning a Public IP Address During Launch in the VPC User Guide.

To	See
Launch your volume gateway AMI into a VPC	Provisioning an Amazon EC2 Host (p. 29)
Launch your tape gateway AMI into a VPC	Provisioning an Amazon EC2 Host (p. 53)

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.

Topics

- Recovering Your Data from an Irrecoverable Tape (p. 293)
- Recovering Your Data from a Malfunctioning Cache Disk (p. 293)
- Recovering Your Data from a Corrupted File System (p. 293)

> **Important**
> AWS Storage Gateway doesn't support recovering a gateway VM from a snapshot that is created by your hypervisor. If your gateway VM malfunctions, activate a new gateway and recover your data to that gateway using the instructions following.

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 (p. 165).
- For cached volumes and tapes setups, when your gateway becomes reachable, your volumes or 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 BOOTSTRAPPING (p.).
- If your gateway malfunctions and issues occur with your volumes or tapes as a result of an unexpected shutdown, you can recover your data. For information about how to recover your data, see the sections following that apply to your scenario.

Recovering Your Data from a Malfunctioning Gateway or VM

If your gateway or virtual machine malfunctions, you can usually recover your data. For volume gateways, you recover data from a recovery snapshot. For tape gateways, you recover data by recovering one or more tapes from a recovery point to a new tape gateway.

If your gateway-cached volume becomes unreachable, you can use the following steps to recover your data from a recovery snapshot:

1. In the AWS Management Console, choose the malfunctioning gateway, choose the volume you want to recover, and then create a recovery snapshot from it.
2. Deploy and activate a new gateway-cached volume. Or, if you have an existing functioning gateway-cached volume, you can use that gateway and volume to recover your volume data.
3. Find the snapshot you created and restore it to a new volume on the functioning gateway.
4. Mount the new volume as an iSCSI device on your on-premises application server. You can access the data on the volume from this server. For gateways hosted on an Amazon Elastic Compute Cloud (Amazon EC2) instance, you can use the snapshot to create an Amazon Elastic Block Store (Amazon EBS) volume and attach it to an EC2 instance.

For detailed information on how to recover gateway-cached data from a recovery snapshot, see Using Recovery Snapshots for Your Gateway-Cached Setup (p. 201).

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 a tape gateway you want to use as the recovery target or create you can create a new one.
2. Disable the malfunctioning gateway.

3. Create recovery tapes for each tape 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 (p. 202).

Retrieving Your Data from an Irrecoverable Volume

If the status of your volume is IRRECOVERABLE, you can no longer use this volume. However, you can use the following steps to retrieve your data from the irrecoverable volume to a new volume:

1. Create a new volume from the disk that was used to create the irrecoverable volume.
2. Preserve existing data when you are creating the new volume.
3. Delete all pending snapshot jobs for the irrecoverable volume.
4. Delete the irrecoverable volume from the gateway.

For detailed information on how to retrieve your data from the irrecoverable volume to a new volume, see The Console Says That Your Volume Is Irrecoverable (p. 198).

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 (p. 204).

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 (p. 203).

Recovering Your Data from a Corrupted File System

If your file system gets corrupted, you can use the **fsck** command to check your file system for errors and repair it. If you can repair the file system, you can then recover your data from the volumes on the file system, as described following:

1. Shut down your virtual machine and use the AWS Storage Gateway Management Console to create a recovery snapshot. This snapshot represents the most current data stored in AWS.

> **Note**
> You use this snapshot as a fallback if your file system can't be repaired or the snapshot creation process can't be completed successfully.

For information about how to create a recovery snapshot, see Using Recovery Snapshots for Your Gateway-Cached Setup (p. 201).

2. Use the **fsck** command to check your file system for errors and attempt a repair.
3. Restart your gateway VM.
4. When your hypervisor host starts to boot up, press and hold down any key (for example, the spacebar) to enter the grub boot menu.
5. From the menu, choose the **CentOS** menu, and then press **`e`** to edit.
6. Choose the kernel line (the second line), and then press **`e`** to edit.
7. Append the following option to the kernel command line: **`init=/bin/bash`**. Use a space to separate the previous option from the option you just appended.
8. Press **`Return`** to save the changes.
9. Press **`b`** to boot your computer with the modified kernel option. Your computer will boot to a `bash#` prompt.
10. Type **`/sbin/fsck`** to run this command manually from the prompt, to check and repair your file system.
11. When the file system check and repair is complete, reboot the instance. The grub settings will revert to the original values, and the gateway will boot up normally.
12. Wait for snapshots that are in-progress from the original gateway to complete, and then validate the snapshot data.

You can continue to use the original volume as-is, or you can create a new gateway with a new volume based on either the recovery snapshot or the completed snapshot. Alternatively, you can create a new volume from any of your completed snapshots from this volume.

Tagging Storage Gateway Resources

In AWS 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 10.
- 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 AWS 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 AWS 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 AWS 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

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:

- https://s3.amazonaws.com/aws-storage-gateway-terms/sources.tar for gateways deployed on VMware ESXi
- https://s3.amazonaws.com/aws-storage-gateway-terms/sources_hyperv.tar for gateways deployed on Microsoft Hyper-V

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

The packages that make up the AWS Storage Gateway VM are tracked and monitored for security vulnerabilities. When updates are issued, they are applied to each gateway and the updated packages will increment their version number, although the major version number of the Linux distribution might not increment.

AWS Storage Gateway Limits

Following, you can find information about volume, virtual tape, configuration, and performance limits for Storage Gateway.

Limits For File Shares, Volumes and Tapes

The following lists file gateway limits.

- You can create 1 file share per bucket. There is a one-to-one mapping between a file share and a bucket.
- You can create up to 10 file shares per gateway.
- The maximum size of an individual file is 5 TB, which is the maximum size of an individual object in Amazon S3. If you write a file larger than 5 TB, you get a "file too large" error message and only the first 5 TB of the file is uploaded.

The following table lists limits for volumes and tapes.

	Cached Volumes	**Stored Volumes**	**Tape Gateway**
Maximum size of a volume	32 TiB	16 TiB	–
Maximum number of volumes for a gateway	32	32	–
Total size of all volumes for a gateway	1,024 TiB	512 TiB	–
Minimum size of a virtual tape	–	–	100 GiB
Maximum size of a virtual tape	–	–	2.5 TiB
Maximum number of virtual tapes for a virtual tape library (VTL)	–	–	1,500
Total size of all tapes in a virtual tape library (VTL)	–	–	1 PiB
Maximum number of virtual tapes in archive	–	–	No limit
Total size of all tapes in a archive	–	–	No limit

Note
If you create a snapshot from a cached volume that is more than 16 TiB in size, you cannot restore it to an Amazon Elastic Block Store (Amazon EBS) volume; however, it can be restored to a Storage Gateway volume. For more information, see Restoring a Snapshot to an Amazon EBS Volume (p. 265).

Configuration and Performance Limits

The following table lists limits for configuration and performance.

	Cached Volumes	Stored Volumes	Tape Gateway
Maximum size of a cache storage	16 TiB	–	16 TiB
Total maximum size of all cache storage for a gateway	16 TiB	–	16 TiB
Maximum size of an upload buffer disk	2 TiB	2 TiB	2 TiB
Total maximum size of all upload buffer disks for a gateway	2 TiB	2 TiB	2 TiB
Maximum upload rate	120 MB/s	120 MB/s	120 MB/s
Maximum download rate	20 MB/s	20 MB/s	20 MB/s

Note
The maximum upload rate was achieved by using 100 percent sequential write operations and 256 KB I/Os. Depending on your I/O mix and network conditions, the actual rate might be lower.

Authentication and Access Control for AWS Storage Gateway

Access to AWS Storage Gateway requires credentials that AWS can use to authenticate your requests. Those credentials must have permissions to access AWS resources, such as a gateway, volume, or tape. The following sections provide details on how you can use AWS Identity and Access Management (IAM) and AWS Storage Gateway to help secure your resources by controlling who can access them:

- Authentication (p. 297)
- Access Control (p. 299)

Authentication

You can access AWS as any of the following types of identities:

- **AWS account root user** – When you sign up for AWS, you provide an email address and password that is associated with your AWS account. These are your *root credentials* and they provide complete access to all of your AWS resources.

> **Important**
> For security reasons, we recommend that you use the root credentials only to create an *administrator user*, which is an *IAM user* with full permissions to your AWS account. Then, you can use this administrator user to create other IAM users and roles with limited permissions. For more information, see IAM Best Practices and Creating an Admin User and Group in the *IAM User Guide*.

- **IAM user** – An IAM user is simply an identity within your AWS account that has specific custom permissions (for example, permissions to create a gateway in AWS Storage Gateway). You can use an IAM user name and password to sign in to secure AWS webpages like the AWS Management Console, AWS Discussion Forums, or the AWS Support Center.

  In addition to a user name and password, you can also generate access keys for each user. You can use these keys when you access AWS services programmatically, either through one of the several SDKs or by using the AWS Command Line Interface (CLI). The SDK and CLI tools use the access keys to cryptographically sign your request. If you don't use the AWS tools, you must sign the request yourself. AWS Storage Gateway supports *Signature Version 4*, a protocol for authenticating inbound API requests. For more information about authenticating requests, see Signature Version 4 Signing Process in the *AWS General Reference*.

- **IAM role** – An IAM role is another IAM identity you can create in your account that has specific permissions. It is similar to an *IAM user*, but it is not associated with a specific person. An IAM role enables you to obtain temporary access keys that can be used to access AWS services and resources. IAM roles with temporary credentials are useful in the following situations:

  - **Federated user access** – Instead of creating an IAM user, you can use preexisting user identities from AWS Directory Service, your enterprise user directory, or a web identity provider. These are known as *federated users*. AWS assigns a role to a federated user when access is requested through an identity provider. For more information about federated users, see Federated Users and Roles in the *IAM User Guide*.

  - **Cross-account access** – You can use an IAM role in your account to grant another AWS account permissions to access your account's resources. For an example, see Tutorial: Delegate Access Across AWS Accounts Using IAM Roles in the *IAM User Guide*.

  - **AWS service access** – You can use an IAM role in your account to grant an AWS service permissions to access your account's resources. For example, you can create a role that allows Amazon Redshift to access an Amazon S3 bucket on your behalf and then load data stored in the bucket into an Amazon Redshift cluster. For more information, see Creating a Role to Delegate Permissions to an AWS Service in the *IAM User Guide*.

  - **Applications running on Amazon EC2** – Instead of storing access keys within the EC2 instance for use by applications running on the instance and making AWS API requests, you can use an IAM role to manage temporary credentials for these applications. To assign an AWS role to an EC2 instance and make it available to all of its applications, you can create an instance profile that is attached to the instance. An instance profile contains the role and enables programs running on the EC2 instance to get temporary credentials. For more information, see Using Roles for Applications on Amazon EC2 in the *IAM User Guide*.

Access Control

You can have valid credentials to authenticate your requests, but unless you have permissions you cannot create or access AWS Storage Gateway resources. For example, you must have permissions to create a gateway in AWS Storage Gateway.

The following sections describe how to manage permissions for AWS Storage Gateway. We recommend that you read the overview first.

- Overview of Managing Access Permissions to Your AWS Storage Gateway (p. 300)
- Identity-Based Policies (IAM Policies) (p. 301)

Overview of Managing Access Permissions to Your AWS Storage Gateway

Every AWS resource is owned by an AWS account, and permissions to create or access a resource are governed by permissions policies. An account administrator can attach permissions policies to IAM identities (that is, users, groups, and roles), and some services (such as AWS Lambda) also support attaching permissions policies to resources.

> **Note**
> An *account administrator* (or administrator user) is a user with administrator privileges. For more information, see IAM Best Practices in the IAM User Guide.

When granting permissions, you decide who is getting the permissions, the resources they get permissions for, and the specific actions that you want to allow on those resources.

Topics

AWS Storage Gateway Resources and Operations

In AWS Storage Gateway, the primary resource is a *gateway*. Storage Gateway also supports the following additional resource types: *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:region:account-id:gateway/gateway-id`
Volume ARN	`arn:aws:storagegateway:region:account-id:gateway/gateway-id/volume/volume-id`
Tape ARN	`arn:aws:storagegateway:region:account-id:tape/tapebarcode`
Target ARN (iSCSI target)	`arn:aws:storagegateway:region:account-id:gateway/gateway-id/target/iSCSItarget`
VTL Device ARN	`arn:aws:storagegateway:region:account-id:gateway/gateway-id/device/vtldevice`

> **Note**
>
> - AWS Storage Gateway resource IDs are in uppercase. 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.

- ARNs for gateways activated prior to September 2, 2015, contain the gateway name instead of the gateway ID. To obtain the ARN for your gateway, use the `DescribeGatewayInformation` API operation.

To grant permissions for specific API operations, such as creating a tape, AWS Storage Gateway provides a set of API actions for you to create and manage these resources and subresources. For a list of API actions, see Actions in the *AWS Storage Gateway API Reference*.

To grant permissions for specific API operations, such as creating a tape, AWS Storage Gateway defines a set of actions that you can specify in a permissions policy to grant permissions for specific API operations. An API operation can require permissions for more than one action. For a table showing all the AWS Storage Gateway API actions and the resources they apply to, see AWS Storage Gateway API Permissions: Actions, Resources, and Conditions Reference (p. 310).

Understanding Resource Ownership

A *resource owner* is the AWS account that created the resource. That is, the resource owner is the AWS account of the *principal entity* (the root account, an IAM user, or an IAM role) that authenticates the request that creates the resource. The following examples illustrate how this works:

- If you use the root account credentials of your AWS account to activate a gateway, your AWS account is the owner of the resource (in AWS Storage Gateway, the resource is the gateway).
- If you create an IAM user in your AWS account and grant permissions to the `ActivateGateway` action to that user, the user can activate a gateway. However, your AWS account, to which the user belongs, owns the gateway resource.
- If you create an IAM role in your AWS account with permissions to activate a gateway, anyone who can assume the role can activate a gateway. Your AWS account, to which the role belongs, owns the gateway resource.

Managing Access to Resources

A permissions policy describes who has access to what. The following section explains the available options for creating permissions policies.

> **Note**
> This section discusses using IAM in the context of AWS Storage Gateway. It doesn't provide detailed information about the IAM service. For complete IAM documentation, see What is IAM in the IAM User Guide. For information about IAM policy syntax and descriptions, see AWS IAM Policy Reference in the IAM User Guide.

Policies attached to an IAM identity are referred to as *identity-based* policies (IAM polices) and policies attached to a resource are referred to as *resource-based* policies. AWS Storage Gateway supports only identity-based policies (IAM policies).

Topics

Identity-Based Policies (IAM Policies)

You can attach policies to IAM identities. For example, you can do the following:

- **Attach a permissions policy to a user or a group in your account** – An account administrator can use a permissions policy that is associated with a particular user to grant permissions for that user to create an AWS Storage Gateway resource, such as a gateway, volume, or tape.
- **Attach a permissions policy to a role (grant cross-account permissions)** – You can attach an identity-based permissions policy to an IAM role to grant cross-account permissions. For example, the administrator in Account A can create a role to grant cross-account permissions to another AWS account (for example, Account B) or an AWS service as follows:
  1. Account A administrator creates an IAM role and attaches a permissions policy to the role that grants permissions on resources in Account A.
  2. Account A administrator attaches a trust policy to the role identifying Account B as the principal who can assume the role.
  3. Account B administrator can then delegate permissions to assume the role to any users in Account B. Doing this allows users in Account B to create or access resources in Account A. The principal in the trust policy can also be an AWS service principal if you want to grant an AWS service permissions to assume the role.

  For more information about using IAM to delegate permissions, see Access Management in the *IAM User Guide*.

The following is an example policy that grants permissions to all `List*` actions on all resources. This action is a read-only action. Thus, the policy doesn't allow the user to change the state of the resources.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowAllListActionsOnAllResources",
            "Effect": "Allow",
            "Action": [
                "storagegateway:List*"
            ],
            "Resource": "*"
        }
    ]
}
```

For more information about using identity-based policies with AWS Storage Gateway, see Using Identity-Based Policies (IAM Policies) for AWS Storage Gateway (p. 303). For more information about users, groups, roles, and permissions, see Identities (Users, Groups, and Roles in the *IAM User Guide*.

Resource-Based Policies

Other services, such as Amazon S3, also support resource-based permissions policies. For example, you can attach a policy to an S3 bucket to manage access permissions to that bucket. AWS Storage Gateway doesn't support resource-based policies.

Specifying Policy Elements: Actions, Effects, Resources, and Principals

For each AWS Storage Gateway resource (see AWS Storage Gateway API Permissions: Actions, Resources, and Conditions Reference (p. 310)), the service defines a set of API operations (see Actions). To grant permissions for these API operations, AWS Storage Gateway defines a set of actions that you can specify in a policy. For example, for the AWS Storage Gateway

gateway resource, the following actions are defined: `ActivateGateway`, `DeleteGateway`, and `DescribeGatewayInformation`. Note that, performing an API operation can require permissions for more than one action.

The following are the most basic policy elements:

- **Resource** – In a policy, you use an Amazon Resource Name (ARN) to identify the resource to which the policy applies. For AWS Storage Gateway resources, you always use the wildcard character `(*)` in IAM policies. For more information, see AWS Storage Gateway Resources and Operations (p. 300).
- **Action** – You use action keywords to identify resource operations that you want to allow or deny. For example, depending on the specified `Effect`, the `storagegateway:ActivateGateway` permission allows or denies the user permissions to perform the AWS Storage Gateway `ActivateGateway` operation.
- **Effect** – You specify the effect when the user requests the specific action—this can be either allow or deny. If you don't explicitly grant access to (allow) a resource, access is implicitly denied. You can also explicitly deny access to a resource, which you might do to make sure that a user cannot access it, even if a different policy grants access.
- **Principal** – In identity-based policies (IAM policies), the user that the policy is attached to is the implicit principal. For resource-based policies, you specify the user, account, service, or other entity that you want to receive permissions (applies to resource-based policies only). AWS Storage Gateway doesn't support resource-based policies.

To learn more about IAM policy syntax and descriptions, see AWS IAM Policy Reference in the *IAM User Guide*.

For a table showing all of the AWS Storage Gateway API actions, see AWS Storage Gateway API Permissions: Actions, Resources, and Conditions Reference (p. 310).

Specifying Conditions in a Policy

When you grant permissions, you can use the IAM policy language to specify the conditions when a policy should take effect when granting permissions. For example, you might want a policy to be applied only after a specific date. For more information about specifying conditions in a policy language, see Condition in the *IAM User Guide*.

To express conditions, you use predefined condition keys. There are no condition keys specific to Storage Gateway. However, there are AWS-wide condition keys that you can use as appropriate. For a complete list of AWS-wide keys, see Available Keys in the *IAM User Guide*.

Using Identity-Based Policies (IAM Policies) for AWS Storage Gateway

This topic provides examples of identity-based policies in which an account administrator can attach permissions policies to IAM identities (that is, users, groups, and roles).

> **Important**
> We recommend that you first review the introductory topics that explain the basic concepts and options available for you to manage access to your AWS Storage Gateway resources. For more information, see Overview of Managing Access Permissions to Your AWS Storage Gateway (p. 300).

The sections in this topic cover the following:

- Permissions Required to Use the Storage Gateway Console (p. 305)
- AWS Managed Policies for AWS Storage Gateway (p. 305)

- Customer Managed Policy Examples (p. 305)

The following shows an example of a permissions policy.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowsSpecifiedActionsOnAllGateways",
            "Effect": "Allow",
            "Action": [
                "storagegateway:ActivateGateway",
                "storagegateway:ListGateways"
            ],
            "Resource": "arn:aws:storagegateway:us-west-2:account-id:gateway/
*"
        },
        {
            "Sid": "AllowsSpecifiedEC2ActionsOnAllGateways",
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeSnapshots",
                "ec2:DeleteSnapshot"
            ],
            "Resource": "*"
        }
    ]
}
```

The policy has two statements (note the `Action` and `Resource` elements in both the statements):

- The first statement grants permissions for two Storage Gateway actions (`storagegateway:ActivateGateway` and `storagegateway:ListGateways`) on a gateway resource using the *Amazon Resource Name (ARN)* for the gateway. The ARN specifies a wildcard character (*) because you don't know the gateway ID until after you create a gateway.

  Note
  ARNs uniquely identify AWS resources. For more information, see Amazon Resource Names (ARNs) and AWS Service Namespaces in the *AWS General Reference*.

  The wildcard character (*) at the end of the gateway ARN means that this statement can match any gateway ID. In this case, the statement allows the `storagegateway:ActivateGateway` and `storagegateway:ListGateways` actions on any gateway in the specified region, `us-west-2`, and the specified ID identifies the account that is owner of the gateway resource. For information about how to use a wildcard character (*) in a policy, see Example 2: Allow Read-Only Access to a Gateway (p. 306).

  To limit permissions for a particular action to a specific gateway only, create a separate statement for that action in the policy and specify the gateway ID in that statement.

- The second statement grants permissions for the `ec2:DescribeSnapshots` and `ec2:DeleteSnapshot` actions. These Amazon Elastic Compute Cloud (Amazon EC2) actions require permissions because snapshots generated from AWS Storage Gateway are stored in Amazon Elastic Block Store (Amazon EBS) and managed as Amazon EC2 resources, and thus they require corresponding EC2 actions. For more information, see Actions in the *Amazon EC2 API Reference*. Because these Amazon EC2 actions don't support resource-level permissions, the policy specifies the wildcard character (*) as the `Resource` value instead of specifying a gateway ARN.

For a table showing all of the AWS Storage Gateway API actions and the resources that they apply to, see AWS Storage Gateway API Permissions: Actions, Resources, and Conditions Reference (p. 310).

Permissions Required to Use the Storage Gateway Console

To use the Storage Gateway console, you need to grant read-only permissions. If you plan to describe snapshots, you also need to grant permissions for additional actions as shown in the following permissions policy:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowsSpecifiedEC2ActionOnAllGateways",
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeSnapshots"
            ],
            "Resource": "*"
        }
    ]
}
```

This additional permission is required because the Amazon EBS snapshots generated from AWS Storage Gateway are managed as Amazon EC2 resources.

To set up the minimum permissions required to navigate the Storage Gateway console, see Example 2: Allow Read-Only Access to a Gateway (p. 306).

AWS Managed Policies for AWS Storage Gateway

AWS addresses many common use cases by providing standalone IAM policies that are created and administered by AWS. Managed policies grant necessary permissions for common use cases so you can avoid having to investigate what permissions are needed. For more information about AWS managed policies, see AWS Managed Policies in the *IAM User Guide*.

The following AWS managed policies, which you can attach to users in your account, are specific to Storage Gateway:

- **AWSStorageGatewayReadOnlyAccess** – Grants read-only access to AWS Storage Gateway resources.
- **AWSStorageGatewayFullAccess** – Grants full access to AWS Storage Gateway resources.

> **Note**
> You can review these permissions policies by signing in to the IAM console and searching for specific policies there.

You can also create your own custom IAM policies to allow permissions for AWS Storage Gateway API actions. You can attach these custom policies to the IAM users or groups that require those permissions.

Customer Managed Policy Examples

In this section, you can find example user policies that grant permissions for various Storage Gateway actions. These policies work when you are using AWS SDKs and the AWS CLI. When you are using the console, you need to grant additional permissions specific to the console, which is discussed in Permissions Required to Use the Storage Gateway Console (p. 305).

> **Note**
>
> All examples use the US West (Oregon) Region (`us-west-2`) and contain fictitious account IDs.

Topics

Example 1: Allow Any AWS Storage Gateway Actions on All Gateways

The following policy allows a user to perform all the AWS Storage Gateway actions. The policy also allows the user to perform Amazon EC2 actions (DescribeSnapshots and DeleteSnapshot) on the Amazon EBS snapshots generated from AWS Storage Gateway.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowsAllAWSStorageGatewayActions",
            "Action": [
                "storagegateway:*"
            ],
            "Effect": "Allow",
            "Resource": "*"
        },
        {
            "Sid": "AllowsSpecifiedEC2Actions",
            "Action": [
                "ec2:DescribeSnapshots",
                "ec2:DeleteSnapshot"
            ],
            "Effect": "Allow",
            "Resource": "*"
        }
    ]
}
```

Example 2: Allow Read-Only Access to a Gateway

The following policy allows all `List*` and `Describe*` actions on all resources. Note that these actions are read-only actions. Thus, the policy doesn't allow the user to change the state of any resources—that is, the policy doesn't allow the user to perform actions such as `DeleteGateway`, `ActivateGateway`, and `ShutdownGateway`.

The policy also allows the `DescribeSnapshots` Amazon EC2 action. For more information, see DescribeSnapshots in the *Amazon EC2 API Reference*.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowReadOnlyAccessToAllGateways",
            "Action": [
                "storagegateway:List*",
                "storagegateway:Describe*"
            ],
            "Effect": "Allow",
            "Resource": "*"
        },
        {
            "Sid": "AllowsUserToDescribeSnapshotsOnAllGateways",
            "Action": [
                "ec2:DescribeSnapshots"
            ],
            "Effect": "Allow",
            "Resource": "*"
        }
    ]
}
```

In the preceding policy, instead of a using a wildcard character (*), you can scope resources covered by the policy to a specific gateway, as shown in the following example. The policy then allows the actions only on the specific gateway.

```
"Resource": [
              "arn:aws:storagegateway:us-west-2:123456789012:gateway/gateway-id/",
              "arn:aws:storagegateway:us-west-2:123456789012:gateway/gateway-id/*"
            ]
```

Within a gateway, you can further restrict the scope of the resources to only the gateway volumes, as shown in the following example:

```
"Resource": "arn:aws:storagegateway:us-west-2:123456789012:gateway/gateway-id/volume/*"
```

Example 3: Allow Access to a Specific Gateway

The following policy allows all actions on a specific gateway. The user is restricted from accessing other gateways you might have deployed.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowReadOnlyAccessToAllGateways",
            "Action": [
                "storagegateway:List*",
                "storagegateway:Describe*"
            ],
            "Effect": "Allow",
            "Resource": "*"
        },
        {
            "Sid": "AllowsUserToDescribeSnapshotsOnAllGateways",
            "Action": [
                "ec2:DescribeSnapshots"
            ],
            "Effect": "Allow",
            "Resource": "*"
        },
        {
            "Sid": "AllowsAllActionsOnSpecificGateway",
            "Action": [
                "storagegateway:*"
            ],
            "Effect": "Allow",
            "Resource": [
                "arn:aws:storagegateway:us-
west-2:123456789012:gateway/gateway-id/",
                "arn:aws:storagegateway:us-
west-2:123456789012:gateway/gateway-id/*"
            ]
        }
    ]
}
```

The preceding policy works if the user to which the policy is attached uses either the API or an AWS SDK to access the gateway. However, if the user is going to use the AWS Storage Gateway console, you must also grant permissions to allow the `ListGateways` action, as shown in the following example:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowsAllActionsOnSpecificGateway",
            "Action": [
                "storagegateway:*"
            ],
            "Effect": "Allow",
            "Resource": [
                "arn:aws:storagegateway:us-
west-2:123456789012:gateway/gateway-id/",
                "arn:aws:storagegateway:us-
west-2:123456789012:gateway/gateway-id/*"
            ]
        },
        {
            "Sid": "AllowsUserToUseAWSConsole",
            "Action": [
                "storagegateway:ListGateways"
            ],
            "Effect": "Allow",
            "Resource": "*"
        }
    ]
}
```

Example 4: Allow a User to Access a Specific Volume

The following policy allows a user to perform all actions to a specific volume on a gateway. Because a user doesn't get any permissions by default, the policy restricts the user to accessing only a specific volume.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "GrantsPermissionsToSpecificVolume",
            "Action": [
                "storagegateway:*"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:storagegateway:us-
west-2:123456789012:gateway/gateway-id/volume/volume-id"
        },
        {
            "Sid": "GrantsPermissionsToUseStorageGatewayConsole",
            "Action": [
                "storagegateway:ListGateways"
            ],
            "Effect": "Allow",
            "Resource": "*"
        }
    ]
}
```

The preceding policy works if the user to whom the policy is attached uses either the API or an AWS SDK to access the volume. However, if this user is going to use the AWS Storage Gateway console, you must also grant permissions to allow the `ListGateways` action, as shown in the following example:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "GrantsPermissionsToSpecificVolume",
            "Action": [
                "storagegateway:*"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:storagegateway:us-
west-2:123456789012:gateway/gateway-id/volume/volume-id"
        },
        {
            "Sid": "GrantsPermissionsToUseStorageGatewayConsole",
            "Action": [
                "storagegateway:ListGateways"
            ],
            "Effect": "Allow",
            "Resource": "*"
        }
    ]
}
```

Example 5: Allow All Actions on Gateways with a Specific Prefix

The following policy allows a user to perform all AWS Storage Gateway actions on gateways with names that start with `DeptX`. The policy also allows the `DescribeSnapshots` Amazon EC2 action which is required if you plan to describe snapshots.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowsActionsGatewayWithPrefixDeptX",
            "Action": [
                "storagegateway:*"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:storagegateway:us-
west-2:123456789012:gateway/DeptX"
        },
        {
            "Sid": "GrantsPermissionsToSpecifiedAction",
            "Action": [
                "ec2:DescribeSnapshots"
            ],
            "Effect": "Allow",
            "Resource": "*"
        }
    ]
}
```

The preceding policy works if the user to whom the policy is attached uses either the API or an AWS SDK to access the gateway. However, if this user plans to use the AWS Storage Gateway console, you must grant additional permissions as described in Example 3: Allow Access to a Specific Gateway (p. 307).

AWS Storage Gateway API Permissions: Actions, Resources, and Conditions Reference

When you are setting up Access Control (p. 299) and writing permissions policies that you can attach to an IAM identity (identity-based policies), you can use the following table as a reference. The table lists each AWS Storage Gateway API operation, the corresponding actions for which you can grant permissions to perform the action, and the AWS resource for which you can grant the permissions. You specify the actions in the policy's `Action` field, and you specify the resource value in the policy's `Resource` field.

You can use AWS-wide condition keys in your AWS Storage Gateway policies to express conditions. For a complete list of AWS-wide keys, see Available keys in the *IAM User Guide*.

> **Note**
> To specify an action, use the `storagegateway:` prefix followed by the API operation name (for example, `storagegateway:ActivateGateway`). For each AWS Storage Gateway action, you can specify a wildcard character (*) as the resource.

For a list of Storage Gateway resources with the ARN format, see AWS Storage Gateway Resources and Operations (p. 300).

AWS Storage Gateway API and Required Permissions for Actions

ActivateGateway

Action(s): `storagegateway:ActivateGateway`

Resource: `*`

AddCache

Action(s): `storagegateway:AddCache`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

AddTagsToResource

Action(s): storagegateway:AddTagsToResource

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

or

arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*/volume/*volume-id*

or

arn:aws:storagegateway:*region*:*account-id*:tape/*tapebarcode*

AddUploadBuffer

Action(s): storagegateway:AddUploadBuffer

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

AddWorkingStorage

Action(s): storagegateway:AddWorkingStorage

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

CancelArchival

Action(s): storagegateway:CancelArchival

Resource: arn:aws:storagegateway:*region*:*account-id*:tape/*tapebarcode*

CancelRetrieval

Action(s): storagegateway:CancelRetrieval

Resource: arn:aws:storagegateway:*region*:*account-id*:tape/*tapebarcode*

CreateCachediSCSIVolume

Action(s): storagegateway:CreateCachediSCSIVolume

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

CreateSnapshot

Action(s): storagegateway:CreateSnapshot

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*/volume/*volume-id*

CreateSnapshotFromVolumeRecoveryPoint

Action(s): storagegateway:CreateSnapshotFromVolumeRecoveryPoint

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*/volume/*volume-id*

CreateStorediSCSIVolume

Action(s): storagegateway:CreateStorediSCSIVolume

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

CreateTapes

Action(s): storagegateway:CreateTapes

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

DeleteBandwidthRateLimit

Action(s): storagegateway:DeleteBandwidthRateLimit

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

DeleteChapCredentials

Action(s): storagegateway:DeleteChapCredentials

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*`/target/`*`iSCSItarget`*

DeleteGateway

Action(s): `storagegateway:DeleteGateway`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

DeleteSnapshotSchedule

Action(s): `storagegateway:DeleteSnapshotSchedule`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*`/volume/`*`volume-id`*

DeleteTape

Action(s): `storagegateway:DeleteTape`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

DeleteTapeArchive

Action(s): `storagegateway:DeleteTapeArchive`

Resource: *

DeleteVolume

Action(s): `storagegateway:DeleteVolume`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*`/volume/`*`volume-id`*

DescribeBandwidthRateLimit

Action(s): `storagegateway:DescribeBandwidthRateLimit`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

DescribeCache

Action(s): `storagegateway:DescribeCache`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

DescribeCachediSCSIVolumes

Action(s): `storagegateway:DescribeCachediSCSIVolumes`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*`/volume/`*`volume-id`*

DescribeChapCredentials

Action(s): `storagegateway:DescribeChapCredentials`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*`/target/`*`iSCSItarget`*

DescribeGatewayInformation

Action(s): `storagegateway:DescribeGatewayInformation`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

DescribeMaintenanceStartTime

Action(s): `storagegateway:DescribeMaintenanceStartTime`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

DescribeSnapshotSchedule

Action(s): `storagegateway:DescribeSnapshotSchedule`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*`/volume/`*`volume-id`*

DescribeStorediSCSIVolumes

Action(s): storagegateway:DescribeStorediSCSIVolumes

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*/volume/*volume-id*

DescribeTapeArchives

Action(s): storagegateway:DescribeTapeArchives

Resource: *

DescribeTapeRecoveryPoints

Action(s): storagegateway:DescribeTapeRecoveryPoints

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

DescribeTapes

Action(s): storagegateway:DescribeTapes

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

DescribeUploadBuffer

Action(s): storagegateway:DescribeUploadBuffer

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

DescribeVTLDevices

Action(s): storagegateway:DescribeVTLDevices

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

DescribeWorkingStorage

Action(s): storagegateway:DescribeWorkingStorage

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

DisableGateway

Action(s): storagegateway:DisableGateway

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

ListGateways

Action(s): storagegateway:ListGateways

Resource: *

ListLocalDisks

Action(s): storagegateway:ListLocalDisks

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

ListTagsForResource

Action(s): storagegateway:ListTagsForResource

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

or

arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*/volume/*volume-id*

or

arn:aws:storagegateway:*region*:*account-id*:tape/*tapebarcode*

ListTapes

Action(s): storagegateway:ListTapes

Resource: arn:aws:storagegateway:*region*:*account-id*:gateway/*gateway-id*

ListVolumeInitiators

Action(s): `storagegateway:ListVolumeInitiators`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*`/volume/`*`volume-id`*

ListVolumeRecoveryPoints

Action(s): `storagegateway:ListVolumeRecoveryPoints`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

ListVolumes

Action(s): `storagegateway:ListVolumes`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

RemoveTagsFromResource

Action(s): `storagegateway:RemoveTagsFromResource`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

or

`arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*`/volume/`*`volume-id`*

or

arn:aws:storagegateway`:`*`region`*`:`*`account-id`*`:tape/`*`tapebarcode`*

ResetCache

Action(s): `storagegateway:ResetCache`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

RetrieveTapeArchive

Action(s): `storagegateway:RetrieveTapeArchive`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

RetrieveTapeRecoveryPoint

Action(s): `storagegateway:RetrieveTapeRecoveryPoint`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

ShutdownGateway

Action(s): `storagegateway:ShutdownGateway`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

StartGateway

Action(s): `storagegateway:StartGateway`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

UpdateBandwidthRateLimit

Action(s): `storagegateway:UpdateBandwidthRateLimit`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*

UpdateChapCredentials

Action(s): `storagegateway:UpdateChapCredentials`

Resource: `arn:aws:storagegateway:`*`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*`/target/`*`iSCSItarget`*

UpdateGatewayInformation

Action(s): `storagegateway:UpdateGatewayInformation`

Resource: `arn:aws:storagegateway:region:account-id:gateway/gateway-id`

UpdateGatewaySoftwareNow
: **Action(s):** `storagegateway:UpdateGatewaySoftwareNow`

  Resource: `arn:aws:storagegateway:region:account-id:gateway/gateway-id`

UpdateMaintenanceStartTime
: **Action(s):** `storagegateway:UpdateMaintenanceStartTime`

  Resource: `arn:aws:storagegateway:region:account-id:gateway/gateway-id`

UpdateSnapshotSchedule
: **Action(s):** `storagegateway:UpdateSnapshotSchedule`

  Resource: `arn:aws:storagegateway:region:account-id:gateway/gateway-id/volume/volume-id`

UpdateVTLDeviceType
: **Action(s):** `storagegateway:UpdateVTLDeviceType`

  Resource: `arn:aws:storagegateway:region:account-id:gateway/gateway-id/device/vtldevice`

Related Topics

- Access Control (p. 299)
- Customer Managed Policy Examples (p. 305)

Tutorial: AWS Storage Gateway Virtual tape Library

In this tutorial, you can find instructions about how to download, deploy, and use a tape gateway. At several steps, you perform tasks outside of the AWS Storage Gateway management console and then return.

Topics

Sign Up for AWS Storage Gateway

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

To sign up for AWS account

1. Open http://aws.amazon.com/, and then choose **Create an AWS Account**.
2. Follow the online instructions.

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

Step 1: Select Gateway Type

After you sign up for AWS, the first thing you decide is the type of gateway you want to create. In this tutorial, you will create a Gateway–virtual tape library (VTL) (p. 6). In tape gateway, you store and archive your data on virtual tapes in AWS and also eliminate the challenges associated with owning and operating an on-premises physical tape infrastructure.

To create a gateway

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Gateways**, and then choose **Create Gateway**.

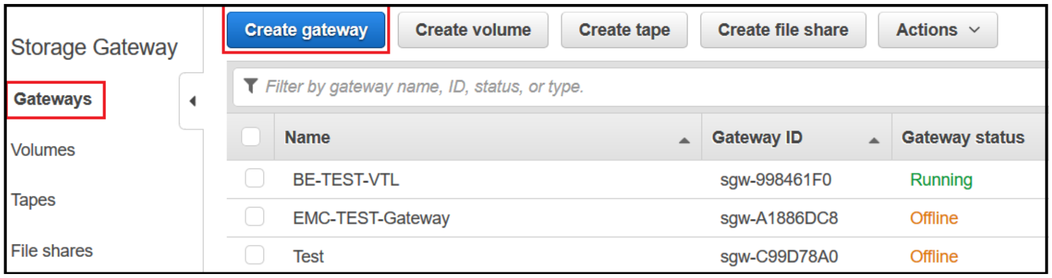

Name	Gateway ID	Gateway status
BE-TEST-VTL	sgw-998461F0	Running
EMC-TEST-Gateway	sgw-A1886DC8	Offline
Test	sgw-C99D78A0	Offline

3. On the **Select Gateway Type** page, choose **Virtual tape library**, and then choose **Next**.

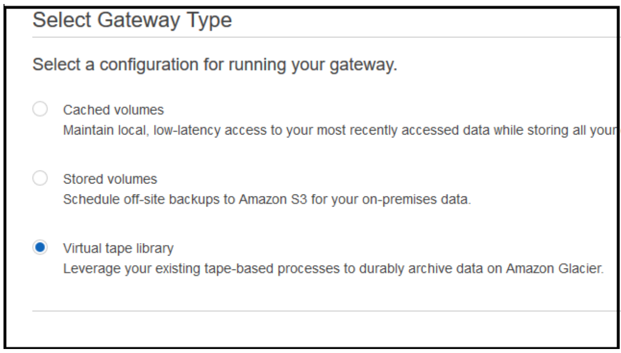

Select Gateway Type

Select a configuration for running your gateway.

○ Cached volumes
Maintain local, low-latency access to your most recently accessed data while storing all your

○ Stored volumes
Schedule off-site backups to Amazon S3 for your on-premises data.

● Virtual tape library
Leverage your existing tape-based processes to durably archive data on Amazon Glacier.

Next Step

- Step 2: Choose Host Platform (p. 316)

Step 2: Choose Host Platform

You can choose to run AWS Storage Gateway either on-premises as a virtual machine (VM) appliance on VMware ESXi or Microsoft Hyper-V, or in AWS as an Amazon Elastic Compute Cloud (Amazon EC2) instance. For information about supported host platforms, see Supported Hypervisors and Host

Requirements (p. 16). If you are creating the gateway on-premises, you download and deploy the gateway VM and then activate the gateway.

> **Note**
> You can also create tape gateway on an Amazon EC2 instance. You launch an Amazon Machine Image (AMI) that contains the gateway VM image and then activate the gateway.

To choose a host platform

1. On the **Choose host platform** page, choose the virtualization platform you want to run your gateway on.

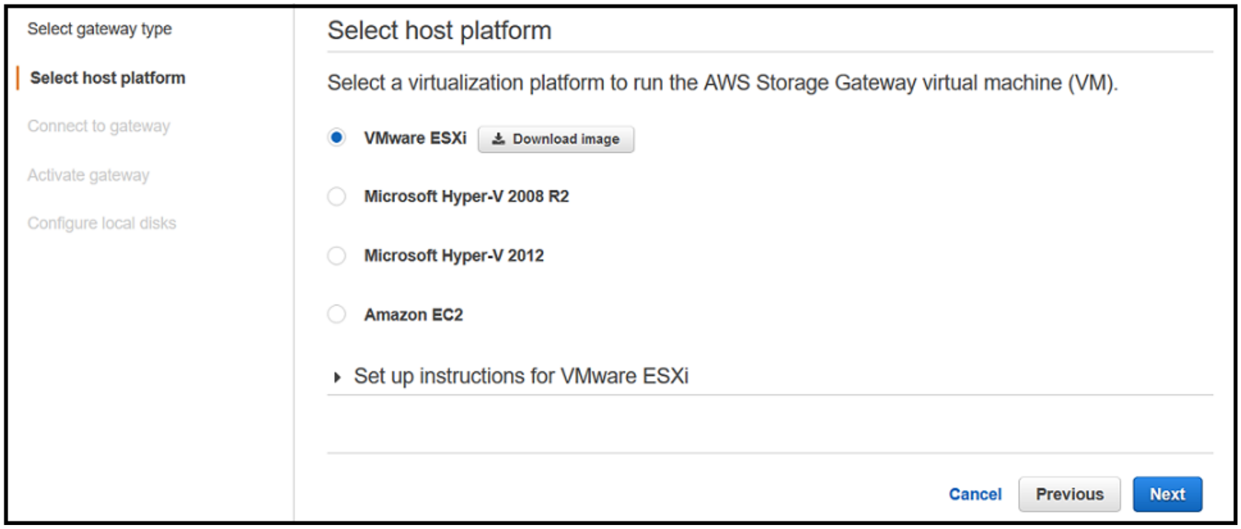

Select host platform

Select a virtualization platform to run the AWS Storage Gateway virtual machine (VM).

● VMware ESXi

○ Microsoft Hyper-V 2008 R2

○ Microsoft Hyper-V 2012

○ Amazon EC2

▸ Set up instructions for VMware ESXi

2. Choose **Download** *Your Virtualization platform* to download a .zip file that contains the .ova file for your virtualization platform.

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

Now that you have downloaded the gateway VM, you provision a host and deploy your gateway VM you downloaded. If you haven't provisioned a host yet, provision it now. You perform this task outside the AWS Storage Gateway console and return when you are done.

Depending on the host platform you chose to run your gateway on, choose one of the following.

Next Step

- Provision a VMware Host (p. 317)
- Provision a Hyper-V Host (p. 319)

Provision a VMware Host

Following, you can find instructions how to provision an on-premises VMware host and deploy your gateway. You perform this task in the VMware vSphere client and return to the AWS Storage Gateway console to continue your gateway setup.

To provision a VMware host

1. Review the minimum host requirements in Requirements (p. 11).

2. Provision a host in your data center with the VMware ESXi hypervisor. For a minimum set of instructions to install the hypervisor, see Configuring a VMware ESXi Host for AWS Storage Gateway (p. 208).

   Note
   If you plan to deploy AWS Storage Gateway using VMware High Availability (HA) for failover protection, see Using AWS Storage Gateway with VMware High Availability (p. 319). In this tutorial exercise, you deploy your AWS Storage Gateway VM on a single host with no clustering or failover.

To deploy your gateway on a VMware host

1. Connect to your gateway host's hypervisor by using your VMware vSphere client.
2. Deploy the OVF template package that you downloaded.
3. Choose **Next** through the following three screens. You might be prompted to select a data store on which to store the `.ova` package.
4. Allocate disks with **Thick provisioned format**.
5. Synchronize the time of your gateway VM to match your gateway host's time.

   You must also ensure that the clock on your gateway host is synchronized with a Network Time Protocol (NTP) server.

This procedure provides minimal instructions to help you deploy your gateway quickly. If you are not familiar with using the VMware vSphere client software, see Deploying the AWS Storage Gateway VM to Your VMWare Host (p. 212) for detailed instructions.

To allocate a local disk from direct-attached storage

Next, you allocate local disks for your deployed gateway VM. After you activate your gateway, you configure the disks as local storage for the gateway's use.

> **Important**
> All gateways require a disk of a minimum size of 150 GiB for use as an upload buffer. Gateway-cached and tape gateway setups require additional disks for use as cache storage. The minimum size of a cache disk is 1.1 times your upload buffer size. For more information, see AWS Storage Gateway Limits (p. 296).

1. Decide the number and size of disks to allocate for your gateway.
2. Start your VMware vSphere client, and then connect to your host.
3. In the client, follow the instructions provided by your host's client and add two virtual disks.
4. Configure the disks according to the sizes you decided for your gateway.

   For your disk configuration, we recommend thick provisioning. We also strongly recommend that you choose the **Specify a datastore or datastore cluster** option and select a data store for these disks different from the one that is used for the gateway VM. In addition, we strongly recommend that you dedicate one data store for the disks that you will use for the upload buffer and a different data store for other disks.
5. Configure your VM to use paravirtualized controllers.

This procedure provides minimal instructions to help you to quickly allocate disks for your gateway. If you are not familiar with using the VMware vSphere client software, see Provisioning Local Disk Storage for the Gateway VM (VMWare) (p. 218) for detailed instructions.

Using AWS 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 AWS Storage Gateway with VMware HA, we recommend doing the following things:

- Deploy the VMware ESX `.ova` downloadable package that contains the AWS 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 (p. 280). For more information on customizing Linux clients' timeout settings, see Customizing Your Linux iSCSI Settings (p. 45).
- 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.

Next Step

- Step 3: Begin Activation (p. 321)

Provision a Hyper-V Host

Following, you can find instructions how to provision an on-premises Microsoft Hyper-V host and deploy your gateway. You perform this task in the Hyper-V Manager and return to the AWS Storage Gateway console to continue your gateway setup.

To provision a Hyper-V host

1. Review the minimum host requirements in Requirements (p. 11).
2. Set up a host in your data center with the Microsoft Hyper-V host. For a minimum set of instructions to install the hypervisor, see Configuring a Hyper-V Host for AWS Storage Gateway (p. 227).

To deploy your gateway on a Hyper-V host

1. Connect to the Microsoft Hyper-V Manager on your Windows client.
2. Create locations on the hypervisor host for the gateway virtual hard disks and VM.

   a. Navigate to the hypervisor drive.

   b. Create a folder with two subfolders, `unzippedSourceVM` and `gateway`.
3. Configure the Hyper-V Manager to point to the `gateway` folder you created. The running VM stores its configuration in this folder.
4. Copy the unzipped source VM files to the folder you created on the host computer. Import the AWS Storage Gateway VM to the host. You must have 75 GiB of disk space for installation of the VM image and system data.

5. Rename the VM to avoid confusion with other VMs that you might import to the host.
6. Confirm that **Time synchronization** is selected for the VM.

   You must also ensure that the clock on your gateway host is synchronized with a Network Time Protocol (NTP) server.

This procedure provides minimal instructions to help you deploy your gateway quickly. If you are not familiar with using the Microsoft Hyper-V manager software, see Deploying a AWS Storage Gateway VM on a Microsoft Hyper-V Host (p. 236) for detailed instructions.

To allocate a local disk from direct-attached storage

Next, you allocate local disks for your deployed gateway VM. After you activate your gateway, you configure the disk as local storage for the gateway's use.

> **Important**
> All gateways require a disk of a minimum size of 150 GiB for use as an upload buffer. Gateway-cached and tape gateway setups require additional disks for use as cache storage. The minimum size of a cache disk is 1.1 times your upload buffer size. For more information, see AWS Storage Gateway Limits (p. 296)

1. Decide the number and size of disks to allocate for your gateway.
2. Start the Microsoft Hyper-V Manager, and then connect to the hypervisor.
3. In the Hyper-V Manager, follow the instructions provided by your host's client and add two virtual disks.
4. Configure the disks according to the sizes that you decided for your gateway.

   For your disk configuration, we recommend thick provisioning. We also strongly recommend that you choose the **Specify a datastore or datastore cluster** option and select a data store for these disks different from the one that is used for the gateway VM. In addition, we strongly recommend that you dedicate one data store for the disks that you will use for the upload buffer and a different data store for other disks.
5. > **Note**
   > AWS Storage Gateway supports the .vhdx file type. This file type enables you to create larger virtual disks than other file types. If you create a .vhdx type virtual disk, make sure that the size of the virtual disks you create does not exceed the recommended disk size for your gateway.

   If you are using the Microsoft Hyper-V 2012 Hypervisor, you will be prompted to choose a disk format (**VHD** or **VHDX**).
6. Configure the disks as **Fixed size** for **Disk Type**.

   When you use fixed-size provisioning, the disk storage is allocated immediately, resulting in better performance. If you don’t use fixed-size provisioning, the storage is allocated on demand. On-demand allocation can have a negative impact on the functioning of AWS Storage Gateway. For AWS Storage Gateway to function properly, the VM disks must be stored in fixed-size provisioned format.

This procedure provides minimal instructions to help you quickly allocate disks. If you are not familiar with using the Microsoft Hyper-V manager software, see Provision Local Storage for the AWS Storage Gateway VM (Hyper-V) (p. 245) for detailed instructions.

Next Step

- Step 3: Begin Activation (p. 321)

Step 3: Begin Activation

The first step is to get to the IP address of your gateway VM. You will need this IP address for activation. For gateways deployed and activated on an on-premises host, you can get the IP address from your hypervisor client or your gateway VM local console. The activation process associates your gateway with your AWS account. You gateway VM must be running for activation to succeed.

To get the IP address for you gateway VM from the local console

1. On your gateway VM local console, navigate to the **AWS Storage Gateway Configuration** main menu, type option **2** to open the **Network configuration** menu. For detailed instructions, see Configuring Your Gateway Network (p. 162).
2. Type option **1** (Describe Adapter) to display the adapter information including the gateway IP address.
3. Take note of the Gateway IP address.

For activation, you can use the public or private IP address assigned to the gateway. You must be able to reach the IP address that you use from the browser from which you perform the activation. In this walkthrough, you will use the public IP address to activate the gateway.

To associate your gateway with your AWS account

1. Open the AWS Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. If the **Begin activation** page isn't already open, choose **Create Gateway**, select a gateway on the **Select gateway** page, and then choose **Next**.
3. On the **Choose host platform** page, choose the host platform that you want to run your gateway on, and then choose **Next**.
4. Type the IP address of your gateway for **IP address**, and then choose the **Begin activation** button.

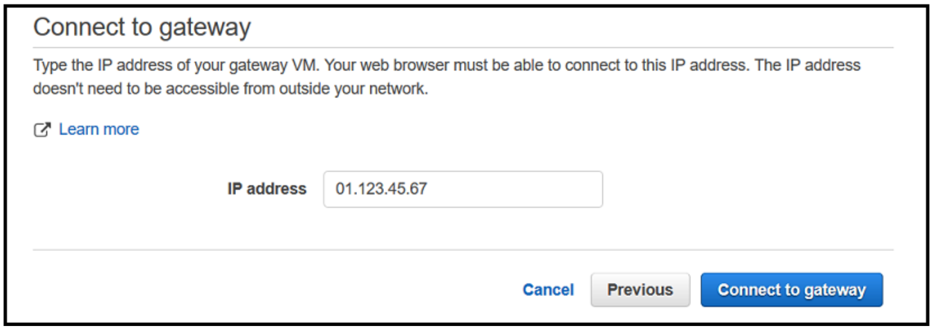

Connect to gateway

Type the IP address of your gateway VM. Your web browser must be able to connect to this IP address. The IP address doesn't need to be accessible from outside your network.

Learn more

IP address: 01.123.45.67

Next Step

- Step 4: Activate Your Gateway (p. 321)

Step 4: Activate Your Gateway

When your gateway VM is deployed and running, you configure your gateway settings and complete the activation process. If activation fails, check that the IP address you entered is correct. If the IP address is correct, confirm that your network is configured to let your browser access the gateway VM.

For more information, see Troubleshooting On-Premises Gateway Issues (p. 191) for troubleshooting guidelines.

To configure your gateway settings

1. To complete the activation process, enter the information listed on the activation page to configure your gateway setting.

   The following screenshot shows the activation page for tape gateways.

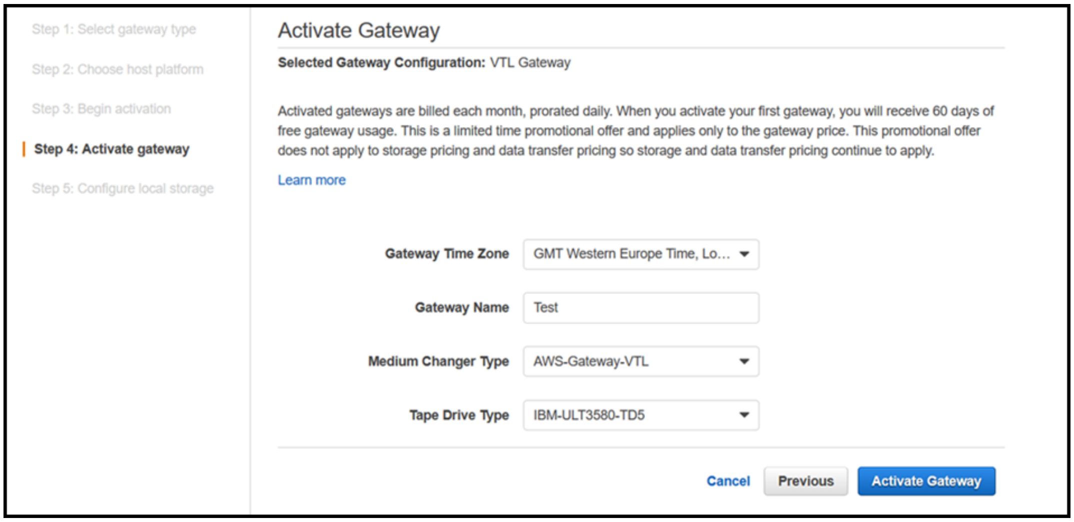
Step 1: Select gateway type
Step 2: Choose host platform
Step 3: Begin activation
Step 4: Activate gateway
Step 5: Configure local storage
Cancel
Previous
Activate Gateway

Activate Gateway

Selected Gateway Configuration: VTL Gateway

Activated gateways are billed each month, prorated daily. When you activate your first gateway, you will receive 60 days of free gateway usage. This is a limited time promotional offer and applies only to the gateway price. This promotional offer does not apply to storage pricing and data transfer pricing so storage and data transfer pricing continue to apply.

Learn more

Gateway Time Zone GMT Western Europe Time, Lo...

Gateway Name Test

Medium Changer Type AWS-Gateway-VTL

Tape Drive Type IBM-ULT3580-TD5

- **Gateway Time Zone** specifies the time zone to use for your gateway.
- **Gateway Name** identifies your gateway. You use this name to manage your gateway in the console; you can change it after the gateway is activated. This name must be unique to your account.
- **Medium Changer Type** specifies the type of medium changer to use for your backup software (available for tape gateways only).

  Important
  The type of medium changer you select depends on the backup software you plan to use. The following table shows which medium changer to select for your backup software. This list includes third-party backup software that has been tested and found to be compatible with tape gateway.

Backup Software	Medium Changer Type
Backup Exec 2012	`STK-L700`
Backup Exec 2014	`AWS-Gateway-VTL`
Backup Exec 15	`AWS-Gateway-VTL`
Backup Exec 16	`AWS-Gateway-VTL`
Dell NetVault Backup 10.0	`STK-L700`
EMC NetWorker 8.x	`STK-L700`
HPE Data Protector 9.x	`AWS-Gateway-VTL`
Microsoft System Center 2012 R2 Data Protection Manager	`STK-L700`

Backup Software	Medium Changer Type
Symantec NetBackup Version 7.x	`AWS-Gateway-VTL`
Veeam Backup & Replication V7	`STK-L700`
Veeam Backup & Replication V8	`STK-L700`
Veeam Backup & Replication V9 Update 2 or later	`AWS-Gateway-VTL`

Note
You must select the medium changer that is recommended for your backup software. Other medium changers might not function properly. You can select a different medium changer after the gateway is activated. For more information, see Selecting a Medium Changer After Gateway Activation (p. 273).

- **Tape Drive Type** specifies the type of tape drive used by this gateway (available for tape gateways only).

2. Choose **Activate Gateway**.
3. When the gateway is successfully activated, the AWS Storage Gateway console displays the **Configure local storage** page.

   If activation is not successful, see Troubleshooting Your Gateway (p. 191) for possible solutions.

Next Step

- Step 5: Configure Local Storage (p. 323)

Step 5: Configure Local Storage

When you deployed the VM, you allocated local disks for your gateway. Now, you will configure your gateway to use these disks. You configure at least one disk for an upload buffer and the other for cache storage.

To configure local storage

1. On the **Configure Local Storage** page, identify the disks you allocated and decide which ones you want to use for upload buffer and cached storage. For information about disk size limits, see Configuration and Performance Limits (p. 297).

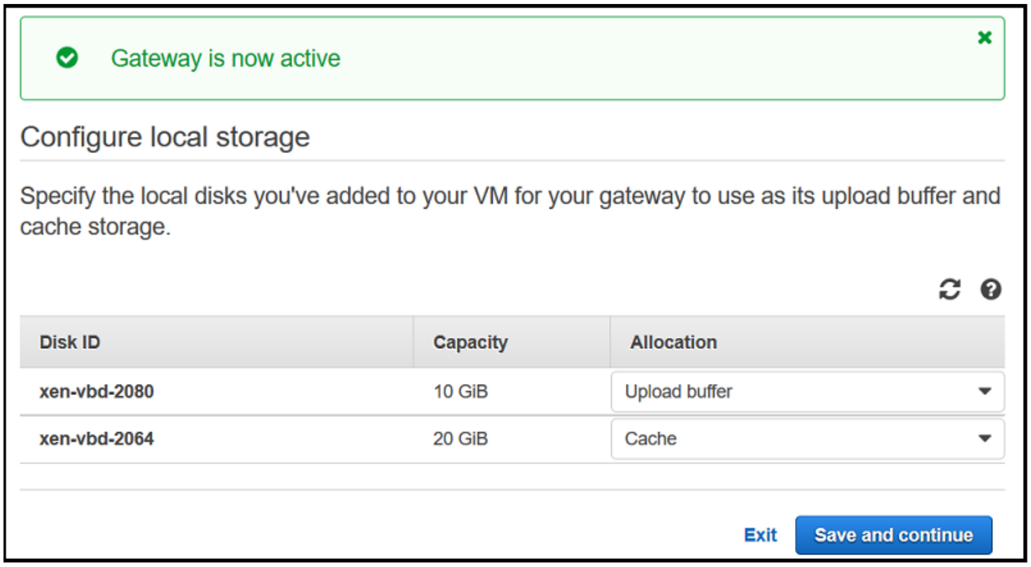

Gateway is now active

Configure local storage

Specify the local disks you've added to your VM for your gateway to use as its upload buffer and cache storage.

Disk ID	Capacity	Allocation
xen-vbd-2080	10 GiB	Upload buffer
xen-vbd-2064	20 GiB	Cache

2. From the list next to your upload buffer disk, choose **Upload Buffer**.
3. Choose **Cache** for the disk you want to configure as a cache storage.

   If you don't see your disks, choose **Refresh**.
4. Choose **Save and continue** to save your configuration settings.

Next Step

You now create virtual tapes and start using your tape gateway.

Step 6: Create Virtual Tapes Using the AWS Storage Gateway Console (p. 324)

Step 6: Create Virtual Tapes Using the AWS Storage Gateway Console

Now that you have activated your tape gateway and configured local storage, you will create virtual tapes.

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

To create virtual tapes

1. In the navigation pane, choose the **Gateways** tab.
2. Choose **Create Tape** to open the **Create Tape** dialog box.

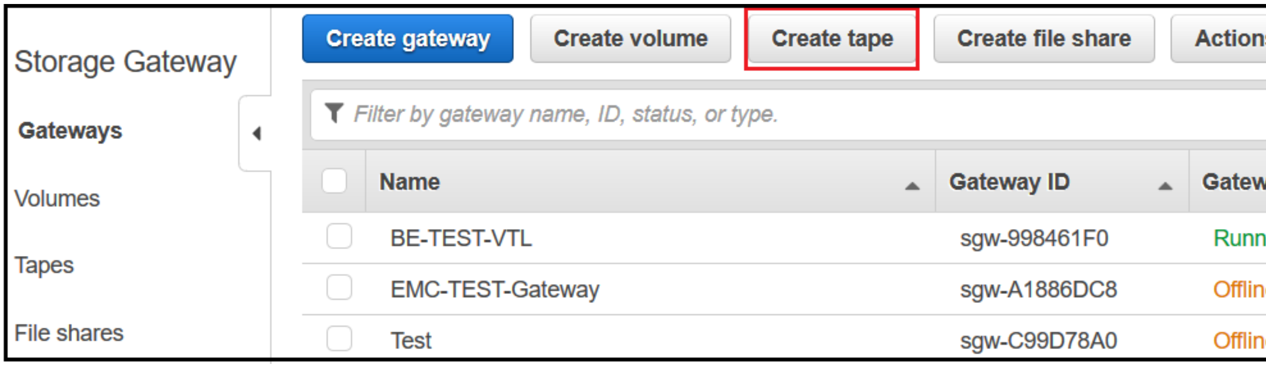

	Name	Gateway ID	Gatew
□	BE-TEST-VTL	sgw-998461F0	Runn
□	EMC-TEST-Gateway	sgw-A1886DC8	Offlin
□	Test	sgw-C99D78A0	Offlin

3. For **Gateway**, choose a gateway. The tape is created for this gateway.

   The following screenshot shows the **Create Virtual Tapes** page.

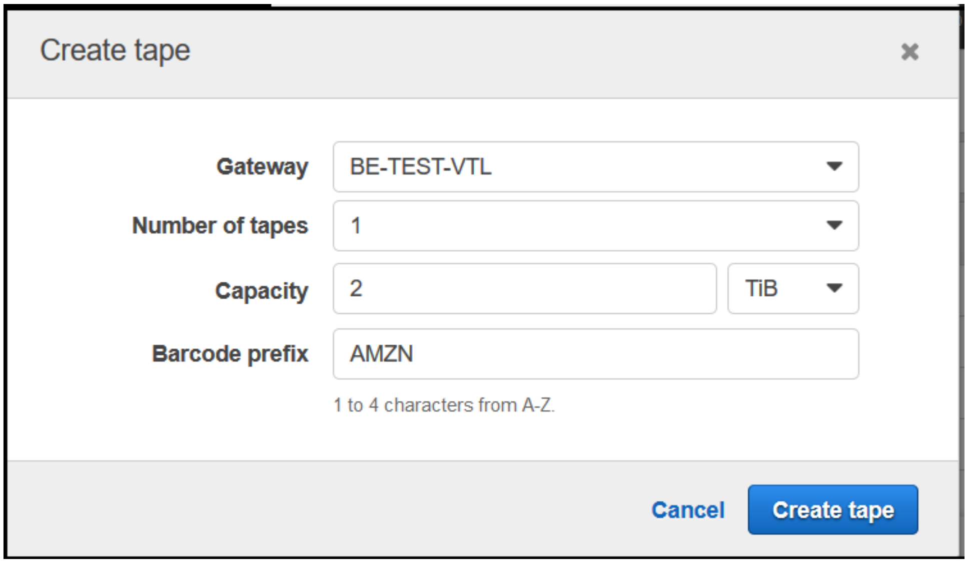

Gateway: BE-TEST-VTL

Number of tapes: 1

Capacity: 2 TiB

Barcode prefix: AMZN

1 to 4 characters from A-Z.

4. For **Number of Tapes**, choose the number of tapes you want to create. You can create a maximum of 10 virtual tapes at a time, up to a maximum of 1500 tapes per gateway. For more information, see AWS Storage Gateway Limits (p. 296).
5. For **Capacity**, type the size of the virtual tape you want to create. Tapes must be larger than 100 GiB. For information about capacity limits, see AWS Storage Gateway Limits (p. 296).
6. For **Barcode Prefix**, type the prefix you want to prepend to the barcode of your virtual tapes. For this exercise, you can accept the default.

   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 for your virtual tapes. The prefix must be uppercase letters (A–Z) and must be one to four characters long.
7. Choose **Create Tape**.
8. In the navigation pane, choose the **Tapes** tab to see your tapes.

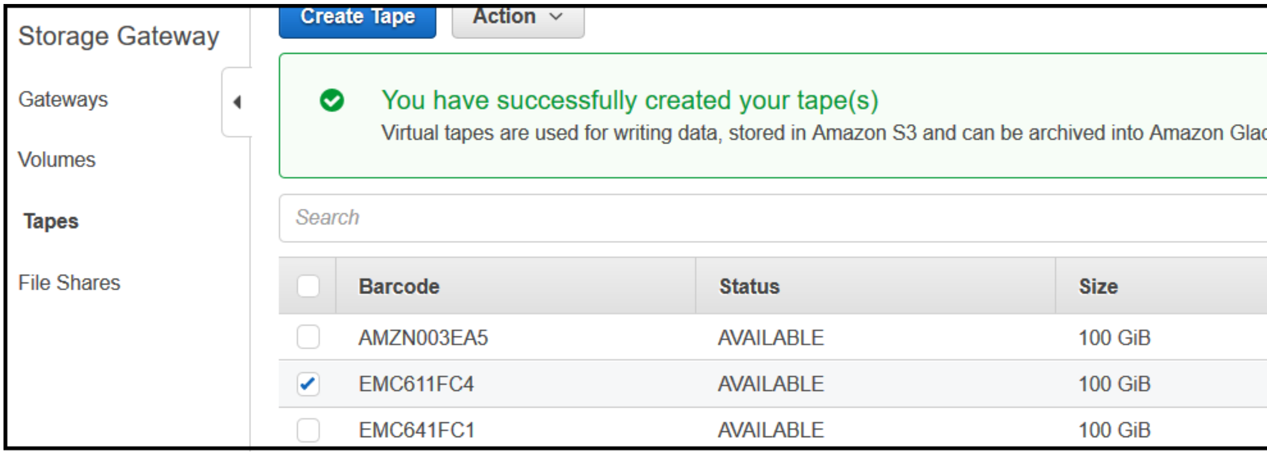

You have successfully created your tape(s)

Virtual tapes are used for writing data, stored in Amazon S3 and can be archived into Amazon Glac

☐	Barcode	Status	Size
☐	AMZN003EA5	AVAILABLE	100 GiB
☑	EMC611FC4	AVAILABLE	100 GiB
☐	EMC641FC1	AVAILABLE	100 GiB

The virtual tapes you have created appear in the AWS Storage Gateway console. The console shows the barcode, capacity, and status of the virtual tapes. 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 Working With Tapes (p. 275).

Next Step

Step 7: Connect Your Tape Gateway Devices to Your Windows Client (p. 326)

Step 7: Connect Your Tape Gateway Devices to Your Windows Client

In this Getting Started exercise, you use the Windows iSCSI initiator to connect to your gateway's VTL devices. At the end of this procedure, the devices become available as local devices on your Windows client. After this step, you configure your backup software to access the devices. For instructions on accessing the iSCSI VTL devices from Linux, see Recommended Red Hat Linux iSCSI Settings (p. 44).

To connect your Windows client to the VTL devices

1. On the **Start** menu of your Windows client computer, type **`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.

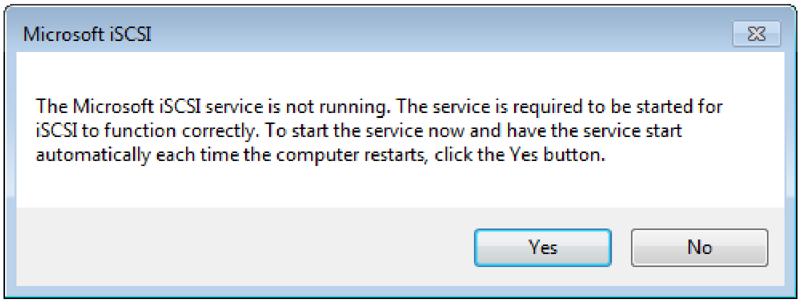

The Microsoft iSCSI service is not running. The service is required to be started for iSCSI to function correctly. To start the service now and have the service start automatically each time the computer restarts, click the Yes button.

3. In the **iSCSI Initiator Properties** dialog box, choose the **Discovery** tab, and then choose the **Discover Portal** button.

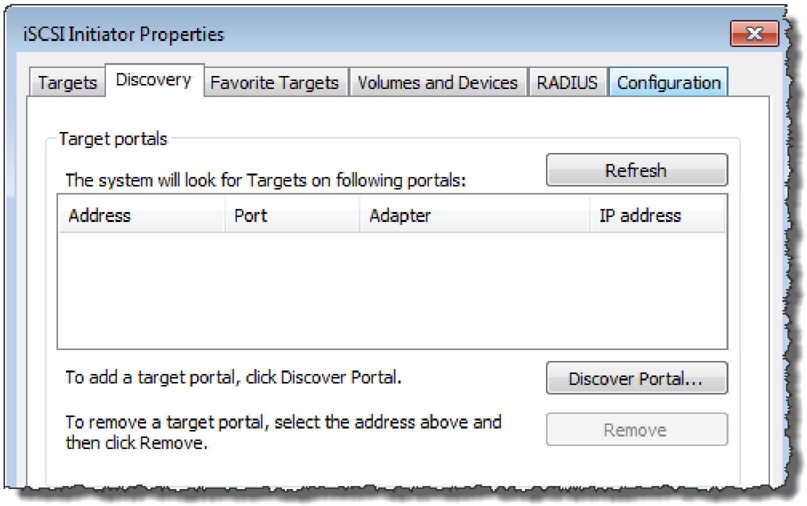

The system will look for Targets on following portals:

To add a target portal, click Discover Portal.

To remove a target portal, select the address above and then click Remove.

4. In the **Discover Target Portal** dialog box, type 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 AWS 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.

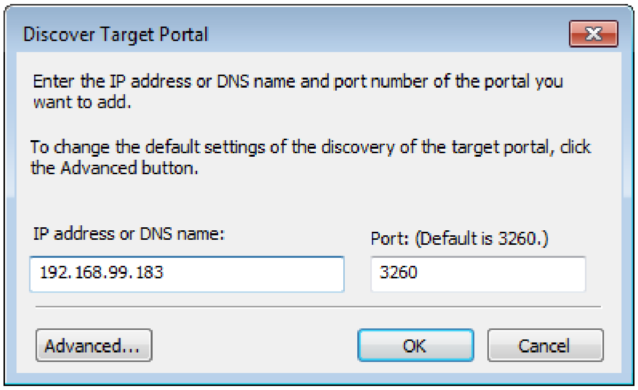

Enter the IP address or DNS name and port number of the portal you want to add.

To change the default settings of the discovery of the target portal, click the Advanced button.

IP address or DNS name: 192.168.99.183

Port: (Default is 3260.) 3260

5. Choose the **Targets** tab, and then choose **Refresh**. All ten tape drives and the medium changer appear in the **Discovered targets** box. The status for the targets is **Inactive**.

   The following screenshot shows the discovered targets.

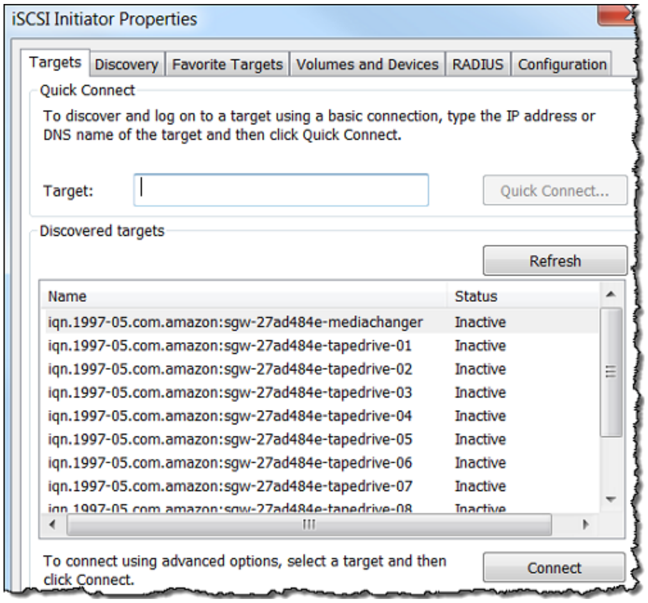

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Name	Status
iqn.1997-05.com.amazon:sgw-27ad484e-mediachanger	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-01	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-02	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-03	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-04	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-05	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-06	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-07	Inactive
iqn.1997-05.com.amazon:sgw-27ad484e-tapedrive-08	Inactive

To connect using advanced options, select a target and then click Connect.

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.
9. 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.

   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**.

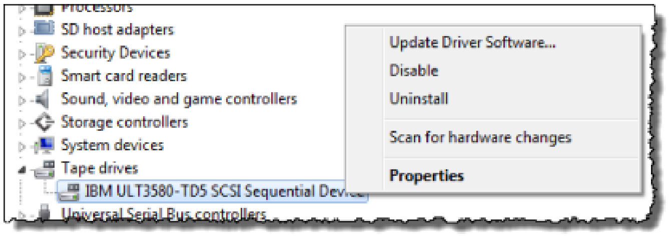
Processors
SD host adapters
Security Devices
Smart card readers
Sound, video and game controllers
Storage controllers
System devices
Tape drives
IBM ULT3580-TD5 SCSI Sequential Device
Universal Serial Bus controllers
Update Driver Software...
Disable
Uninstall
Scan for hardware changes
Properties

   3. In the **Driver** tab of the **Device Properties** dialog box, verify **Driver Provider** is Microsoft.

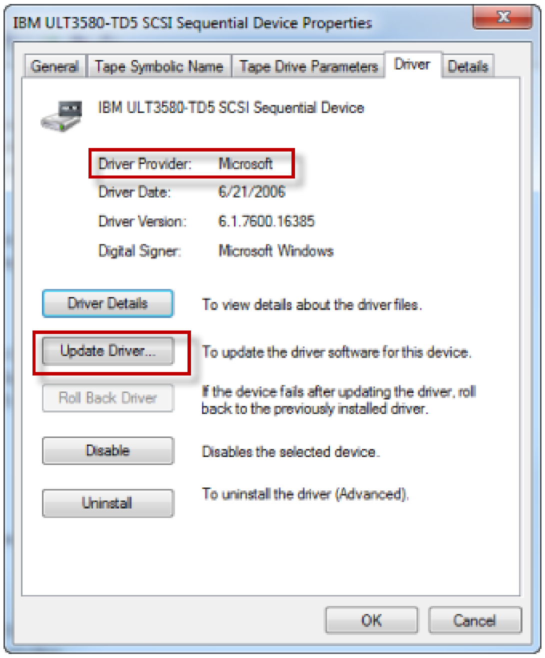

IBM ULT3580-TD5 SCSI Sequential Device

Driver Provider: Microsoft

Driver Date: 6/21/2006

Driver Version: 6.1.7600.16385

Digital Signer: Microsoft Windows

To view details about the driver files.

To update the driver software for this device.

If the device fails after updating the driver, roll back to the previously installed driver.

Disables the selected device.

To uninstall the driver (Advanced).

4. If **Driver Provider** is not Microsoft, set the value as follows:

   1. Choose **Update Driver**.
   2. In the **Update Driver Software** dialog box, choose **Browse my computer for driver software**.

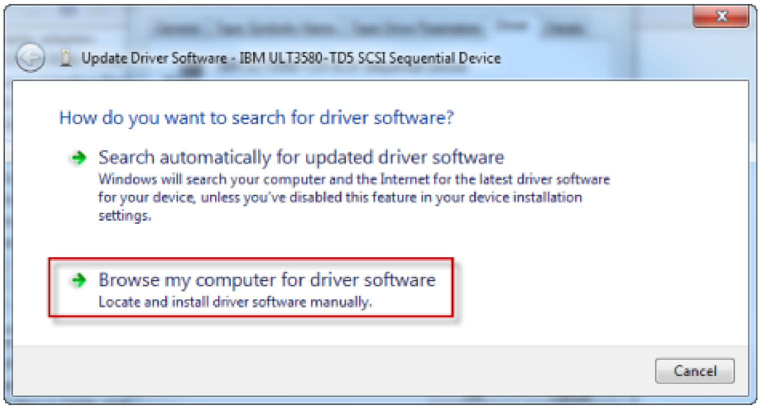
Update Driver Software - IBM ULT3580-TD5 SCSI Sequential Device
Cancel

How do you want to search for driver software?

Search automatically for updated driver software
Windows will search your computer and the Internet for the latest driver software for your device, unless you've disabled this feature in your device installation settings.

Browse my computer for driver software
Locate and install driver software manually.

   3. In the **Update Driver Software** dialog box, choose **Let me pick from a list of device drivers on my computer**.

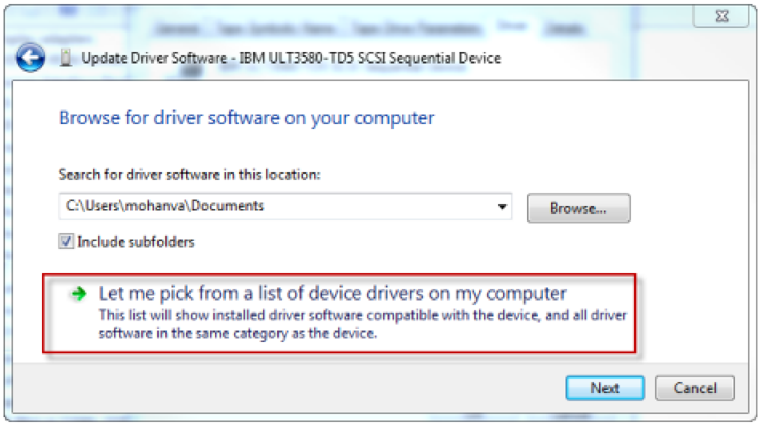

Update Driver Software - IBM ULT3580-TD5 SCSI Sequential Device

Browse for driver software on your computer

Search for driver software in this location:

C:\Users\mohanva\Documents

☑ Include subfolders

Let me pick from a list of device drivers on my computer

This list will show installed driver software compatible with the device, and all driver software in the same category as the device.

4. Select **LTO Tape drive** and choose **Next**.

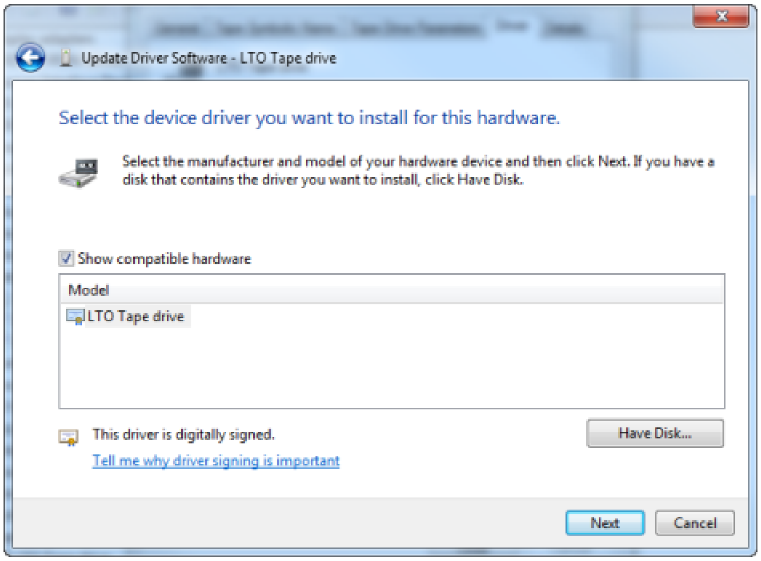

Update Driver Software - LTO Tape drive

Select the device driver you want to install for this hardware.

Select the manufacturer and model of your hardware device and then click Next. If you have a disk that contains the driver you want to install, click Have Disk.

☑ Show compatible hardware

Model
LTO Tape drive

This driver is digitally signed.

Tell me why driver signing is important

5. Choose **Close** to close the **Update Driver Software** window, and verify that the **Driver Provider** value is now set to Microsoft.
6. Repeat steps 9.2 through 9.5 to update all the tape drives.

Next Step

Step 8: Back Up Your Data to Virtual Tapes (p. 330)

Step 8: Back Up Your Data to Virtual Tapes

You test your tape gateway setup by performing the following tasks using your backup software:

1. Configure the backup software to detect your storage devices.
2. Back up data to a tape.
3. Archive the tape.
4. Retrieve the tape from the archive.
5. Restore data from the tape.

You can test your setup with one of the following types of compatible backup software:

- Testing Your Setup by Using Symantec NetBackup Version 7.x (p. 331).
- Testing Your Setup by Using Symantec Backup Exec (p. 344).
- Testing Your Setup by Using Microsoft System Center 2012 R2 Data Protection Manager (p. 347).
- Testing Your Setup by Using Veeam Backup & Replication (p. 350).
- Testing Your Setup by Using Dell NetVault Backup (p. 353).
- Testing Your Setup by Using EMC NetWorker (p. 356).

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Testing Your Setup by Using Symantec NetBackup Version 7.x

You can back up your data to virtual tapes, archive the tapes in a virtual tape shelf (VTS), and manage your virtual tape library (VTL) devices by using Symantec NetBackup version 7.x. In this topic, you can find basic documentation on how to configure the NetBackup software for a tape gateway and perform a backup. For detailed information about how to use NetBackup, see the NetBackup documentation on the Symantec website. For Symantec support information on hardware compatibility, see the Symantec NetBackup Enterprise Server and Server 7.x Hardware Compatibility List on the Symantec website.

In this topic, you can find out how to configure storage, write data to a tape, archive a tape in the VTS, and restore the data.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

Configuring NetBackup Storage Devices

After you have connected the virtual tape library (VTL) devices to the Windows client, you configure Symantec NetBackup version 7.x storage to recognize your devices. For information about how to connect VTL devices to the Windows client, see Step 7: Connect Your Tape Gateway Devices to Your Windows Client (p. 326).

To configure NetBackup to use storage devices on your tape gateway

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

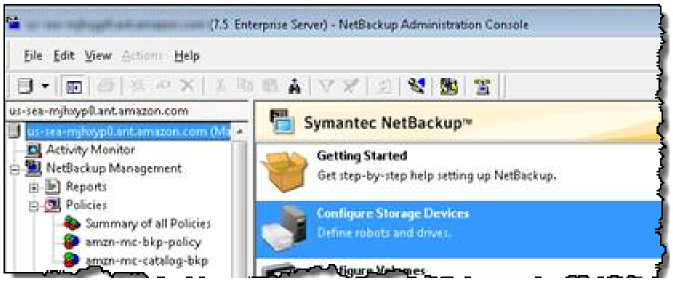

[redacted] (7.5 Enterprise Server) - NetBackup Administration Console
File
Edit
View
Actions
Help
us-sea-mjhxyp0.ant.amazon.com
us-sea-mjhxyp0.ant.amazon.com (Ma
Activity Monitor
NetBackup Management
Reports
Policies
Summary of all Policies
amzn-mc-bkp-policy
amzn-mc-catalog-bkp
Symantec NetBackup™
Getting Started
Get step-by-step help setting up NetBackup.
Configure Storage Devices
Define robots and drives.
figure Volumes

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

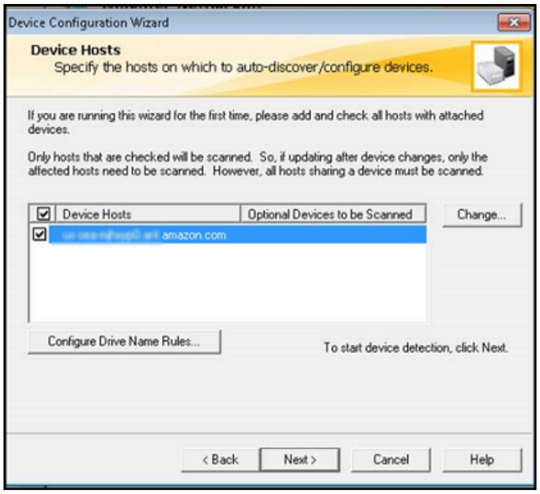

Device Hosts

Specify the hosts on which to auto-discover/configure devices.

If you are running this wizard for the first time, please add and check all hosts with attached devices.

Only hosts that are checked will be scanned. So, if updating after device changes, only the affected hosts need to be scanned. However, all hosts sharing a device must be scanned.

☑	Device Hosts	Optional Devices to be Scanned
☑	[redacted] amazon.com	

To start device detection, click Next.

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

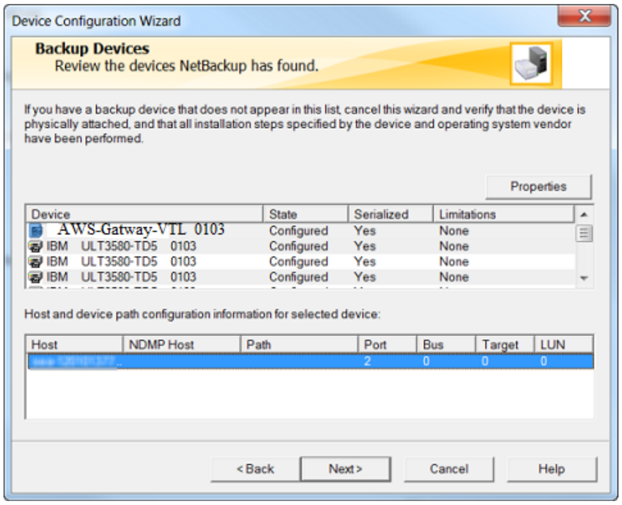

Backup Devices

Review the devices NetBackup has found.

If you have a backup device that does not appear in this list, cancel this wizard and verify that the device is physically attached, and that all installation steps specified by the device and operating system vendor have been performed.

Device	State	Serialized	Limitations
AWS-Gatway-VTL 0103	Configured	Yes	None
IBM ULT3580-TD5 0103	Configured	Yes	None
IBM ULT3580-TD5 0103	Configured	Yes	None
IBM ULT3580-TD5 0103	Configured	Yes	None

Host and device path configuration information for selected device:

Host	NDMP Host	Path	Port	Bus	Target	LUN
[redacted]			2	0	0	0

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

To verify your devices in the NetBackup software

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.

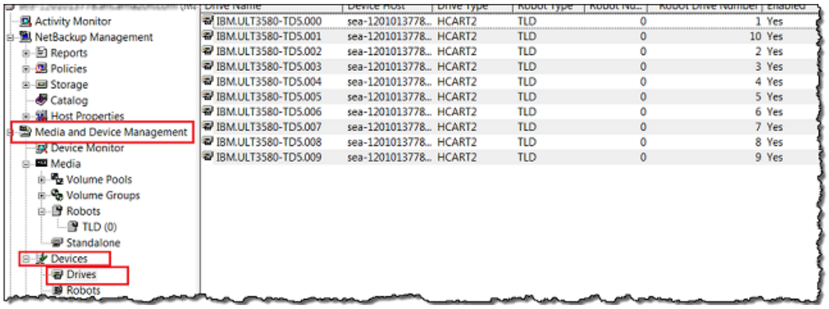

Drive Name	Device Host	Drive Type	Robot Type	Robot Nu...	Robot Drive Number	Enabled
IBM.ULT3580-TD5.000	sea-1201013778...	HCART2	TLD	0	1	Yes
IBM.ULT3580-TD5.001	sea-1201013778...	HCART2	TLD	0	10	Yes
IBM.ULT3580-TD5.002	sea-1201013778...	HCART2	TLD	0	2	Yes
IBM.ULT3580-TD5.003	sea-1201013778...	HCART2	TLD	0	3	Yes
IBM.ULT3580-TD5.004	sea-1201013778...	HCART2	TLD	0	4	Yes
IBM.ULT3580-TD5.005	sea-1201013778...	HCART2	TLD	0	5	Yes
IBM.ULT3580-TD5.006	sea-1201013778...	HCART2	TLD	0	6	Yes
IBM.ULT3580-TD5.007	sea-1201013778...	HCART2	TLD	0	7	Yes
IBM.ULT3580-TD5.008	sea-1201013778...	HCART2	TLD	0	8	Yes
IBM.ULT3580-TD5.009	sea-1201013778...	HCART2	TLD	0	9	Yes

2. In the **Devices** node, choose **Robots** to display all your medium changers. In the NetBackup software, 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**.

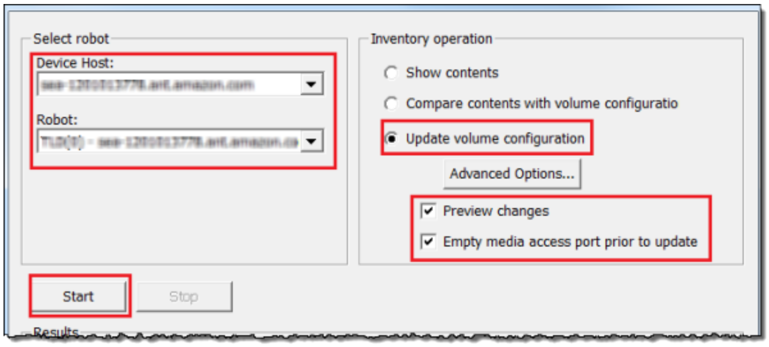

Device Host: [redacted]

Robot: [redacted]

○ Show contents
○ Compare contents with volume configuratio
● Update volume configuration

☑ Preview changes
☑ Empty media access port prior to update

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.

7. 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.

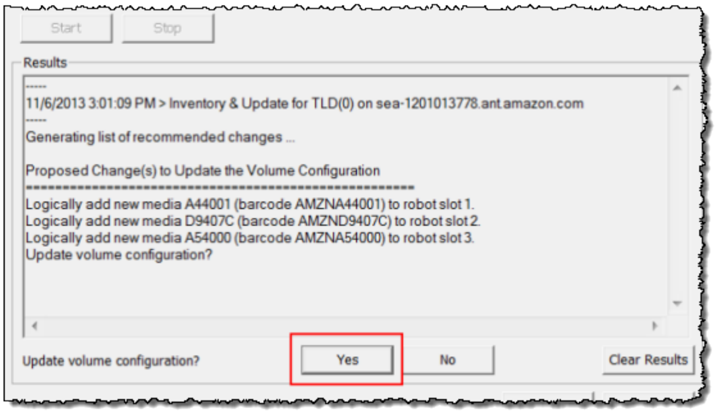

```
-----
11/6/2013 3:01:09 PM > Inventory & Update for TLD(0) on sea-1201013778.ant.amazon.com
-----
Generating list of recommended changes ...

Proposed Change(s) to Update the Volume Configuration
=========================================================
Logically add new media A44001 (barcode AMZNA44001) to robot slot 1.
Logically add new media D9407C (barcode AMZND9407C) to robot slot 2.
Logically add new media A54000 (barcode AMZNA54000) to robot slot 3.
Update volume configuration?
```

Update volume configuration?

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

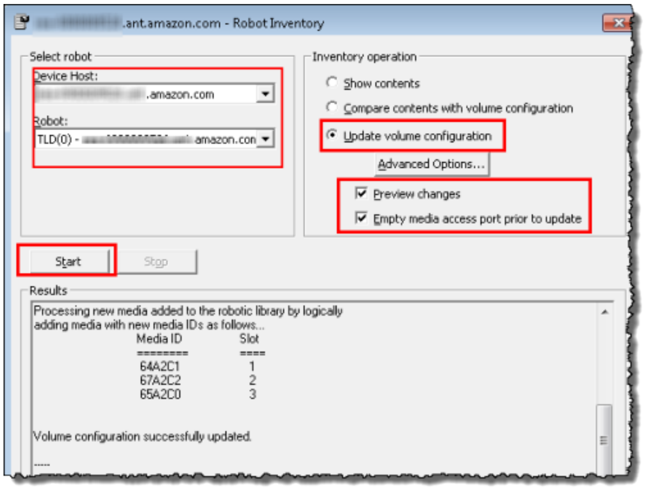

[redacted].ant.amazon.com - Robot Inventory

Select robot

Device Host: [redacted].amazon.com

Robot: TLD(0) - [redacted].amazon.con

Inventory operation

○ Show contents

○ Compare contents with volume configuration

● Update volume configuration

☑ Preview changes

☑ Empty media access port prior to update

```
Processing new media added to the robotic library by logically
adding media with new media IDs as follows...
                Media ID          Slot
                ========          ====
                 64A2C1             1
                 67A2C2             2
                 65A2C0             3

Volume configuration successfully updated.
-----
```

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 software, you might have multiple robots. Make sure you select the right robot.

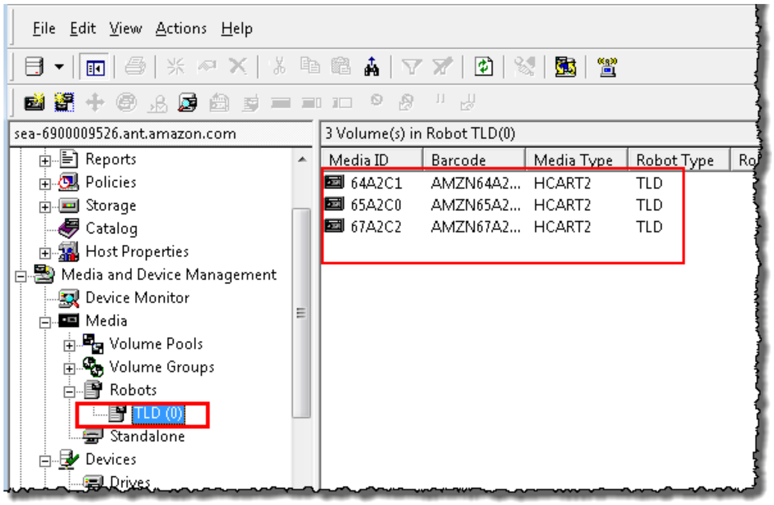

3 Volume(s) in Robot TLD(0)

Media ID	Barcode	Media Type	Robot Type	Ro
64A2C1	AMZN64A2...	HCART2	TLD	
65A2C0	AMZN65A2...	HCART2	TLD	
67A2C2	AMZN67A2...	HCART2	TLD	

Now that you have connected your devices and made them available to your backup software, 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 in the virtual tape shelf (VTS), which is backed by Amazon Glacier.

Backing Up Sample Data to a Tape on Your Tape Gateway

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 AWS Storage Gateway detail page.

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.

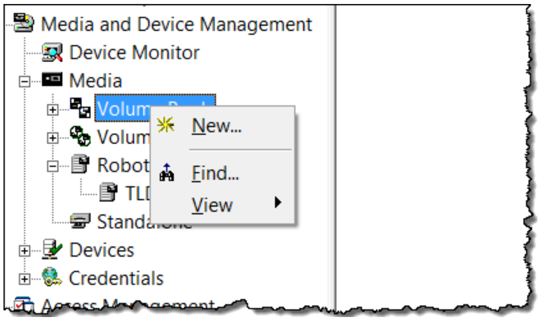

3. For **Name**, type a name for your volume pool.
4. For **Description**, type a description for the volume pool, and then choose **OK**. The volume pool you just created is added to the volume pool list.

   The following screenshot shows a list of volume pools.

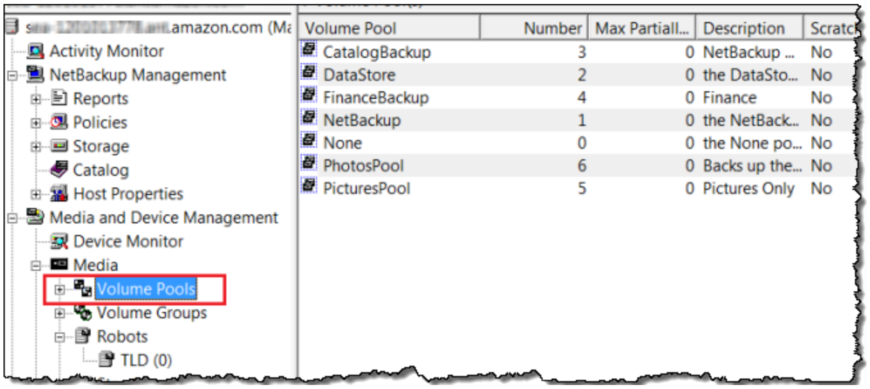

Volume Pool	Number	Max Partiall...	Description	Scratc
CatalogBackup	3	0	NetBackup ...	No
DataStore	2	0	the DataSto...	No
FinanceBackup	4	0	Finance	No
NetBackup	1	0	the NetBack...	No
None	0	0	the None po...	No
PhotosPool	6	0	Backs up the...	No
PicturesPool	5	0	Pictures Only	No

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.

   Note that 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.

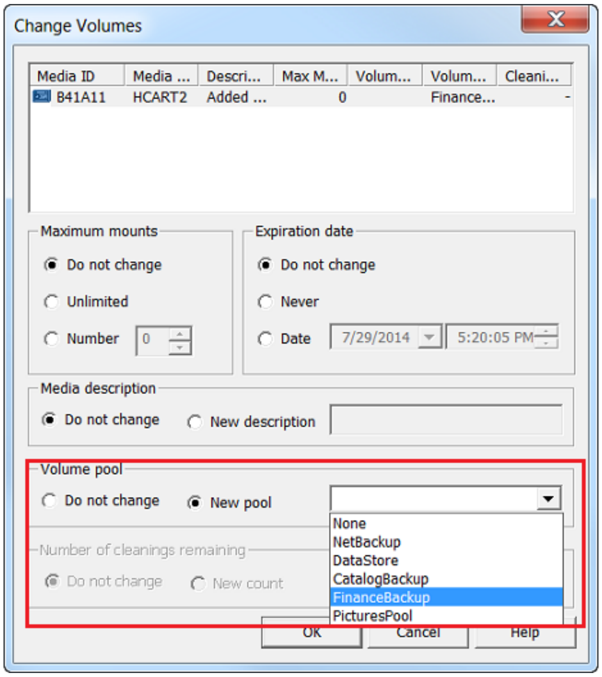

Media ID	Media ...	Descri...	Max M...	Volum...	Volum...	Cleani...
B41A11	HCART2	Added ...	0		Finance...	-

Maximum mounts: ● Do not change ○ Unlimited ○ Number 0

Expiration date: ● Do not change ○ Never ○ Date 7/29/2014 5:20:05 PM

Media description: ● Do not change ○ New description

Volume pool: ○ Do not change ● New pool

None
NetBackup
DataStore
CatalogBackup
FinanceBackup
PicturesPool

Number of cleanings remaining: ● Do not change ○ New count

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 Symantec NetBackup console.

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

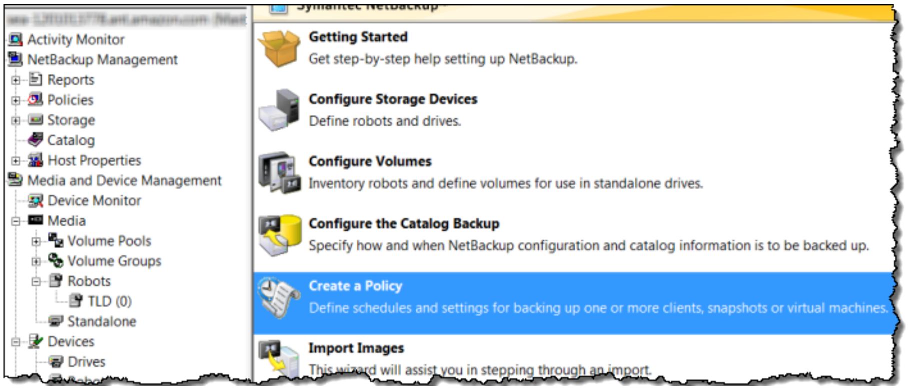

Getting Started
Get step-by-step help setting up NetBackup.

Configure Storage Devices
Define robots and drives.

Configure Volumes
Inventory robots and define volumes for use in standalone drives.

Configure the Catalog Backup
Specify how and when NetBackup configuration and catalog information is to be backed up.

Create a Policy
Define schedules and settings for backing up one or more clients, snapshots or virtual machines.

Import Images
This wizard will assist you in stepping through an import.

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).

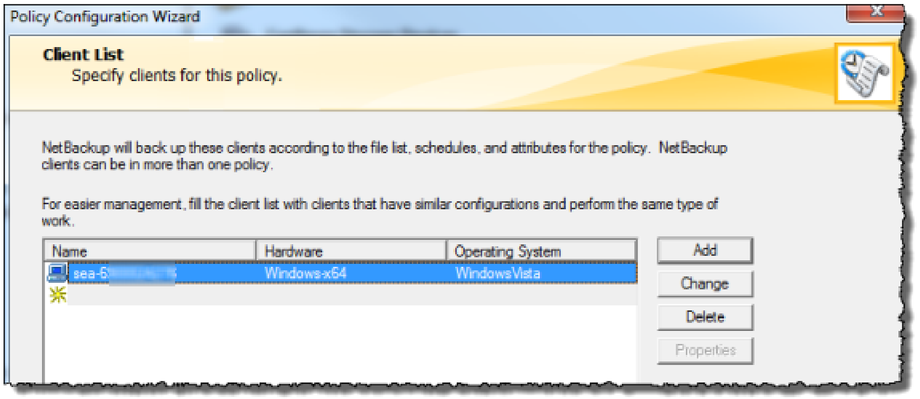

Policy Configuration Wizard

Client List
Specify clients for this policy.

NetBackup will back up these clients according to the file list, schedules, and attributes for the policy. NetBackup clients can be in more than one policy.

For easier management, fill the client list with clients that have similar configurations and perform the same type of work.

Name	Hardware	Operating System
sea-6[redacted]	Windows-x64	WindowsVista

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

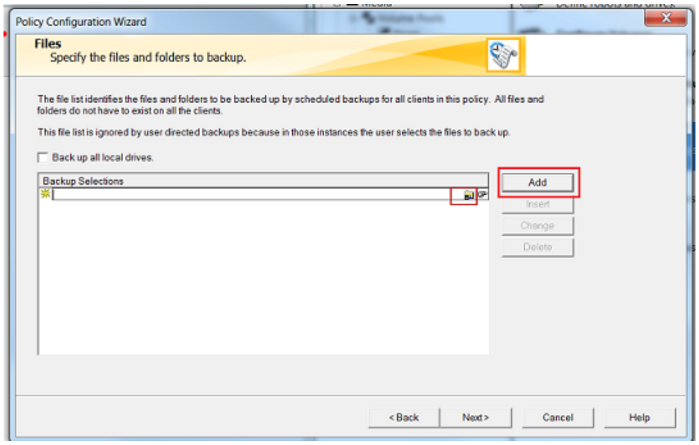

Policy Configuration Wizard

Files
Specify the files and folders to backup.

The file list identifies the files and folders to be backed up by scheduled backups for all clients in this policy. All files and folders do not have to exist on all the clients.

This file list is ignored by user directed backups because in those instances the user selects the files to back up.

☐ Back up all local drives.

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**.

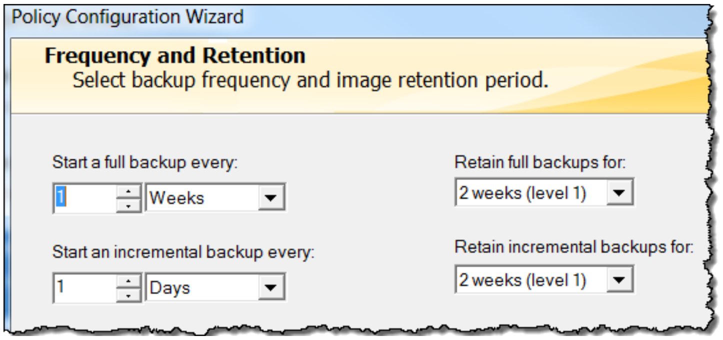

Policy Configuration Wizard

Frequency and Retention
Select backup frequency and image retention period.

Start a full backup every: 1 Weeks

Retain full backups for: 2 weeks (level 1)

Start an incremental backup every: 1 Days

Retain incremental backups for: 2 weeks (level 1)

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.

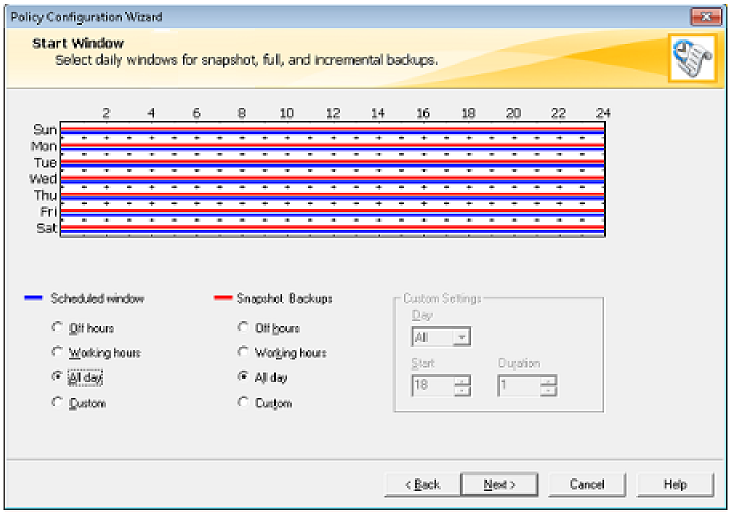

11. In the **Policy Configuration** wizard, choose **Finish**.

The policy runs the backups according to the schedule. You can also perform a manual backup at any time, which we will 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**.

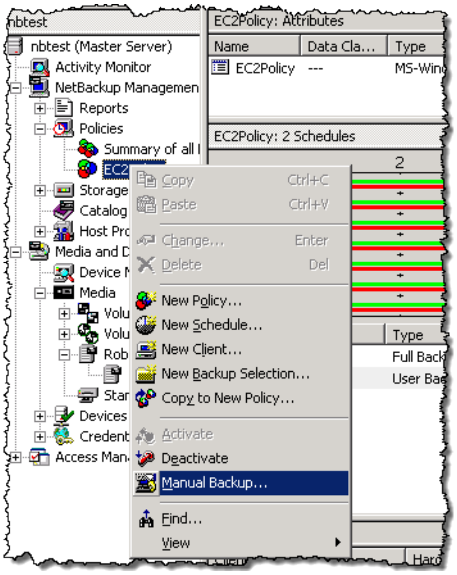

4. In the **Manual Backup** window, select a schedule, select a client, and then choose **OK**.

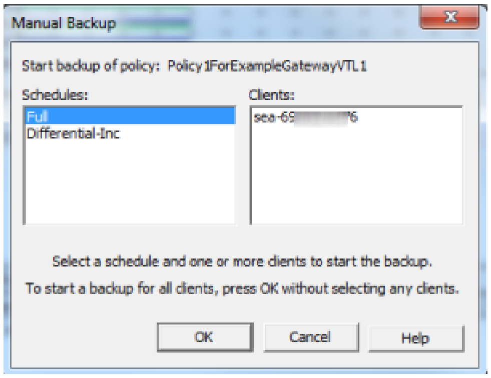

Start backup of policy: Policy1ForExampleGatewayVTL1

Schedules:
- Full
- Differential-Inc

Clients:
- sea-69[redacted]6

Select a schedule and one or more clients to start the backup.

To start a backup for all clients, press OK without selecting any clients.

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.

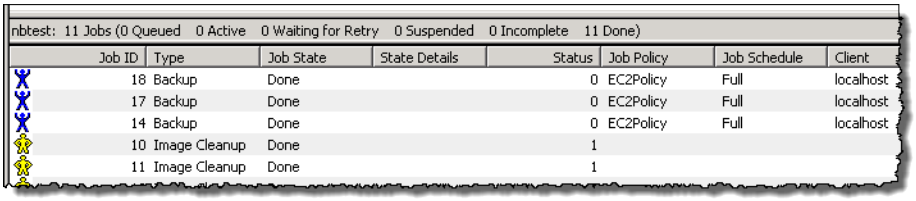

nbtest: 11 Jobs (0 Queued 0 Active 0 Waiting for Retry 0 Suspended 0 Incomplete 11 Done)

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 will need this barcode in the procedure in the next section, where you archive the tape to the virtual tape shelf (VTS).

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.

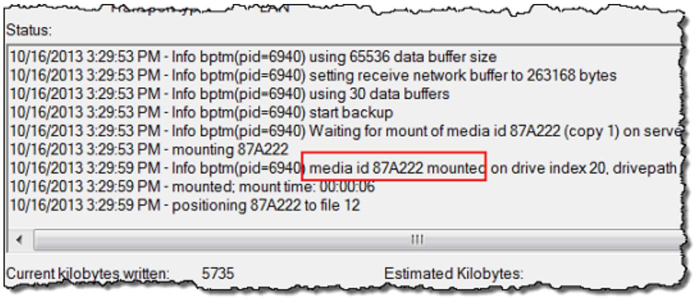

Status:

```
10/16/2013 3:29:53 PM - Info bptm(pid=6940) using 65536 data buffer size
10/16/2013 3:29:53 PM - Info bptm(pid=6940) setting receive network buffer to 263168 bytes
10/16/2013 3:29:53 PM - Info bptm(pid=6940) using 30 data buffers
10/16/2013 3:29:53 PM - Info bptm(pid=6940) start backup
10/16/2013 3:29:53 PM - Info bptm(pid=6940) Waiting for mount of media id 87A222 (copy 1) on serve
10/16/2013 3:29:53 PM - mounting 87A222
10/16/2013 3:29:59 PM - Info bptm(pid=6940) media id 87A222 mounted on drive index 20, drivepath
10/16/2013 3:29:59 PM - mounted; mount time: 00:00:06
10/16/2013 3:29:59 PM - positioning 87A222 to file 12
```

Current kilobytes written: 5735 Estimated Kilobytes:

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 VTS.

Archiving the Tape

When you archive a tape, AWS Storage Gateway moves the tape from your gateway's virtual tape library (VTL) to the virtual tape shelf (VTS), 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**.

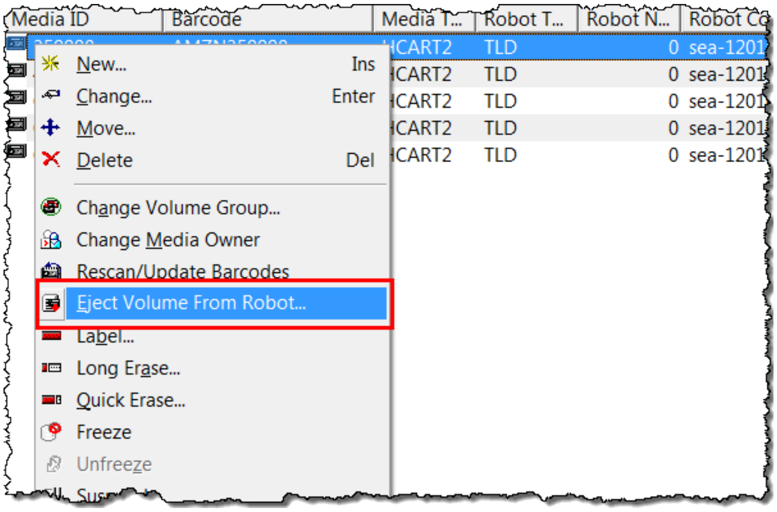

4. In the **Eject Volumes** window, make sure the **Media ID** matches the virtual tape you want to eject, and then choose **Eject**.

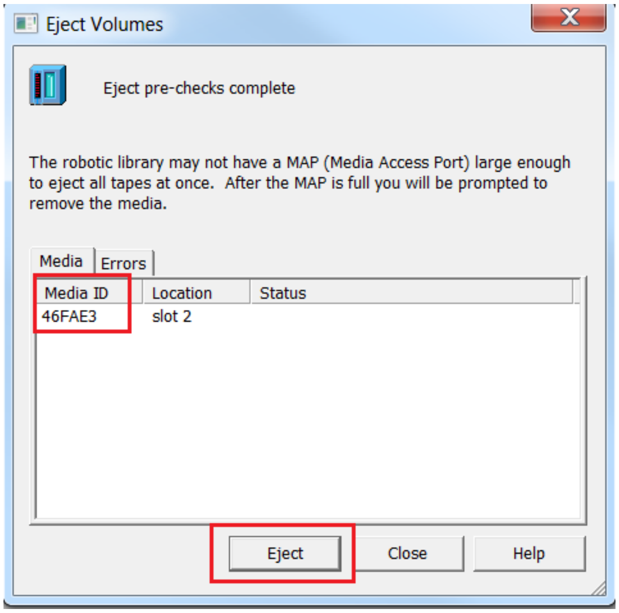

Eject pre-checks complete

The robotic library may not have a MAP (Media Access Port) large enough to eject all tapes at once. After the MAP is full you will be prompted to remove the media.

Media ID	Location	Status
46FAE3	slot 2	

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

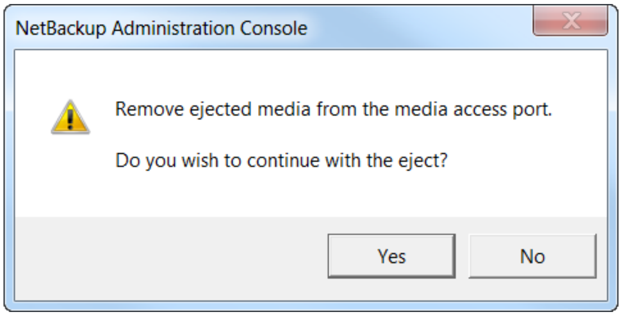

NetBackup Administration Console

Remove ejected media from the media access port.

Do you wish to continue with the eject?

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

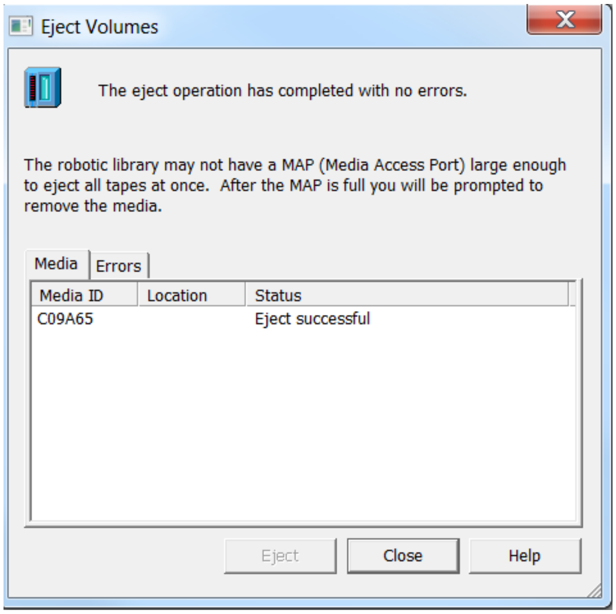

Eject Volumes

The eject operation has completed with no errors.

The robotic library may not have a MAP (Media Access Port) large enough to eject all tapes at once. After the MAP is full you will be prompted to remove the media.

Media ID	Location	Status
C09A65		Eject successful

6. Choose **Close** to close the **Eject Volumes** window.
7. In the AWS 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 will be listed in the gateway's VTL with the status IN TRANSIT TO VTS. When archiving starts, the status will be ARCHIVING. Once data upload has completed, the ejected tape will no longer be listed in the VTL.
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 AWS Storage Gateway console, choose the VTS. Verify that your archived tape is listed in the VTS and that its status is ARCHIVED.

Retrieving the Archived Tape from the VTS Back to Your Tape Gateway

Archived tapes are stored in a virtual tape shelf (VTS), which provides offline storage. If you want to access tape data, you must first retrieve the tape from the VTS back to your gateway. In this step, you will retrieve the tape that you archived in the preceding step.

Note

It takes about 24 hours to retrieve a tape from the VTS to a gateway.

To retrieve an archived tape

1. In the navigation pane of the AWS Storage Gateway console, choose **Virtual Tape Shelf**. The console displays all virtual tapes that have been archived by all your gateways.
2. Select the virtual tape you want to retrieve, and choose **Retrieve Tape**.

   Note
   The status of the virtual tape you want to retrieve must be ARCHIVED.
3. In the **Tape Barcode** field of the **Retrieve Tape** wizard, verify that the barcode identifies the virtual tape you want to retrieve.
4. For **Gateway**, choose the gateway you want to retrieve the archived tape to, and then choose **Proceed**.

The status of the tape changes from ARCHIVED to RETRIEVING. After all the data is moved, the status of the virtual tape in the VTS changes to RETRIEVED, and the tape appears in your gateway's VTL.

Note
Retrieved virtual tapes are read-only.

Restoring Data from the Tape

In this step, you restore data from the virtual tape to your client computer.

Note
This step does not use the NetBackup Administration Console that you used in the previous steps. Instead, you will use the Backup, Archive, and Restore software installed with Symantec NetBackup software.

To restore data

1. Start the Backup, Archive, and Restore software, and run it as an administrator.
2. Choose the **Select for Restore** tab.

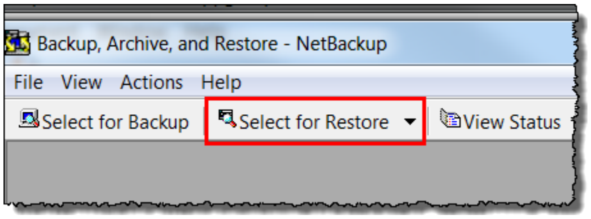

3. If you have previously backed up data, you will see backup icons in the **NetBackup History** pane. In this example, there is only one backup icon.

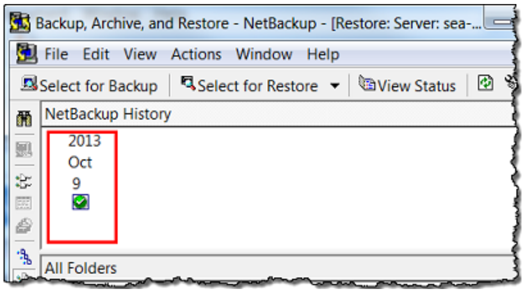
Backup, Archive, and Restore - NetBackup - [Restore: Server: sea-...
File Edit View Actions Window Help
Select for Backup
Select for Restore
View Status
NetBackup History
2013
Oct
9
All Folders

4. Select the backup icon that represents the backup you want to restore.
5. In the **All Folders** pane, select the folder you want to restore, and then choose the **Restore Marked Files** icon in the left pane. This icon is not labeled. The name appears when you rest your mouse on the icon.

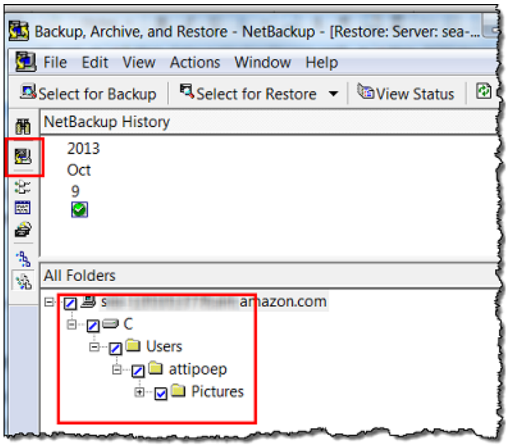

6. In the **Restore Marked Files** window, select the **Restore everything to a different location (maintaining existing structure)** button. This selection avoids overwriting your original data.

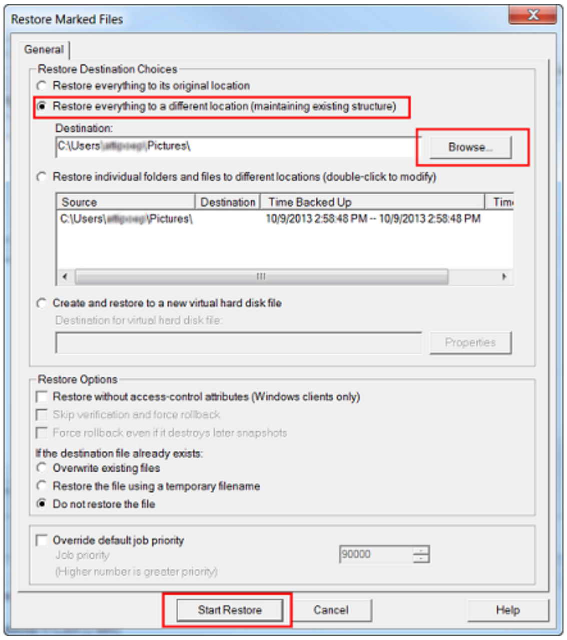

Restore Destination Choices

○ Restore everything to its original location

● Restore everything to a different location (maintaining existing structure)

Destination:
C:\Users\[redacted]\Pictures\

○ Restore individual folders and files to different locations (double-click to modify)

Source	Destination	Time Backed Up	Tim
C:\Users\[redacted]\Pictures\		10/9/2013 2:58:48 PM -- 10/9/2013 2:58:48 PM	

○ Create and restore to a new virtual hard disk file

Destination for virtual hard disk file:

Restore Options

□ Restore without access-control attributes (Windows clients only)

□ Skip verification and force rollback

□ Force rollback even if it destroys later snapshots

If the destination file already exists:

○ Overwrite existing files

○ Restore the file using a temporary filename

● Do not restore the file

□ Override default job priority

Job priority 90000

(Higher number is greater priority)

7. For **Destination**, browse to the folder you want to restore the data to, and then choose **Start Restore**.
8. In the dialog box that appears, choose **Yes** to view the progress of the restore process.

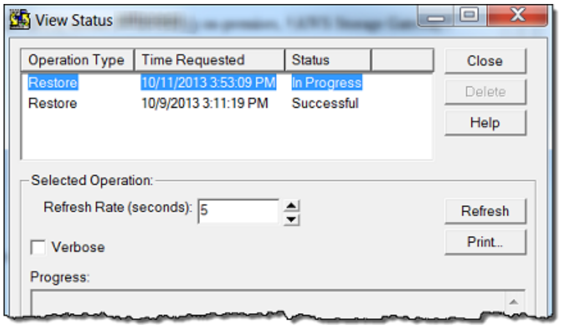

Operation Type	Time Requested	Status
Restore	10/11/2013 3:53:09 PM	In Progress
Restore	10/9/2013 3:11:19 PM	Successful

Selected Operation:

Refresh Rate (seconds): 5

□ Verbose

Progress:

In the **View Status** window, you can see the status of the restore process. If the restore succeeds, the status changes to **Successful**.

Next Step

Clean Up Resources You Don't Need (p. 359)

Testing Your Setup by Using Symantec Backup Exec

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

- Symantec Backup Exec 2014
- Symantec Backup Exec 15

The procedure for using these versions of Backup Exec with AWS Storage Gateway–VTL is the same. For detailed information about how to use Backup Exec, see the How to Create Secure Backups with Backup Exec video on the Symantec website. For Symantec support information on hardware compatibility, see the Software Compatibility Lists (SCL), Hardware Compatibility Lists (HCL), and Administrator Guides for Backup Exec (all versions) on the Symantec website. For information about best practices, see Best Practices for using Symantec Backup products (NetBackup, Backup Exec) with the Amazon Web Services (AWS) Tape Gateway on the Symantec website.

Using Symantec Backup Exec, you can configure storage, write data to a tape, archive a tape in the VTS, and restore the data.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

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 Step 7: Connect Your Tape Gateway Devices to Your Windows Client (p. 326).

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.

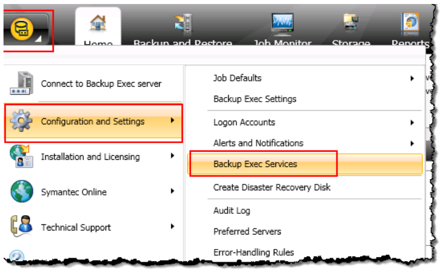

3. Choose **Restart All Services**. Backup Exec then recognizes the VTL devices (that is, the medium changer and tape drives). The restart process might take a few minutes.

   Note
   AWS Storage Gateway–VTL provides 10 tape drives. However, your Backup Exec license agreement might require your backup software to work with fewer than 10 tape drives. In that case, you must disable tape drives in the Backup Exec robotic library to leave only the number of tape drives allowed by your license agreement enabled. For instructions, see Disabling a Tape Drive in Backup Exec (p. 347).

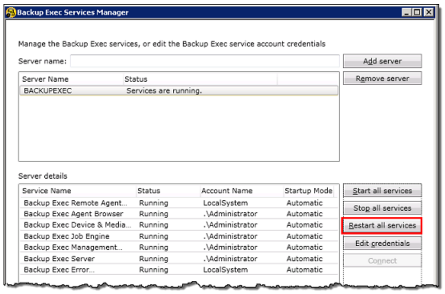

Manage the Backup Exec services, or edit the Backup Exec service account credentials

Server name:

Server Name	Status
BACKUPEXEC	Services are running.

Server details

Service Name	Status	Account Name	Startup Mode
Backup Exec Remote Agent...	Running	LocalSystem	Automatic
Backup Exec Agent Browser	Running	.\Administrator	Automatic
Backup Exec Device & Media...	Running	.\Administrator	Automatic
Backup Exec Job Engine	Running	.\Administrator	Automatic
Backup Exec Management...	Running	.\Administrator	Automatic
Backup Exec Server	Running	.\Administrator	Automatic
Backup Exec Error...	Running	LocalSystem	Automatic

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
   Symantec Backup Exec software requires the AWS 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 software. For information about how to select a different medium changer type, see Selecting a Medium Changer After Gateway Activation (p. 273).

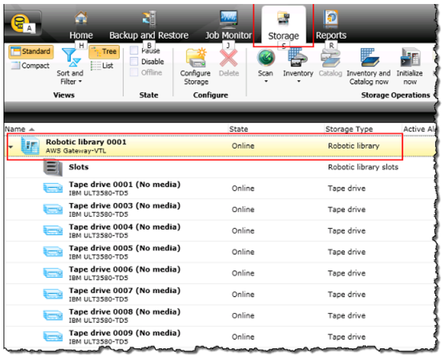

Name	State	Storage Type	Active Al
Robotic library 0001 AWS Gateway-VTL	Online	Robotic library	
Slots		Robotic library slots	
Tape drive 0001 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0003 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0004 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0005 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0006 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0007 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0008 (No media) IBM ULT3580-TD5	Online	Tape drive	
Tape drive 0009 (No media) IBM ULT3580-TD5	Online	Tape drive	

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 will be moved into the tape drives.

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

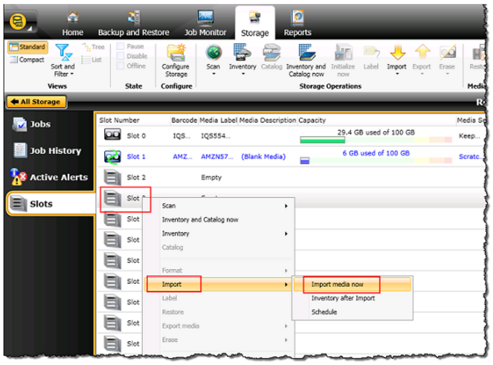

Slot Number	Barcode	Media Label	Media Description	Capacity	Media S
Slot 0	IQS...	IQS554...		29.4 GB used of 100 GB	Keep...
Slot 1	AMZ...	AMZN57...	(Blank Media)	6 GB used of 100 GB	Scratc...
Slot 2		Empty			
Slot					
Slot					
Slot					
Slot					
Slot					
Slot					
Slot					
Slot					

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

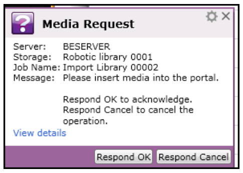

Media Request

Server: BESERVER
Storage: Robotic library 0001
Job Name: Import Library 00002
Message: Please insert media into the portal.

Respond OK to acknowledge.
Respond Cancel to cancel the operation.

View details

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

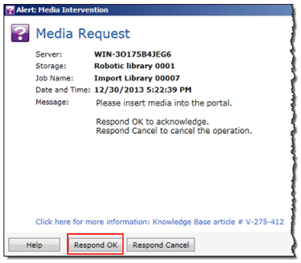

Alert: Media Intervention

Media Request

Server: WIN-3O175B4JEG6
Storage: Robotic library 0001
Job Name: Import Library 00007
Date and Time: 12/30/2013 5:22:39 PM
Message: Please insert media into the portal.

Respond OK to acknowledge.
Respond Cancel to cancel the operation.

Click here for more information: Knowledge Base article # V-275-412

The tape appears in the slot you selected.

> **Note**
> Tapes that are imported include empty tapes and tapes that have been retrieved from the VTS 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.

Archiving a Tape Using Backup Exec

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

To archive your tape in VTS

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.

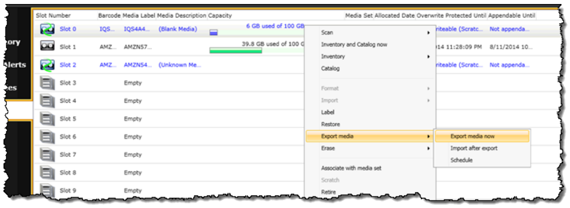

Slot Number	Barcode	Media Label	Media Description	Capacity	Media Set	Allocated Date	Overwrite Protected Until	Appendable Until
Slot 0	IQS...	IQS4A4...	(Blank Media)	6 GB used of 100 GB			riteable (Scratc...	Not appenda...
Slot 1	AMZ...	AMZN57...		39.8 GB used of 100 G			14 11:28:09 PM	8/11/2014 10...
Slot 2	AMZ...	AMZN54...	(Unknown Me...				riteable (Scratc...	Not appenda...
Slot 3		Empty						
Slot 4		Empty						
Slot 5		Empty						
Slot 6		Empty						
Slot 7		Empty						
Slot 8		Empty						
Slot 9		Empty						

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

   In the AWS 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 will be listed in the gateway's 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 will no longer be listed in the VTL.
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 AWS Storage Gateway console, choose **Virtual Tape Shelf (VTS)**. Verify that your archived tape is listed in the VTS and that the 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 from your VTS to a tape gateway. For instructions, see Retrieving the Archived Tape from Archive Back to Your Tape Gateway (p. 92).
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.

Disabling a Tape Drive in Backup Exec

AWS Storage Gateway–VTL provides 10 tape drives, but you might decide to use fewer tape drives. In that case, you disable the tape drives you will not 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 disable, and then choose **Disable**.

Next Step

Clean Up Resources You Don't Need (p. 359)

Testing Your Setup by Using Microsoft System Center 2012 R2 Data Protection Manager

You can back up your data to virtual tapes, archive the tapes in a virtual tape shelf (VTS), and manage your virtual tape library (VTL) devices by using Microsoft System Center 2012 R2 Data Protection Manager (DPM). In this topic, you can find basic documentation on how to configure the DPM backup software for a tape gateway and perform a backup. For detailed information about how to use DPM, see the DPM documentation on the Microsoft System Center website. In this topic, you can find out how to configure storage, write data to a tape, archive a tape in the VTS, and restore the data.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

- Configuring DPM to Recognize VTL Devices (p. 348)
- Importing a Tape into DPM (p. 349)
- Writing Data to a Tape in DPM (p. 350)
- Archiving a Tape by Using DPM (p. 350)
- Restoring Data from a Tape Archived in DPM (p. 350)

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 Step 7: Connect Your Tape Gateway Devices to Your Windows Client (p. 326).

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 (p. 274).

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

To map VTL 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 Step 6: Create Virtual Tapes Using the AWS Storage Gateway Console (p. 324).
2. Import the tape into the DPM library. For information on how to do this, see Importing a Tape into DPM (p. 349).
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 2012 R2\DPM\DPM\Config\DPMLA.xml`.

   Note
   If this file exists, the DPMDriveMappingTool will overwrite 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 2012 R2\DPM\DPM\Bin`, and run the following command.

   ```
   C:\Microsoft System Center 2012 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 software library.

To import tapes into the DPM backup software 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.

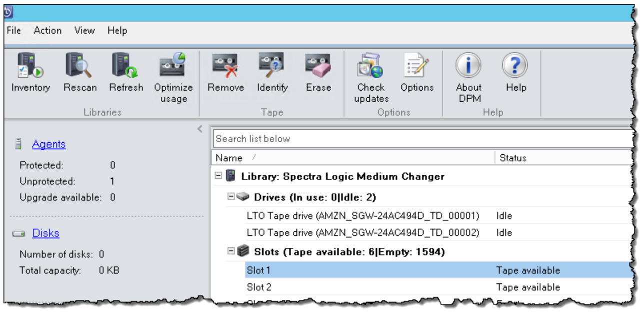

Protected:	0
Unprotected:	1
Upgrade available:	0
Number of disks:	0
Total capacity:	0 KB

Name	Status
Library: Spectra Logic Medium Changer	
Drives (In use: 0|Idle: 2)	
LTO Tape drive (AMZN_SGW-24AC494D_TD_00001)	Idle
LTO Tape drive (AMZN_SGW-24AC494D_TD_00002)	Idle
Slots (Tape available: 6|Empty: 1594)	
Slot 1	Tape available
Slot 2	Tape available

2. 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.

3. 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.

   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.

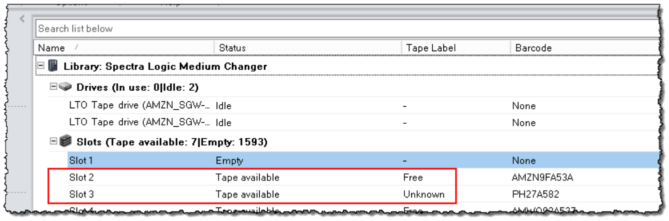

Name	Status	Tape Label	Barcode
Library: Spectra Logic Medium Changer			
Drives (In use: 0|Idle: 2)			
LTO Tape drive (AMZN_SGW-...	Idle	-	None
LTO Tape drive (AMZN_SGW-...	Idle	-	None
Slots (Tape available: 7|Empty: 1593)			
Slot 1	Empty	-	None
Slot 2	Tape available	Free	AMZN9FA53A
Slot 3	Tape available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

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.

Archiving a Tape by Using DPM

When you archive a tape, AWS Storage Gateway moves the tape from the DPM tape library to the VTS—that is, the offline storage. You begin tape archival by removing the tape from the slot using your backup software—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)**.

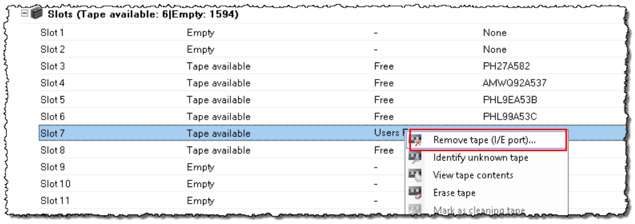

Slots (Tape available: 6|Empty: 1594)

Slot 1	Empty	-	None
Slot 2	Empty	-	None
Slot 3	Tape available	Free	PH27A582
Slot 4	Tape available	Free	AMWQ92A537
Slot 5	Tape available	Free	PHL9EA53B
Slot 6	Tape available	Free	PHL99A53C
Slot 7	Tape available	Users F	
Slot 8	Tape available	Free	
Slot 9	Empty	-	
Slot 10	Empty	-	
Slot 11	Empty	-	

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 to the VTS for archiving.
3. On the AWS 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 your VTS to a tape gateway. For instructions, see Retrieving the Archived Tape from Archive Back to Your Tape Gateway (p. 92).
2. Use the DPM backup software 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

Clean Up Resources You Don't Need (p. 359)

Testing Your Setup by Using Veeam Backup & Replication

You can back up your data to virtual tapes, archive the tapes in a virtual tape shelf (VTS), and manage your virtual tape library (VTL) devices by using Veeam Backup & Replication V7, V8, or V9 Update

2 or later. 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. For detailed information about how to use the Veeam software, see the Veeam Backup & Replication documentation in the Veeam Help Center. In this topic, you can find out how to configure storage, write data to a tape, archive a tape in the VTS, and restore the data.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

Configuring the Veeam Software 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 Step 7: Connect Your Tape Gateway Devices to Your Windows Client (p. 326).

Updating VTL Device Drivers

By default, the Veeam V7 and V8 backup software does not recognize tape gateway devices. 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 (p. 274).

Discovering VTL Devices

For the Veeam 9 backup software, 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 Working with Tape Libraries.

To discover the VTL devices

1. In the Veeam software, choose **Backup Infrastructure**. When the tape gateway is connected, virtual tapes will be listed in the **Backup Infrastructure** tab.

   Note
   Depending on the version of the Veeam Backup & Replication backup software you are using, the user interface might differ somewhat from that shown in the screenshots in this documentation.

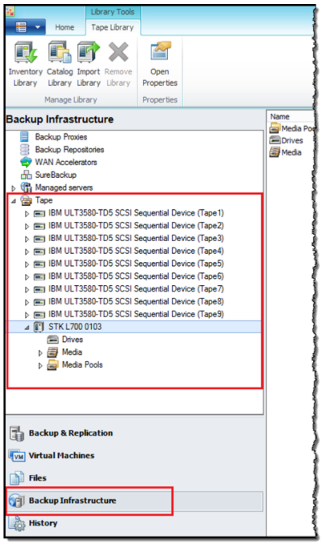
Library Tools
Home
Tape Library
Inventory Library
Catalog Library
Import Library
Remove Library
Open Properties
Manage Library
Properties
Backup Infrastructure
Backup Proxies
Backup Repositories
WAN Accelerators
SureBackup
Managed servers
Tape
IBM ULT3580-TD5 SCSI Sequential Device (Tape1)
IBM ULT3580-TD5 SCSI Sequential Device (Tape2)
IBM ULT3580-TD5 SCSI Sequential Device (Tape3)
IBM ULT3580-TD5 SCSI Sequential Device (Tape4)
IBM ULT3580-TD5 SCSI Sequential Device (Tape5)
IBM ULT3580-TD5 SCSI Sequential Device (Tape6)
IBM ULT3580-TD5 SCSI Sequential Device (Tape7)
IBM ULT3580-TD5 SCSI Sequential Device (Tape8)
IBM ULT3580-TD5 SCSI Sequential Device (Tape9)
STK L700 0103
Drives
Media
Media Pools
Name
Media Po
Drives
Media
Backup & Replication
Virtual Machines
Files
Backup Infrastructure
History

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 will 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 the Veeam Software

You are now ready to import tapes from your tape gateway into the Veeam backup software 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 the Veeam Software

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 Veeam documentation in the Veeam Help Center.

Archiving a Tape by Using the Veeam Software

When you archive a tape, AWS Storage Gateway moves the tape from the Veeam tape library to the VTS—that is, 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 VTS by using your backup software—that is, the Veeam software.

To archive a tape in the Veeam library

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

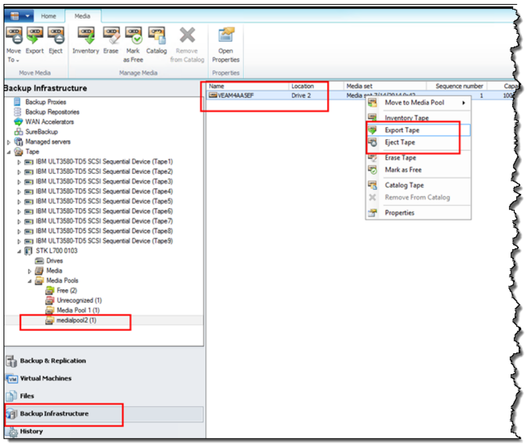

2. Open the context (right–click) menu for the tape that you want to archive, and then choose **Eject Tape**.
3. For **Ejecting tape** box, choose **Close**. The location of the tape changes from a tape drive to a slot.
4. Open the context (right–click) menu for the tape again, and then choose **Export**.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

Restoring your archived data is a two-step process.

To restore data from an archived tape

1. Retrieve the archived tape from your VTS to a tape gateway. For instructions, see Retrieving the Archived Tape from Archive Back to Your Tape Gateway (p. 92).
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 Data from Tape in the Veeam Help Center.

Next Step

Clean Up Resources You Don't Need (p. 359)

Testing Your Setup by Using Dell NetVault Backup

You can back up your data to virtual tapes, archive the tapes in a virtual tape shelf (VTS), and manage your virtual tape library (VTL) devices by using Dell NetVault Backup version 10.0. In this topic, you can find basic documentation on how to configure the NetVault Backup software for a tape gateway and perform a backup. For detailed information about how to use the NetVault Backup software, see the NetVault Backup 10.0.1 – Administration Guide. In this topic, you can find out how to configure storage devices, write data to a tape, archive a tape in the VTS, and restore the data. For additional setup information, see Backing up to Amazon AWS with Dell NetVault Backup on the Dell website.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

Configuring the NetVault Backup Software to Work with VTL Devices

After you have connected the virtual tape library (VTL) devices to the Windows client, you configure NetVault Backup to recognize your devices. For information about how to connect VTL devices to the Windows client, see Step 7: Connect Your Tape Gateway Devices to Your Windows Client (p. 326).

By default, the NetVault Backup software does not automatically recognize tape gateway devices. You must manually add the devices to expose them to the NetVault Backup software and then discover the VTL devices.

Adding VTL Devices

To add the VTL devices

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

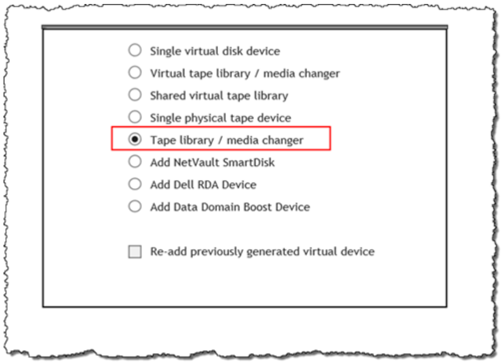

○ Single virtual disk device
○ Virtual tape library / media changer
○ Shared virtual tape library
○ Single physical tape device
● Tape library / media changer
○ Add NetVault SmartDisk
○ Add Dell RDA Device
○ Add Data Domain Boost Device

□ Re-add previously generated virtual device

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 will be displayed. In this case, your medium changer is displayed in the device box.
6. Select your medium changer and choose **Next**. Detailed information about the device is displayed in the wizard.
7. On the Add Tapes to Bays page, select **Scan For Devices**, choose your client machine, and then choose **Next**.

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

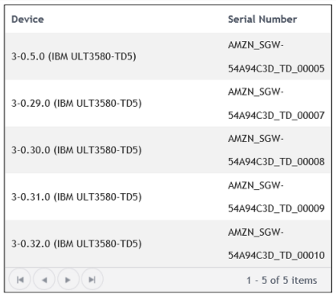

Device	Serial Number
3-0.5.0 (IBM ULT3580-TD5)	AMZN_SGW-54A94C3D_TD_00005
3-0.29.0 (IBM ULT3580-TD5)	AMZN_SGW-54A94C3D_TD_00007
3-0.30.0 (IBM ULT3580-TD5)	AMZN_SGW-54A94C3D_TD_00008
3-0.31.0 (IBM ULT3580-TD5)	AMZN_SGW-54A94C3D_TD_00009
3-0.32.0 (IBM ULT3580-TD5)	AMZN_SGW-54A94C3D_TD_00010

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 you added.

Backing Up Data to a Tape in the NetVault Backup Software

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 Dell NetVault Backup documentation.

Archiving a Tape by Using the NetVault Backup Software

When you archive a tape, AWS Storage Gateway moves the tape from the NetVault Backup tape library to the VTS— that is, 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 VTS by using your backup software —that is, the NetVault Backup software.

To archive a tape in NetVault Backup

1. In the 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.

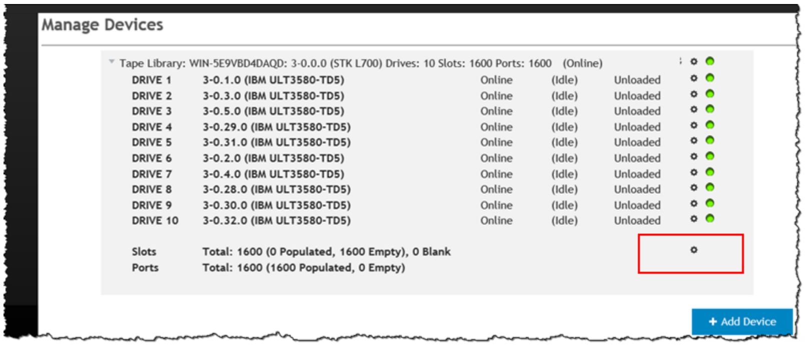

Manage Devices

Tape Library: WIN-5E9VBD4DAQD: 3-0.0.0 (STK L700) Drives: 10 Slots: 1600 Ports: 1600 (Online)

DRIVE 1	3-0.1.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 2	3-0.3.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 3	3-0.5.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 4	3-0.29.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 5	3-0.31.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 6	3-0.2.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 7	3-0.4.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 8	3-0.28.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 9	3-0.30.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded
DRIVE 10	3-0.32.0 (IBM ULT3580-TD5)	Online	(Idle)	Unloaded

Slots Total: 1600 (0 Populated, 1600 Empty), 0 Blank
Ports Total: 1600 (1600 Populated, 0 Empty)

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

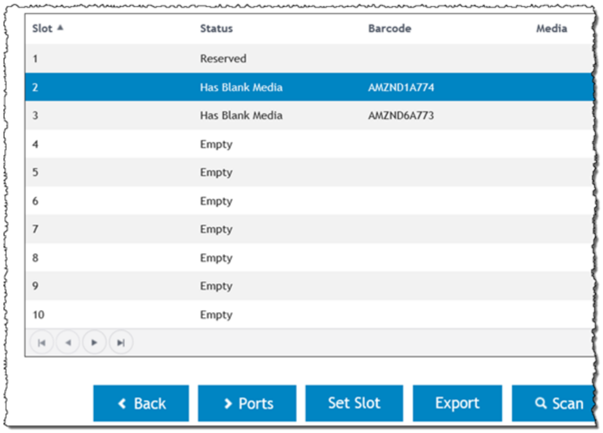

Slot ▲	Status	Barcode	Media
1	Reserved		
2	Has Blank Media	AMZND1A774	
3	Has Blank Media	AMZND6A773	
4	Empty		
5	Empty		
6	Empty		
7	Empty		
8	Empty		
9	Empty		
10	Empty		

Restoring Data from a Tape Archived in the NetVault Backup Software

Restoring your archived data is a two-step process.

To restore data from an archived tape

1. Retrieve the archived tape from your VTS to a tape gateway. For instructions, see Retrieving the Archived Tape from Archive Back to Your Tape Gateway (p. 92).
2. Use the NetVault Backup 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 NetVault Backup 10.0.1 – Administration Guide (Creating a restore job) in the NetVault Backup documentation.

Next Step

Clean Up Resources You Don't Need (p. 359)

Testing Your Setup by Using EMC NetWorker

You can back up your data to virtual tapes, archive the tapes in a virtual tape shelf (VTS), and manage your virtual tape library (VTL) devices by using EMC NetWorker version 8.1 or 8.2. In this topic, you can find basic documentation on how to configure the 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 in the VTS, and restore data from a tape.

For detailed information about how to install and use the EMC NetWorker software, see the *EMC NetWorker Administration Guide*.

For more information about compatible backup software, see Compatible Third-Party Backup Software for Tape Gateway (p. 17).

Topics

Configuring the EMC NetWorker Software to Work with VTL Devices

After you have connected your virtual tape library (VTL) devices to your Microsoft Windows client, you configure EMC NetWorker to recognize your devices. For information about how to connect VTL devices to the Windows client, see Step 7: Connect Your Tape Gateway Devices to Your Windows Client (p. 326).

EMC NetWorker 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 EMC NetWorker software and that you are familiar with the EMC NetWorker Management Console. For more information about the EMC NetWorker Management Console, see the NetWorker Management Console interface section of the *EMC NetWorker Administration Guide*.

The following screenshot shows the EMC NetWorker Management Console.

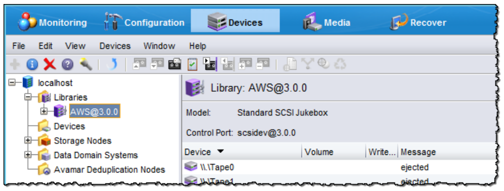

Library: AWS@3.0.0

Model: Standard SCSI Jukebox

Control Port: scsidev@3.0.0

Device	Volume	Write...	Message
\\.\Tape0			ejected
\\.\Tape1			ejected

To configure the EMC NetWorker software for VTL devices

1. Start the 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. This process might take a few seconds to load the devices into the library. In the example shown preceding, we have one library (`AWS@.3.0.0`).
6. Open the context (right-click) menu for your library, and then choose **Configure All Libraries**.
7. In the **Provide General Configuration Information** box, choose the configuration settings you want, and then choose **Next**.
8. In the **Select Target Storage Nodes** box, verify that a storage node is selected, and then choose **Start Configuration**.
9. In the Start Configuration wizard, choose **Finish**.
10. Choose your library to see your tapes in the left pane and the corresponding empty volume slots list in the right pane.

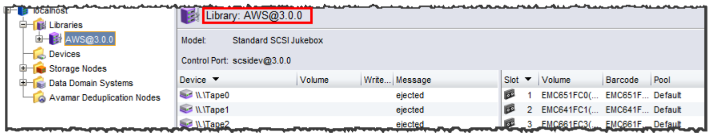

Library: AWS@3.0.0

Model: Standard SCSI Jukebox

Control Port: scsidev@3.0.0

Device	Volume	Write...	Message
\\.\Tape0			ejected
\\.\Tape1			ejected
\\.\Tape2			ejected

Slot	Volume	Barcode	Pool
1	EMC651FC0(...	EMC651F...	Default
2	EMC641FC1(...	EMC641F...	Default
3	EMC661FC3(...	EMC661F...	Default

11. In the volume list, select the volumes you want to enable (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.
12. In the dialog box that appears, choose **Yes**, and then in the **Load the Cartridges into** dialog box, choose **Yes**.
13. If you don't have any more tapes to deposit, choose **No** or **Ignore**. Otherwise, choose **Yes** to deposit additional tapes.

Enabling Import of WORM Tapes into EMC NetWorker

You are now ready to import tapes from your tape gateway into the EMC NetWorker library.

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

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

> **Note**
> You use the 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.

Archiving a Tape in EMC NetWorker

When you archive a tape, AWS Storage Gateway moves the tape from the EMC NetWorker tape library to the VTS—that is, the offline storage. You begin tape archival by ejecting from the tape drive to the storage slot and then withdrawing the tape from the slot to the VTS by using your backup software—that is, the EMC NetWorker software.

To archive a tape by using 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 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 an Archived Tape in EMC NetWorker

Restoring your archived data is a two-step process:

1. Retrieve the archived tape from your VTS to a tape gateway. For instructions, see Retrieving the Archived Tape from Archive Back to Your Tape Gateway (p. 92).
2. Use the 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 *EMC NetWorker Administration Guide.*

Next Step

Clean Up Resources You Don't Need (p. 359)

Clean Up Resources You Don't Need

If you created the gateway as 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? (p. 97).

To clean up resources you don't need

1. Delete tapes from both your gateway's virtual tape library (VTL) and from the virtual tape shelf (VTS): For more information, see Deleting Your Gateway by Using the AWS Storage Gateway Console and Removing Associated Resources (p. 186).

   a. Cancel retrieval for any tapes that have the **RETRIEVING** status in your gateway's VTL. For instructions, see Canceling Tape Retrieval (p. 277).

   b. Archive any tapes that have the **RETRIEVED** status in your gateway's VTL. For instructions, see Archiving Tapes (p. 276).

   c. Delete any remaining tapes from your gateway's VTL. For instructions, see Deleting Tapes (p. 117).

   d. Delete any tapes you have in the VTS. For instructions, see Deleting Tapes (p. 117).
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 (p. 186).
3. Delete the AWS Storage Gateway VM from your on-premises host. If you created your gateway on an Amazon EC2 instance, terminate the instance.

Next Step

Where Do I Go from Here? (p. 97)

API Reference for AWS 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 AWS Storage Gateway, see Regions and Endpoints.

> **Note**
> You can also use the AWS SDKs when developing applications with AWS Storage Gateway. The AWS SDKs for Java, .NET, and PHP wrap the underlying AWS Storage Gateway API, simplifying your programming tasks. For information about downloading the SDK libraries, see Sample Code Libraries.

Topics

AWS Storage Gateway Required Request Headers

This section describes the required headers that you must send with every POST request to AWS Storage Gateway. You include HTTP headers to identify key information about the request including the operation you want to invoke, the date of the request, and information that indicates the authorization of you as the sender of the request. Headers are case insensitive and the order of the headers is not important.

The following example shows headers that are used in the ActivateGateway operation.

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
Content-Type: application/x-amz-json-1.1
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-
east-1/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-
date;x-amz-target,
 Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.ActivateGateway
```

The following are the headers that must include with your POST requests to AWS Storage Gateway. Headers shown below that begin with "x-amz" are AWS-specific headers. All other headers listed are common header used in HTTP transactions.

Header	Description

`Authorization`	The authorization header contains several of pieces of information about the request that enable AWS Storage Gateway to determine if the request is a valid action for the requester. The format of this header is as follows (line breaks added for readability):
	`Authorization: AWS4-HMAC_SHA456`
	`Credentials=`*`YourAccessKey`*`/`*`yyymmdd`*`/`*`region`*`/storagegateway/aws4_request,`
	`SignedHeaders=content-type;host;x-amz-date;x-amz-target,`
	`Signature=`*`CalculatedSignature`*
	In the preceding syntax, you specify *YourAccessKey*, the year, month, and day (*yyyymmdd*), the *region*, and the *CalculatedSignature*. 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. 362).

`Content-Type`	Use `application/x-amz-json-1.1` as the content type for all requests to AWS Storage Gateway. `Content-Type: application/x-amz-json-1.1`
`Host`	Use the host header to specify the AWS Storage Gateway endpoint where you send your request. For example, `storagegateway.us-east-1.amazonaws.com` is the endpoint for the US East (N. Virginia) region. For more information about the endpoints available for AWS Storage Gateway, see Regions and Endpoints. `Host: storagegateway.`*`region`*`.amazonaws.com`
`x-amz-date`	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 AWS 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	Description
	header are used, the format of the Date header does not have to be ISO8601. `x-amz-date: YYYYMMDD'T'HHMMSS'Z'`
`x-amz-target`	This header specifies the version of the API and the operation that you are requesting. The target header values are formed by concatenating the API version with the API name and are in the following format. `x-amz-target: StorageGateway_APIversion.operationName` The *operationName* value (e.g. "ActivateGateway") can be found from the API list, API Reference for AWS Storage Gateway (p. 360).

Signing Requests

AWS Storage Gateway requires that you authenticate every request you send by signing the request. To sign a request, you calculate a digital signature using a cryptographic hash function. A cryptographic hash is a function that returns a unique hash value based on the input. The input to the hash function includes the text of your request and your secret access key. The hash function returns a hash value that you include in the request as your signature. The signature is part of the `Authorization` header of your request.

After receiving your request, AWS Storage Gateway recalculates the signature using the same hash function and input that you used to sign the request. If the resulting signature matches the signature in the request, AWS Storage Gateway processes the request. Otherwise, the request is rejected.

AWS Storage Gateway supports authentication using AWS Signature Version 4. The process for calculating a signature can be broken into three tasks:

- Task 1: Create a Canonical Request

  Rearrange your HTTP request into a canonical format. Using a canonical form is necessary because AWS Storage Gateway uses the same canonical form when it recalculates a signature to compare with the one you sent.
- Task 2: Create a String to Sign

  Create a string that you will use as one of the input values to your cryptographic hash function. The string, called the *string to sign*, is a concatenation of the name of the hash algorithm, the request date, a *credential scope* string, and the canonicalized request from the previous task. The *credential scope* string itself is a concatenation of date, region, and service information.
- Task 3: Create a Signature

  Create a signature for your request by using a cryptographic hash function that accepts two input strings: your *string to sign* and a *derived key*. The *derived key* is calculated by starting with your secret access key and using the *credential scope* string to create a series of Hash-based Message Authentication Codes (HMACs).

Example Signature Calculation

The following example walks you through the details of creating a signature for ListGateways. 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 (N. Virginia) region.

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

```
POST / HTTP/1.1
Host: storagegateway.us-east-1.amazonaws.com
x-amz-Date: 20120910T000000Z
Authorization: SignatureToBeCalculated
Content-type: application/x-amz-json-1.1
x-amz-target: StorageGateway_20120630.ListGateways
{}
```

The canonical form of the request calculated for Task 1: Create a Canonical Request (p. 362) is:

```
POST
/

content-type:application/x-amz-json-1.1
host:storagegateway.us-east-1.amazonaws.com
x-amz-date:20120910T000000Z
x-amz-target:StorageGateway_20120630.ListGateways

content-type;host;x-amz-date;x-amz-target
44136fa355b3678a1146ad16f7e8649e94fb4fc21fe77e8310c060f61caaff8a
```

The last line of the canonical request is the hash of the request body. Also, note the empty third line in the canonical request. This is because there are no query parameters for this API (or any AWS Storage Gateway APIs).

The *string to sign* for Task 2: Create a String to Sign (p. 362) is:

```
AWS4-HMAC-SHA256
20120910T000000Z
20120910/us-east-1/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 (p. 362), the *derived key* can be represented as:

```
derived key = HMAC(HMAC(HMAC(HMAC("AWS4" +
 YourSecretAccessKey,"20120910"),"us-
east-1"),"storagegateway"),"aws4_request")
```

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

```
6d4c40b8f2257534dbdca9f326f147a0a7a419b63aff349d9d9c737c9a0f4c81
```

The final step is to construct the `Authorization` header. For the demonstration access key AKIAIOSFODNN7EXAMPLE, the header (with line breaks added for readability) is:

```
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120910/us-
east-1/storagegateway/aws4_request,
SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=6d4c40b8f2257534dbdca9f326f147a0a7a419b63aff349d9d9c737c9a0f4c81
```

Error Responses

Topics

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, AWS Storage Gateway may return only just an exception, or it may return both an exception and an operation error code. Examples of error responses are shown in the Error Responses (p. 377).

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 (p. 365) message codes that give the specific operation error code.

Exception	Message	HTTP Status Code
`IncompleteSignatureException`	The specified signature is incomplete.	400 Bad Request
`InternalFailure`	The request processing has failed due to some unknown error, exception or failure.	500 Internal Server Error
`InternalServerError`	One of the operation error code messages Operation Error Codes (p. 365).	500 Internal Server Error
`InvalidAction`	The requested action or operation is invalid.	400 Bad Request
`InvalidClientTokenId`	The X.509 certificate or AWS Access Key ID provided does not exist in our records.	403 Forbidden

Exception	Message	HTTP Status Code
InvalidGatewayRequestException	One of the operation error code messages in Operation Error Codes (p. 365).	400 Bad Request
InvalidSignatureException	The request signature we calculated does not match the signature you provided. Check your AWS Access Key and signing method.	400 Bad Request
MissingAction	The request is missing an action or operation parameter.	400 Bad Request
MissingAuthenticationToken	The request must contain either a valid (registered) AWS Access Key ID or X.509 certificate.	403 Forbidden
RequestExpired	The request is past the expiration date or the request date (either with 15 minute padding), or the request date occurs more than 15 minutes in the future.	400 Bad Request
SerializationException	An error occurred during serialization. Check that your JSON payload is well-formed.	400 Bad Request
ServiceUnavailable	The request has failed due to a temporary failure of the server.	503 Service Unavailable
SubscriptionRequiredException	The AWS Access Key Id needs a subscription for the service.	400 Bad Request
ThrottlingException	Rate exceeded.	400 Bad Request
UnknownOperationException	An unknown operation was specified. Valid operations are listed in Operations in AWS Storage Gateway (p. 379).	400 Bad Request
UnrecognizedClientException	The security token included in the request is invalid.	400 Bad Request
ValidationException	The value of an input parameter is bad or out of range.	400 Bad Request

Operation Error Codes

The following table shows the mapping between AWS Storage Gateway operation error codes and APIs that can return the codes. All operation error codes are returned with one of two general exceptions—InternalServerError and InvalidGatewayRequestException—described in Exceptions (p. 364).

Operation Error Code	Message	Operations That Return this Error Code
`ActivationKeyExpired`	The specified activation key has expired.	ActivateGateway
`ActivationKeyInvalid`	The specified activation key is invalid.	ActivateGateway
`ActivationKeyNotFound`	The specified activation key was not found.	ActivateGateway
`BandwidthThrottleScheduleNotFound`	The specified bandwidth throttle was not found.	DeleteBandwidthRateLimit
`CannotExportSnapshot`	The specified snapshot cannot be exported.	CreateCachediSCSIVolume CreateStorediSCSIVolume
`InitiatorNotFound`	The specified initiator was not found.	DeleteChapCredentials
`DiskAlreadyAllocated`	The specified disk is already allocated.	AddCache AddUploadBuffer AddWorkingStorage CreateStorediSCSIVolume
`DiskDoesNotExist`	The specified disk does not exist.	AddCache AddUploadBuffer AddWorkingStorage CreateStorediSCSIVolume
`DiskSizeNotGigAligned`	The specified disk is not gigabyte-aligned.	CreateStorediSCSIVolume
`DiskSizeGreaterThanVolumeMaxSize`	The specified disk size is greater than the maximum volume size.	CreateStorediSCSIVolume
`DiskSizeLessThanVolumeSize`	The specified disk size is less than the volume size.	CreateStorediSCSIVolume
`DuplicateCertificateInfo`	The specified certificate information is a duplicate.	ActivateGateway

Operation Error Code	Message	Operations That Return this Error Code
`GatewayInternalError`	A gateway internal error occurred.	AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateStorediSCSIVolume CreateSnapshotFromVolumeRecoveryPoint DeleteBandwidthRateLimit DeleteChapCredentials DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes DescribeWorkingStorage ListLocalDisks 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 ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewaySoftwareNow UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
`GatewayNotFound`	The specified gateway was not found.	AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteGateway DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes DescribeWorkingStorage 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 ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewaySoftwareNow UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
`InternalError`	An internal error occurred.	ActivateGateway AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteGateway DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes DescribeWorkingStorage ListLocalDisks ListGateways ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime

Operation Error Code	Message	Operations That Return this Error Code
		UpdateGatewayInformation UpdateGatewaySoftwareNow UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
`InvalidParameters`	The specified request contains invalid parameters.	ActivateGateway AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteGateway DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes DescribeWorkingStorage ListLocalDisks ListGateways ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime

Operation Error Code	Message	Operations That Return this Error Code
		UpdateGatewayInformation UpdateGatewaySoftwareNow UpdateSnapshotSchedule
`LocalStorageLimitExceeded`	The local storage limit was exceeded.	AddCache AddUploadBuffer AddWorkingStorage
`LunInvalid`	The specified LUN is invalid.	CreateStorediSCSIVolume
`MaximumVolumeCountExceeded`	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 DescribeWorkingStorage ListLocalDisks ListGateways ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime

Operation Error Code	Message	Operations That Return this Error Code
		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 invalid.	CreateCachediSCSIVolume CreateStorediSCSIVolume
`StagingAreaFull`	The staging area is full.	CreateCachediSCSIVolume CreateStorediSCSIVolume
`TargetAlreadyExists`	The specified target already exists.	CreateCachediSCSIVolume CreateStorediSCSIVolume
`TargetInvalid`	The specified target is invalid.	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 invalid.	DeleteVolume
VolumeInUse	The specified volume is already in use.	DeleteVolume
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 (p. 364).

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 AWS Storage Gateway calculates a signature that does not match the signature sent with a request.

```
{
  "__type": "InvalidSignatureException",
  "message": "The request signature we calculated does not match the
 signature you provided."
}
```

Operations in AWS Storage Gateway

For a list of AWS Storage Gateway operations, see Actions in the *AWS Storage Gateway API Reference*.

Document History for AWS Storage Gateway

The following table describes important changes to the documentation since the last release of the *AWS Storage Gateway User Guide*.

- **API version**: 2013-06-30
- **Latest documentation update**: December 08, 2016

Change	Description	Date Changed
New region	AWS Storage Gateway is now available in the Canada (Central) region. For detailed information, see Regions (p. 10).	In this release
Support for File Gateway	In addition to volume gateways and tape gateway, AWS Storage Gateway now provides File Gateway. File Gateway combines a service and virtual software appliance, enabling 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 a 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 Amazon Glacier. For more information, see Testing Your Setup by Using Backup Exec (p. 64).	November 7, 2016
Compatibility with HPE Data Protector 9.x	Tape Gateway is now compatible with HPE Data Protector 9.x. You can now use HPE Data Protector to back up your data to Amazon S3 and archive directly to Amazon Glacier. For more information, see Testing Your Setup by Using HPE Data Protector (p. 73).	November 2, 2016
New region	AWS Storage Gateway is now available in the US East (Ohio) region. For detailed information, see Regions (p. 10).	October 17, 2016

Change	Description	Date Changed
AWS Storage Gateway console redesign	The AWS 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 (p. 10).	August 30, 2016
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 Amazon Glacier. For more information, see Testing Your Setup by Using Veeam Backup & Replication (p. 94).	August 15, 2016
Longer volume and snapshot IDs	AWS Storage Gateway is introducing longer IDs for volumes and snapshots. You can enable the longer ID format for your volumes, snapshots, and other supported AWS resources. For more information, see Understanding AWS Storage Gateway Resources and Resource IDs (p. 290).	April 25, 2016
New region Support for storage up to 512 TiB in size for stored volumes Other gateway updates and enhancements to the AWS Storage Gateway local console	Tape Gateway is now available in the Asia Pacific (Seoul) region. For more information, see Regions (p. 10). 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 Volume Architecture (p. 5) and AWS Storage Gateway Limits (p. 296). Total size of all tapes in a virtual tape library is increased to 1 PiB. For more information, see AWS Storage Gateway Limits (p. 296). You can now set the password for your VM local console on the AWS Storage Gateway Console. For information, see Setting the Local Console Password from the Storage Gateway Console (p. 161).	March 21, 2016
Compatibility with for EMC NetWorker 8.x	Tape Gateway is now compatible with EMC NetWorker 8.x. You can now use EMC NetWorker to back up your data to Amazon S3 and archive directly to Amazon Glacier. For more information, see Testing Your Setup by Using EMC NetWorker (p. 70).	February 29, 2016

Change	Description	Date Changed
Support for VMware ESXi Hypervisor version 6.0 and Red Hat Enterprise Linux 7 iSCSI initiator Content restructure	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 (p. 16) and Supported iSCSI Initiators (p. 16). 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 (p. 99).	October 20, 2015
Support for storage up to 1,024 TiB in size for cached volumes Support for the VMXNET3 (10 GbE) network adapter type in VMware ESXi hypervisor Performance enhancements Miscellaneous enhancements and updates to the AWS Storage Gateway local console	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 Volume Architecture (p. 4) and AWS Storage Gateway Limits (p. 296). 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. 171). The maximum upload rate for AWS Storage Gateway has increased to 120 MB a second, and the maximum download rate has increased to 20 MB a second. For more information, see Configuration and Performance Limits (p. 297). The AWS 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. 162).	September 16, 2015
Support for tagging	AWS 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. 294).	September 2, 2015
Compatibility with Dell NetVault Backup 10.0	Tape Gateway is now compatible with Dell NetVault Backup 10.0. You can now use Dell NetVault Backup 10.0 to back up your data to Amazon S3 and archive directly to Amazon Glacier. For more information, see Testing Your Setup by Using Dell NetVault Backup (p. 67).	June 22, 2015

Change	Description	Date Changed
Support for 16 TiB storage volumes for Gateway-stored setups Support for system resource checks on the AWS Storage Gateway local console Support for the Red Hat Enterprise Linux 6 iSCSI initiator	AWS Storage Gateway now supports 16 TiB storage volumes for Gateway-stored setups. You can now create 12 16 TiB storage volumes for a maximum of 192 TiB of storage. For more information, see Stored Volume Architecture (p. 5). 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 (p. 170) or Viewing Your Gateway System Resource Status (p. 170). AWS Storage Gateway now supports the Red Hat Enterprise Linux 6 iSCSI initiator. For more information, see Requirements (p. 11). This release includes the following AWS Storage Gateway improvements and updates: • From the AWS 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 (p. 156). • AWS Storage Gateway now provides an API you can use to list iSCSI initiators connected to your storage volumes. For more information, see ListVolumeInitiators in the API reference.	June 3, 2015
Support for Microsoft Hyper-V hypervisor versions 2012 and 2012 R2	AWS 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 (p. 16).	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 Amazon Glacier. For more information, see Testing Your Setup by Using Backup Exec (p. 64).	April 6, 2015
CHAP authentication support for storage volumes	AWS Storage Gateway now supports configuring CHAP authentication for storage volumes. For more information, see Creating Volumes (p. 32).	April 2, 2015
Support for VMware ESXi Hypervisor version 5.1 and 5.5	AWS 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 (p. 16).	March 30, 2015

Change	Description	Date Changed
Support for Windows CHKDSK utility	AWS 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 (p. 198).	March 04, 2015
Integration with AWS CloudTrail to capture API calls	AWS Storage Gateway is now integrated with AWS CloudTrail. AWS CloudTrail captures API calls made by or on behalf of AWS Storage Gateway in your AWS account and delivers the log files to an Amazon S3 bucket that you specify. For more information, see Logging AWS Storage Gateway API Calls by Using AWS CloudTrail (p. 138). This release includes the following AWS Storage Gateway improvement and update: • 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 (p. 202).	December 16, 2014
Compatibility with additional backup software and medium changer	Tape Gateway is now compatible with the following backup software: • Symantec Backup Exec 2014 • Microsoft System Center 2012 R2 Data Protection Manager • Veeam Backup & Replication V7 • Veeam Backup & Replication V8 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 Amazon Glacier. For more information, see Testing Your Gateway Setup (p. 63). AWS Storage Gateway now provides an additional medium changer that works with the new backup software. This release includes miscellaneous AWS Storage Gateway improvements and updates.	November 3, 2014
EU (Frankfurt) region	AWS Storage Gateway is now available in the EU (Frankfurt) region. For detailed information, see Regions (p. 10).	October 23, 2014

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 gateway-stored, gateway-cached, and tape gateway setups. For more information, see Creating a Tape Gateway (p. 48).	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 Amazon Glacier. For more information, see Testing Your Setup by Using Backup Exec (p. 64).	April 28, 2014
Support for Windows Server Failover Clustering Support for VMware ESX initiator Support for performing configuration tasks on AWS Storage Gateway local console	• Volume gateways can now be used as storage for applications configured using Windows Server Failover Clustering (WSFC). This enables coordinated iSCSI access to AWS Storage Gateway volumes by applications clustered using WSFC. • AWS Storage Gateway now enables you to manage storage connectivity directly through your ESX host. This provides an alternative to using initiators resident in the guest OS of your VMs. • AWS Storage Gateway now provides support for performing configuration tasks in the AWS Storage Gateway local console. For information about performing configuration tasks on gateways deployed on-premises, see Performing Maintenance Tasks on the VM Local Console (p. 156) or Performing Maintenance Tasks on the VM Local Console (p. 156). For information about performing configuration tasks on gateways deployed on an EC2 instance, see Performing Maintenance Tasks on the Amazon EC2 Gateway Local Console (p. 180) or Performing Maintenance Tasks on the Amazon EC2 Gateway Local Console (p. 180).	January 31, 2014

Change	Description	Date Changed
Support for virtual tape library (VTL) and introduction of API version 2013-06-30	AWS 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 (gateway-cached and gateway-stored), AWS 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 *AWS Storage Gateway User Guide*. • For an architectural overview, see Tape Gateway (p. 6). • To get started with tape gateway, see Creating a Tape Gateway (p. 48).	November 5, 2013
Support for Microsoft Hyper-V	AWS Storage Gateway now provides the ability to deploy an on-premises gateway on the Microsoft Hyper-V virtualization platform. Gateways deployed on Microsoft Hyper-V have all the same functionality and features as the existing on-premises storage gateway. To get started deploying a gateway with Microsoft Hyper-V, see Provisioning a Hyper-V Host (p. 51).	April 10, 2013
Support for deploying a gateway on Amazon EC2	AWS Storage Gateway now provides the ability to deploy a gateway in Amazon Elastic Compute Cloud (Amazon EC2). You can launch a gateway instance in Amazon EC2 using the AWS Storage Gateway AMI available in AWS Marketplace. To get started deploying a gateway using the AWS Storage Gateway AMI, see Provisioning an Amazon EC2 Host (p. 53).	January 15, 2013

Change	Description	Date Changed
Support for cached volumes and introduction of API Version 2012-06-30	In this release, AWS 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 gateway-cached volumes is stored in Amazon Simple Storage Service (Amazon S3), with only a cache of recently written and recently read data stored locally on your on-premises storage hardware. 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. In this release, AWS Storage Gateway also introduces a new API version that, in addition to supporting the current operations, provides new operations to support cached volumes. For more information on the two AWS Storage Gateway solutions, see How AWS Storage Gateway Works (Architecture) (p. 3). You can also try a test setup. For instructions, see Creating a Tape Gateway (p. 48).	October 29, 2012
API and IAM Support	In this release, AWS Storage Gateway introduces API support as well as support for AWS Identity and Access Management(IAM). • **API support—**You can now programmatically configure and manage your AWS Storage Gateway resources. For more information about the APIs, see API Reference for AWS Storage Gateway (p. 360) in the *AWS Storage Gateway User Guide*. • **IAM Support** – AWS Identity and Access Management (IAM) enables you create users and manage user access to your AWS Storage Gateway resources by means of IAM policies. For examples of IAM policies, see Authentication and Access Control for AWS Storage Gateway (p. 297). 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 (p. 162).	March 5, 2012
New Guide	This is the first release of *AWS Storage Gateway User Guide*.	January 24, 2012