

Drobo How-To Guide


Cloud Storage Using Amazon Storage Gateway with Drobo iSCSI SAN



The Amazon Web Services (AWS) Storage Gateway uses an on-premises virtual appliance to replicate a portion of your local Drobo iSCSI SAN (Drobo B1200i, left below, and Drobo B800i, right below) to cloud-based storage. This solution provides an option for redundant backup and disaster recovery in the cloud for critical data, particularly for users that may not have a secondary site with IT equipment.



The AWS Storage Gateway is installed on a VMware host that is linked to an AWS account—providing an on-ramp to cloud storage. Volumes of up to 1TB in size (12 volumes per gateway maximum) are created and mounted by on-premises servers as iSCSI devices. Based on functionality available in the Beta version of AWS Storage Gateway, data written to your gateway-stored volumes is recorded on Drobo storage and asynchronously stored in Amazon S3 (Simple Storage Service) in the form of Amazon EBS (Elastic Block Store) snapshots.



This guide provides the tested steps to sync data on Drobo iSCSI storage to the cloud for offsite backup using gateway-stored volumes. The information in this guide can be the foundation for leveraging other Amazon Web Services (AWS) for recovery in the cloud. For example, Amazon Route 53 service enables DNS updating so that users are pointed to the instance running in the cloud. It is also possible to create an Amazon EBS instance from an S3 backup, mount it with an Amazon EC2 instance, and restart the application completely from the cloud.



Topics

- Deploying the AMS Storage Gateway appliance
- Configuring the AMS Storage Gateway appliance
- Activating the AWS Storage Gateway appliance
- Creating AWS storage volumes
- Mounting Storage Gateway volumes
- Restoring storage volume snapshots

What You Will Need

- Drobo iSCSI SAN B800i or B1200i
- Drobo Dashboard management software (latest version)
- Amazon Web Services account: sign up @ <http://aws.amazon.com>
- AWS Storage Gateway virtual machine (VM) image, available from the [AWS Management Console](#)
- Sufficient Internet connectivity to support file synchronization to the cloud
- VMware ESXi Hypervisor v4.1 or above (can work with the free vSphere Hypervisor (ESXi))
- Gateway appliance resources: at least 7.5GB of RAM, 4 virtual processors, and 75GB of disk space
 - Disk space for OVA appliance and system data can be stored on a Drobo iSCSI volume

Prerequisites

This document assumes that the Drobo iSCSI storage has been deployed and that a datastore has been configured in your VMware virtualization infrastructure. If not, to do so, read *How To Deploy VMware and Drobo as a Complete and Cost-Effective Virtualization Solution* @ www.drobo.com/downloads/how-to/HT-0045-00_vmware_drobo_virt_solution.pdf.



Deploying the AWS Storage Gateway Appliance

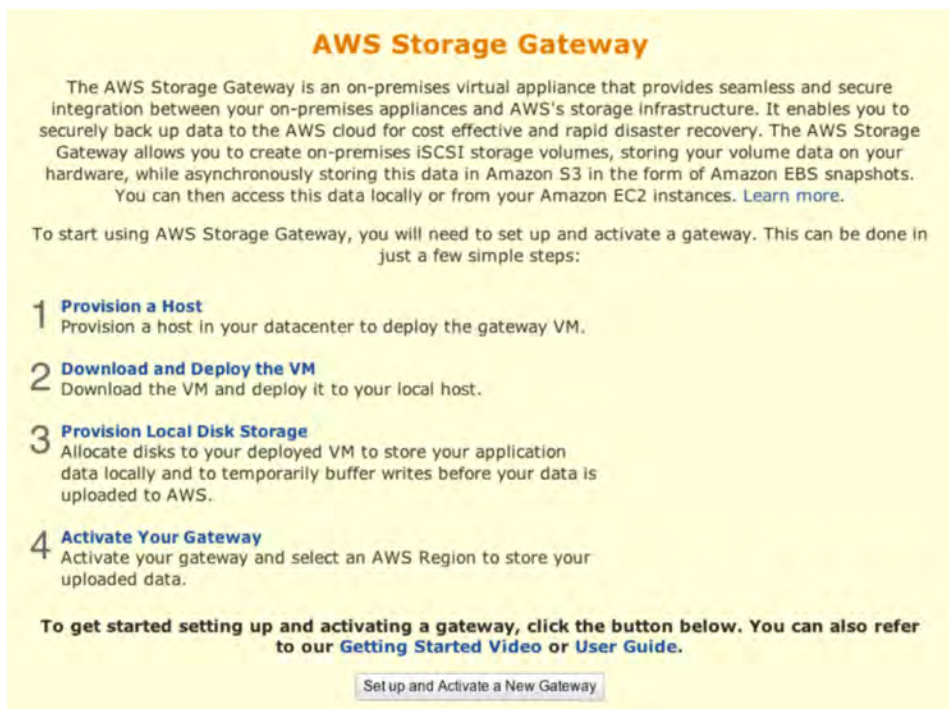
STEP 1



Log in to the AWS Management console:
<http://aws.amazon.com>

Navigate to the **Storage Gateway** tab from the AWS Management console.

STEP 2



To begin the Storage Gateway activation wizard, click **Set up and Activate a New Gateway** at the bottom of the screen.

STEP 3




Click **Download** to save the Storage Gateway OVA Package.

NOTE: The OVA appliance package is over 500MB and may take some time to



download.

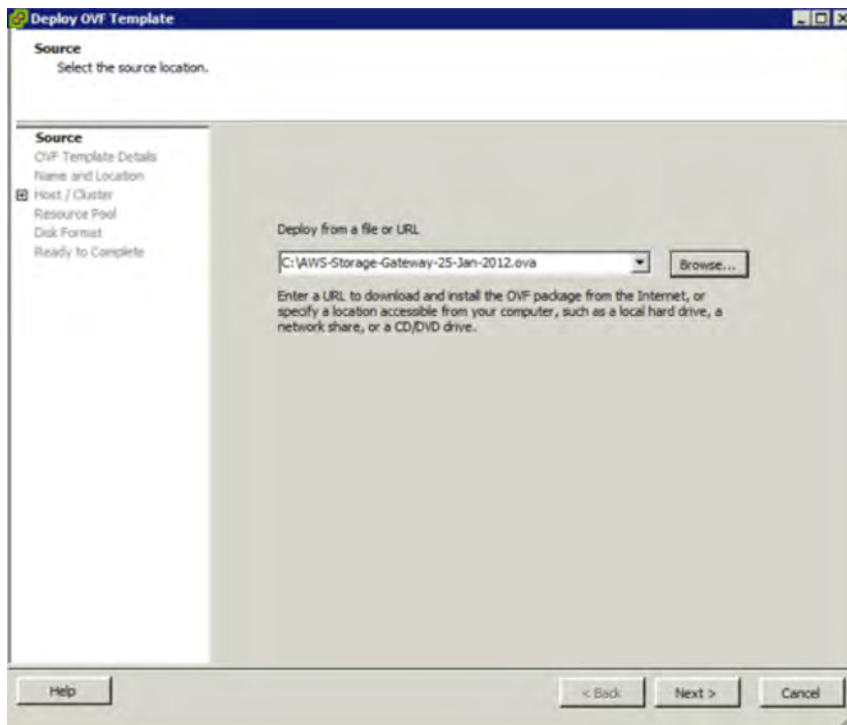
STEP 4



Load the vSphere client and log into the ESXi host or vCenter server that will be used to deploy the appliance.

Once authenticated, select **File > Deploy OVF Template**.

STEP 5

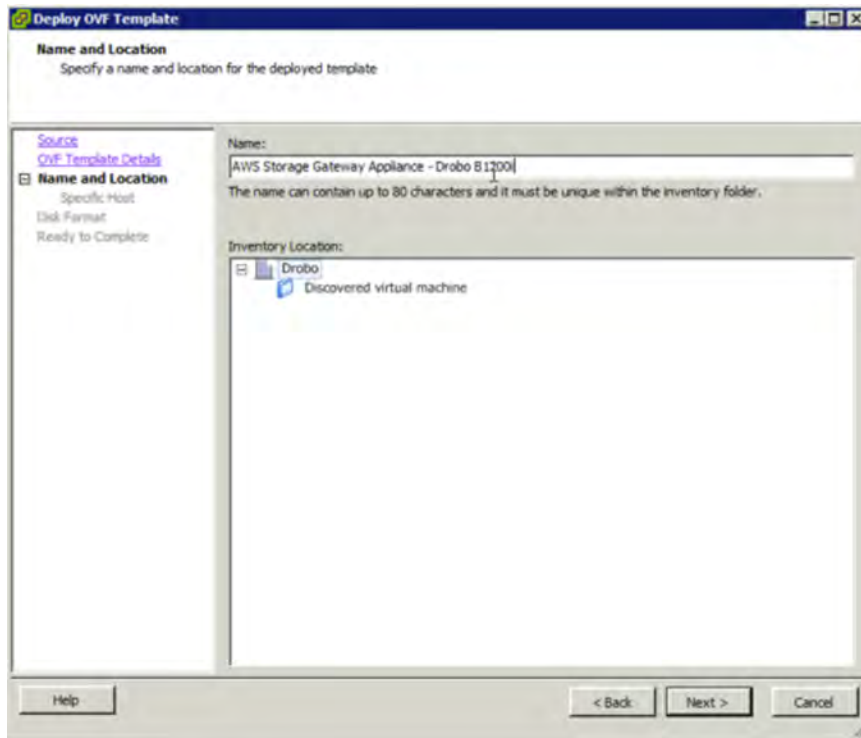


Provide the path to the AWS Storage Gateway OVA package and click **Next**.

On the next screen the details of the appliance are shown. Click **Next** to continue.



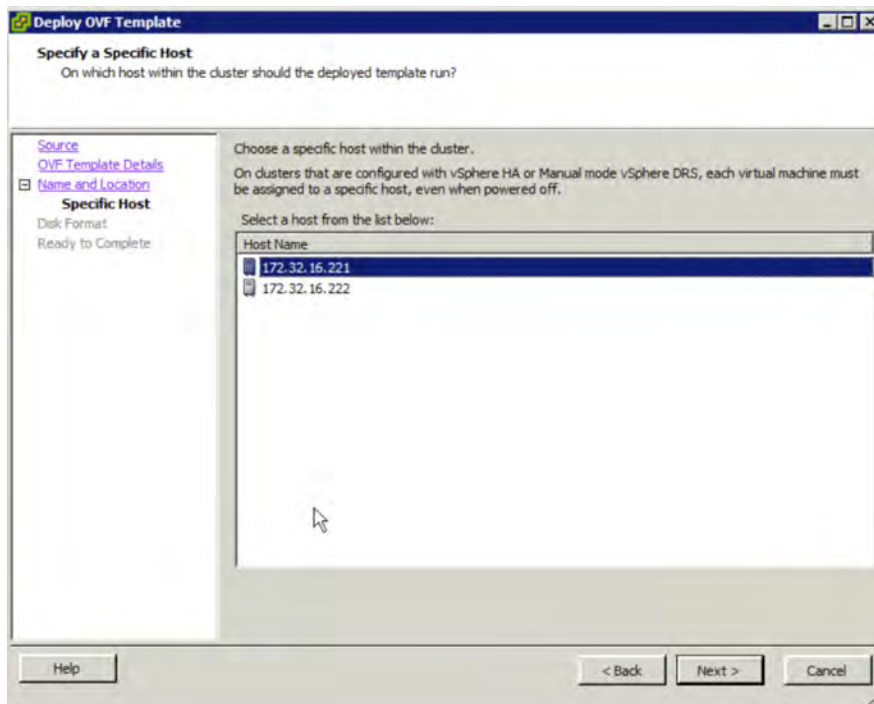
STEP 6



Enter a name for the VM that will be displayed within the vSphere client and click **Next**.

The name provided will only be used in the VM inventory list within VMware.

STEP 7

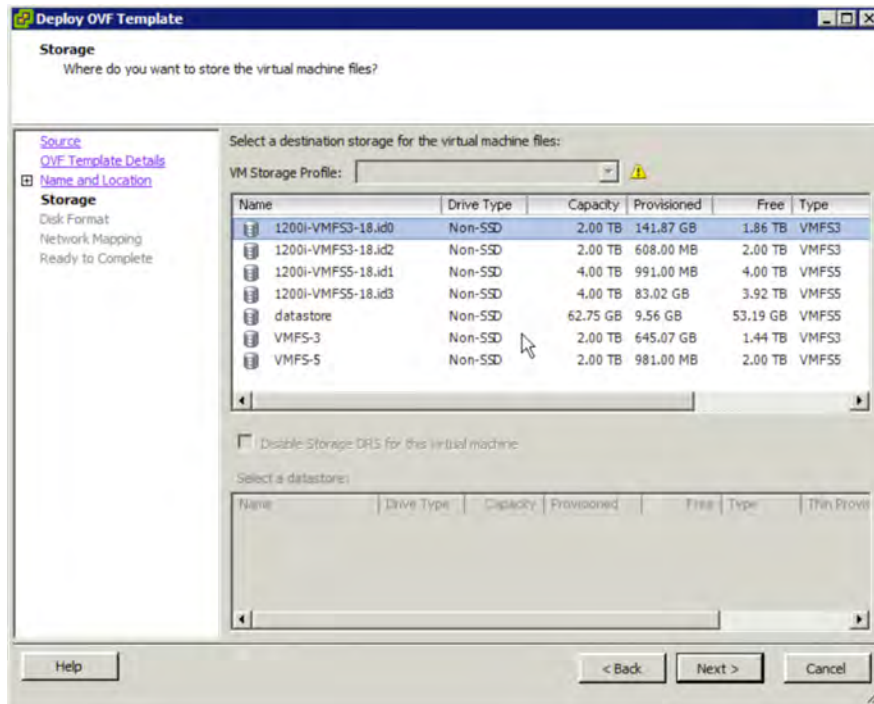


Enter a name for the VM that will be displayed within the vSphere client and click **Next**.

The name provided will only be used in the VM inventory list within VMware.



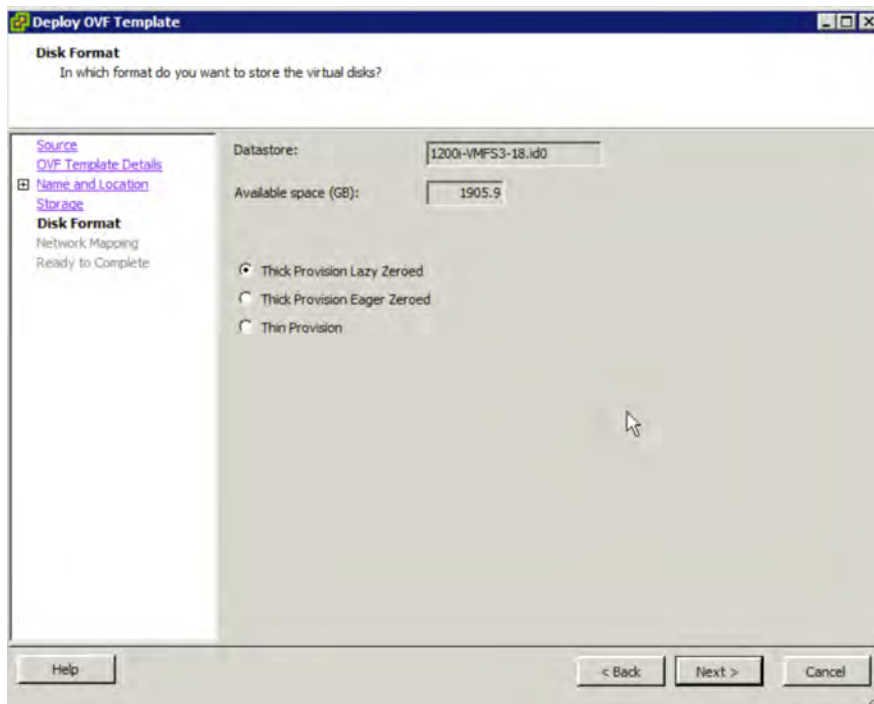
STEP 8



Select the datastore on the Drobo that will be used to store the appliance VM.

Typically this store will also be used to store all disk images associated with the appliance.

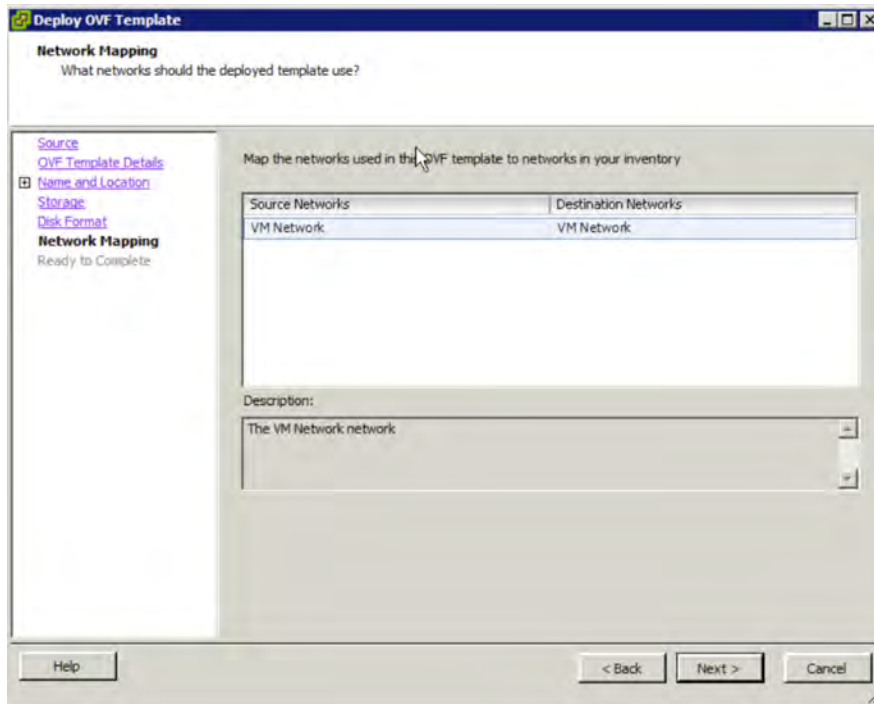
STEP 9



Select **Thick Provision Lazy Zeroed** to allocate the disk storage immediately and click **Next**.

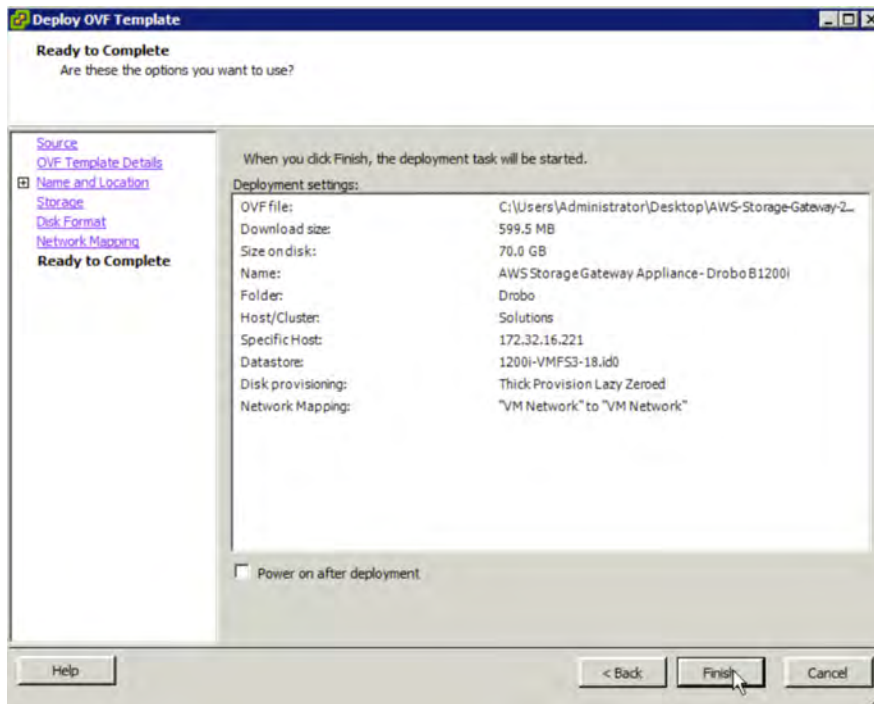


STEP 10



If multiple networks are configured on the ESXi server, select the appropriate network and click **Next**.

STEP 11

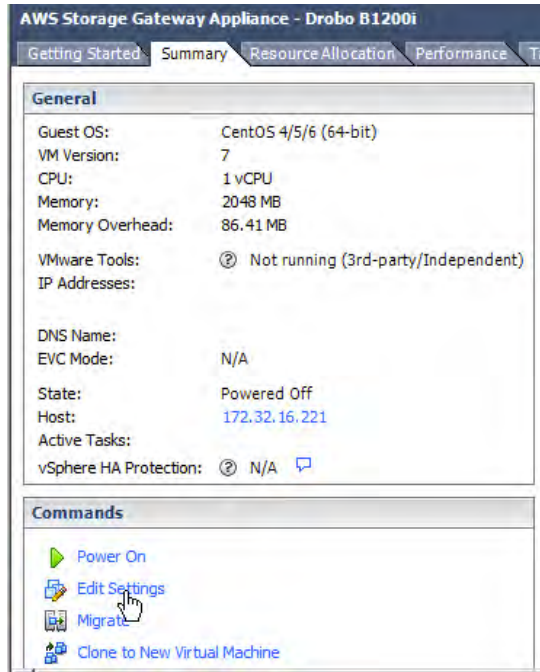


Click **Finish** to complete the wizard and deploy the VM to the host server.



Configuring the AWS Storage Gateway Appliance

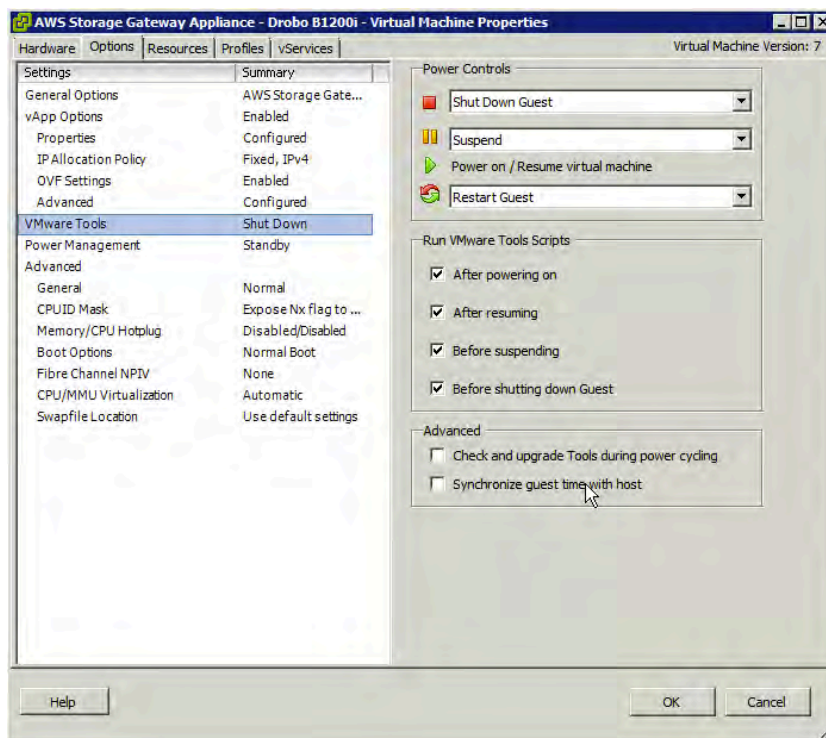
STEP 1



Once the appliance has been deployed, select it from the inventory list.

On the summary pane for the gateway virtual machine select **Edit Settings** to modify the properties of the VM.

STEP 2

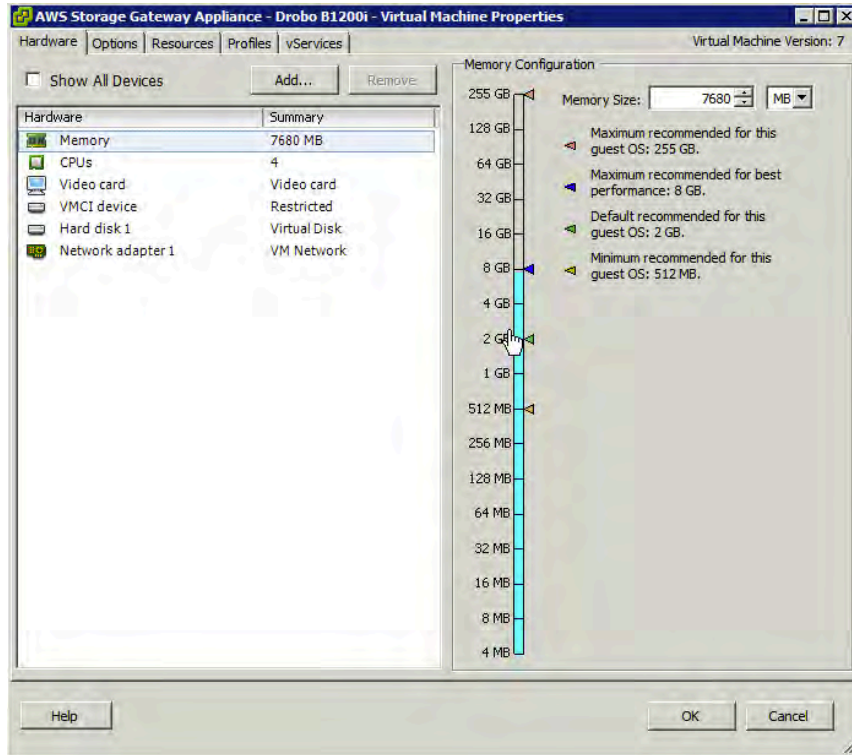


In the Options tab, select **VMware Tools** on the left, check the box **Synchronize guest time with host** in the Advanced section at the right bottom, and select the **Hardware** tab.



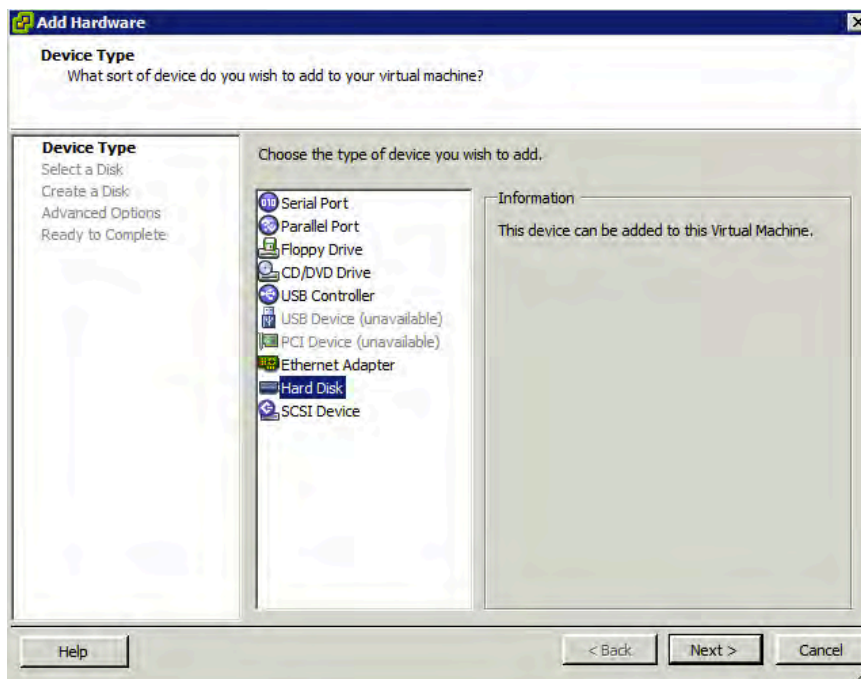
STEP 3

A minimum of two disks must be created for use by the gateway appliance. The first disk will store application data and the second will be used as buffer space for data waiting to be synchronized to AWS.



In the hardware tab, click **Add** to create the disk that will be used by the gateway for data storage.

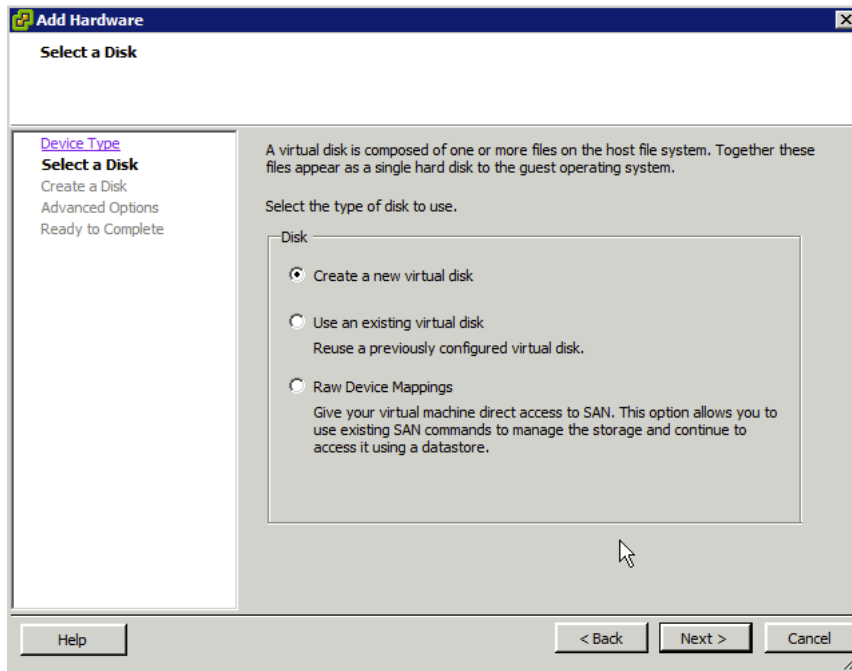
STEP 4



Select **Hard Disk** and click **Next**.



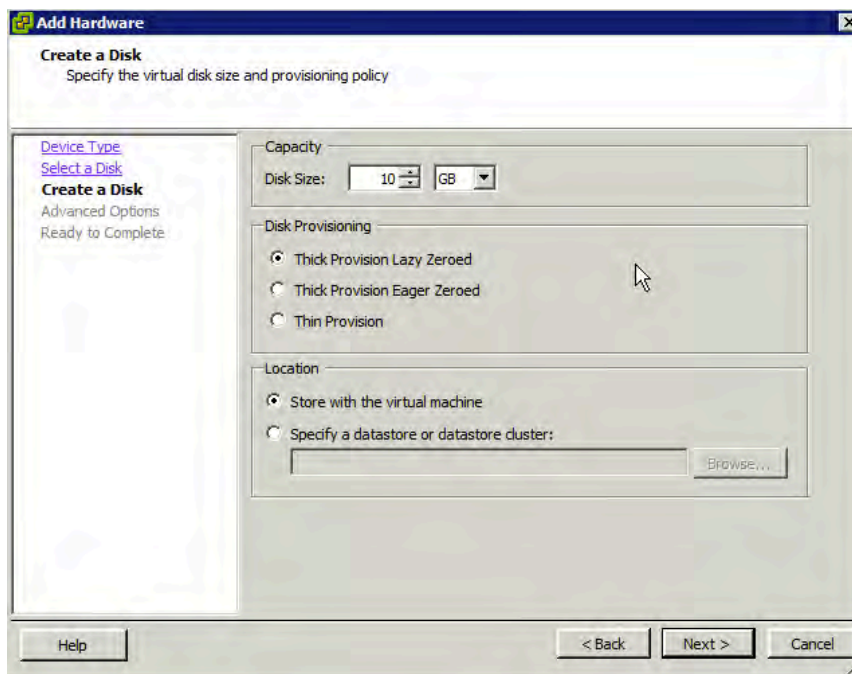
STEP 5



Select **Create a new virtual disk** and click **Next**.

STEP 6

The size of the data disk to be replicated to the AWS should be sufficient to store all data but without being so large there will be a very large amount of unused space. AWS storage is billed per gigabyte and is not thinly provisioned like a Drobo. Data is billed on the size provisioned.



Enter the size for the data disk to be replicated to AWS.

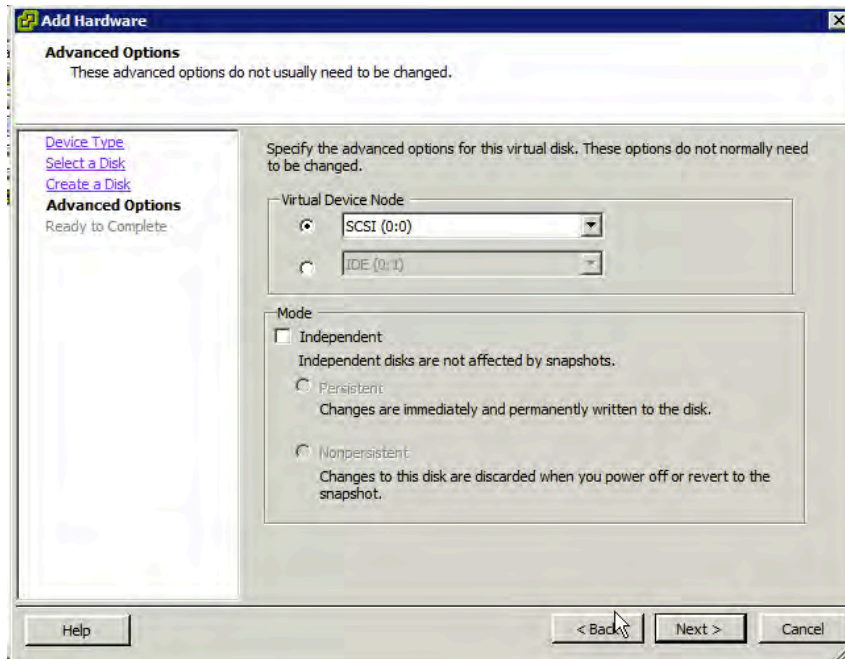
Storage volumes can be any size between 1 and 1000GB.



NOTE: Data can only be restored from AWS in bulk and no granular restoration option is available. When creating data disks, do not store all data in one large volume. Since there is no additional cost, create multiple volumes to break data into manageable chunks.

For example, if you are using the appliance to synchronize file shares, create a volume for each department or group drive. When a snapshot needs to be restored to the on-premises appliance, this will provide the most efficient utilization of limited Internet bandwidth and lower recovery times for higher priority data.

STEP 7



Leave all of the Advanced options set to the default values and click **Next**.

The next screen displays a summary of hardware changes. Click **Finish** to complete the wizard.

STEP 8

Repeat the steps above to create a second disk that will be used as a buffer by the storage gateway. Configure its size relative to the amount of data backed up and the capacity of your upstream Internet connection. *The lower the bandwidth, the larger the buffer needs to be.* As network outages and degradation can occur, allow sufficient additional buffer space to accommodate unexpected network performance issues.

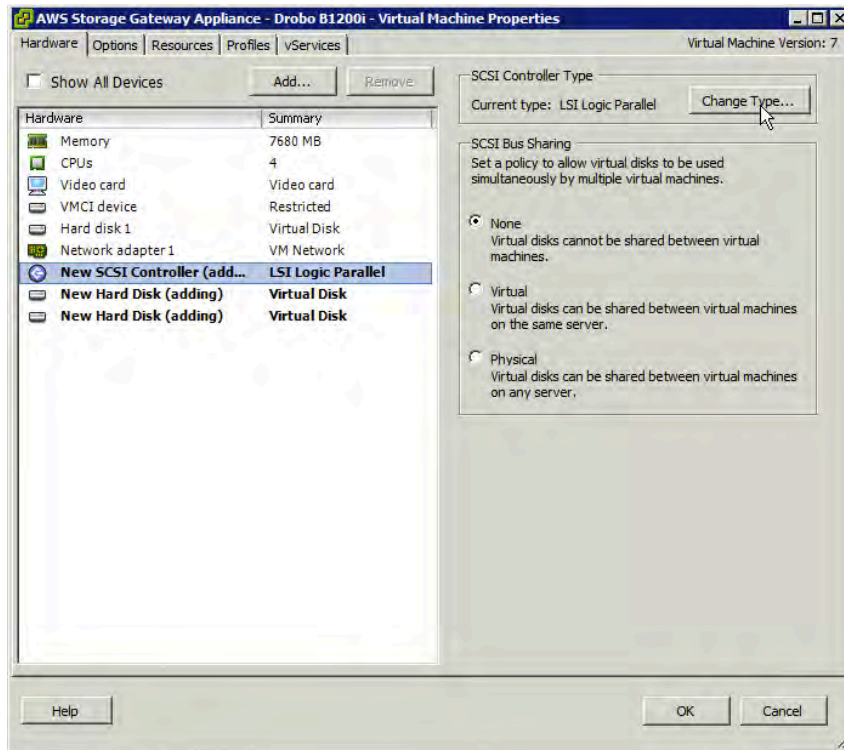
In this example, 20% is used for the daily changes of the data disk, or 2GB.

NOTE: If the size of the initial dataset to be synchronized is very large, the buffer disk may need to be as large as the dataset until the initial push can be completed over time. Once the initial synchronization is complete, the buffer volume can be resized to represent the average working state.



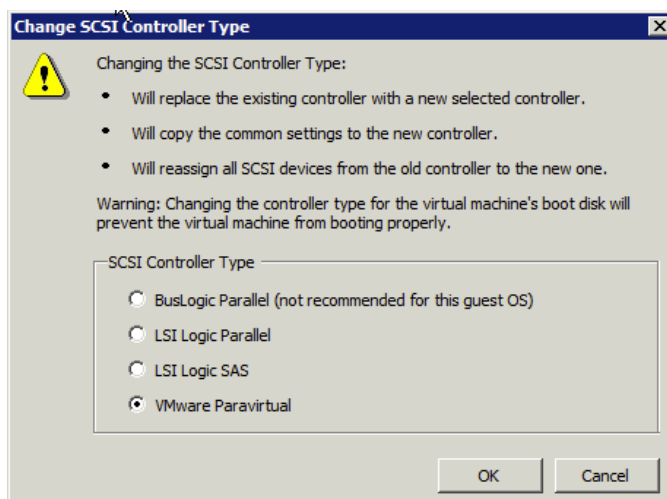
STEP 9

The newly created disks must be configured to use the VMware paravirtualized SCSI controller.



Select **New SCSI Controller** under devices and then select **Change Type** to modify the controller type.

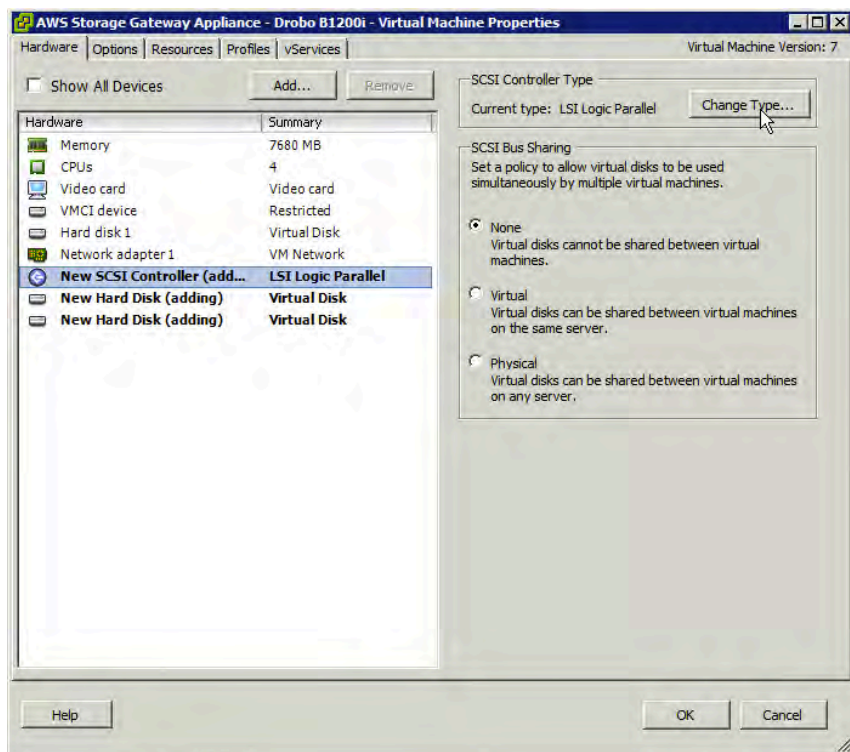
STEP 10



Under the **SCSI Controller Type**, select **VMware Paravirtual** from the drop-down menu and click **OK**.



STEP 11




Verify that the two additional hard drives and the new paravirtualized SCSI controller have been added properly. Click **OK** to save the hardware changes.



Activating the AWS Storage Appliance

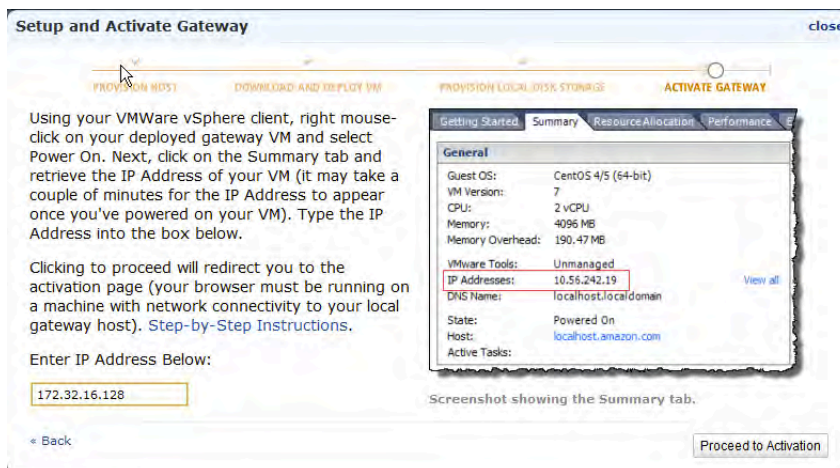
STEP 1



Power on the gateway virtual appliance and record the IP address assigned to the virtual machine.

It may take several minutes after powering on the appliance for the IP address to appear in the summary tab.

STEP 2



Enter the IP address of the gateway VM in the AWS wizard and click **Proceed to Activation**.

NOTE: Your browser must have network connectivity to the gateway VM. If the activation process fails, verify connectivity between your computer and the gateway VM.



STEP 3

The **AWS Region** is the location your data and snapshots will be stored. *This cannot be changed once the gateway is deployed.* Pricing varies between regions.

The **Gateway Name** is the name referenced in the AWS Console. This can be changed after deployment.

AWS Storage Gateway

Activating Your AWS Storage Gateway Virtual Machine (VM)

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

IP Address: 172.32.16.128

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

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

AWS Region:

Gateway Time Zone:

Gateway Name:

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

To complete the setup process, complete the activation form.

Navigation

Deploy a new Gateway

Region:

Gateways

- DroboAWSGateway
- Snapshots

DroboAWSGateway

Volumes | Gateway | Snapshot Schedules

You can create storage volumes to begin using the AWS Storage Gateway.

Once the activation process is complete, the new gateway is displayed in the AWS management console.

NOTE: Deployed gateways are displayed by region. If no gateways are displayed, verify that the correct region is selected from the drop-down menu.



STEP 4

The screenshot shows the 'DroboAWSGateway' configuration window with the 'Gateway' tab selected. The window has three buttons at the top: 'Shut Down', 'Delete Gateway', and 'Configure Local Working Storage'. The configuration details are as follows:

Name:	DroboAWSGateway (edit name)
Gateway ID:	sgw-84A540ED
IP Addresses:	172.32.16.128
Time Zone:	(GMT -8:00) Pacific Time (US & Canada) (edit time zone)
Rate Limit on Upload to AWS:	No Limit (edit limit)
Rate Limit on Download from AWS:	No Limit (edit limit)
Working Storage Used:	N/A
Maintenance Start Time:	Mondays 18:19 GMT-8:00 (edit time)

To define the disk to be used as a buffer when synchronizing data with AWS, select the **Gateway** tab and then select **Configure Local Working Storage**.

STEP 5

The 'Configure Local Working Storage' dialog box prompts the user to select disks for working storage. It includes a 'close' button in the top right corner and a link to 'Step-by-Step Instructions'.

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

Local Disks			
SCSI (0:0)	10.00 GiB	<input type="checkbox"/>	Use for Working Storage Space
SCSI (0:1)	2.00 GiB	<input checked="" type="checkbox"/>	Use for Working Storage Space

Buttons: Cancel, Next

Select the buffer storage from the available disks and click **Next**.

STEP 6

The 'Configure Local Working Storage' dialog box shows the confirmation step. It lists the selected disk and asks for confirmation to format it for working storage, warning that existing data will be lost.

The following disk(s) will be used for working storage on your gateway. Check the box below and click "Confirm" to proceed or "Back" to change your selection. Once confirmed, you can proceed with configuring your client iSCSI initiators to connect to your storage volumes.

SCSI (0:1) - (2.00 GiB capacity)

☒ Confirm that you want AWS Storage Gateway to format these disks to use for working storage. This format will result in the loss of existing data on these disks.

Buttons: Back, Confirm

Click **Confirm** to verify that you understand that the disk will be formatted for use by the gateway and all data will be lost.

(This disk was configured specifically for the appliance so no data will be lost.)

The AWS storage appliance setup is now complete and you are ready to provision storage.



Creating AWS Storage Volumes

The AWS Storage Gateway is now configured and ready to map storage to appliance servers.

STEP 1



Click **Create Volumes** to create a storage volume.

In the next step, the Disk drop-down menu lists all drives available to the storage gateway. In this example only the 10GB storage disk is listed.

STEP 2

close

Disk: ☐ Preserve existing data

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

Based on Snapshot ID:

Size: 10 GiB

Host IP:

Port:

Select the drive from the **Disk** drop-down menu.

Enter an **iSCSI target name**, which can be any combination of lowercase letters, numbers, periods, or hyphens.

Leave the **Based on Snapshot ID** field blank.

Click **Create Volume** to activate the volume.

The storage volume is now active and linked to the storage gateway.

Any data written to this device will be replicated to cloud-based AWS storage.

Volume ID	Capacity	Source Snapshot	Status
vol-FE484E86	10.00 GiB	none	AVAILABLE

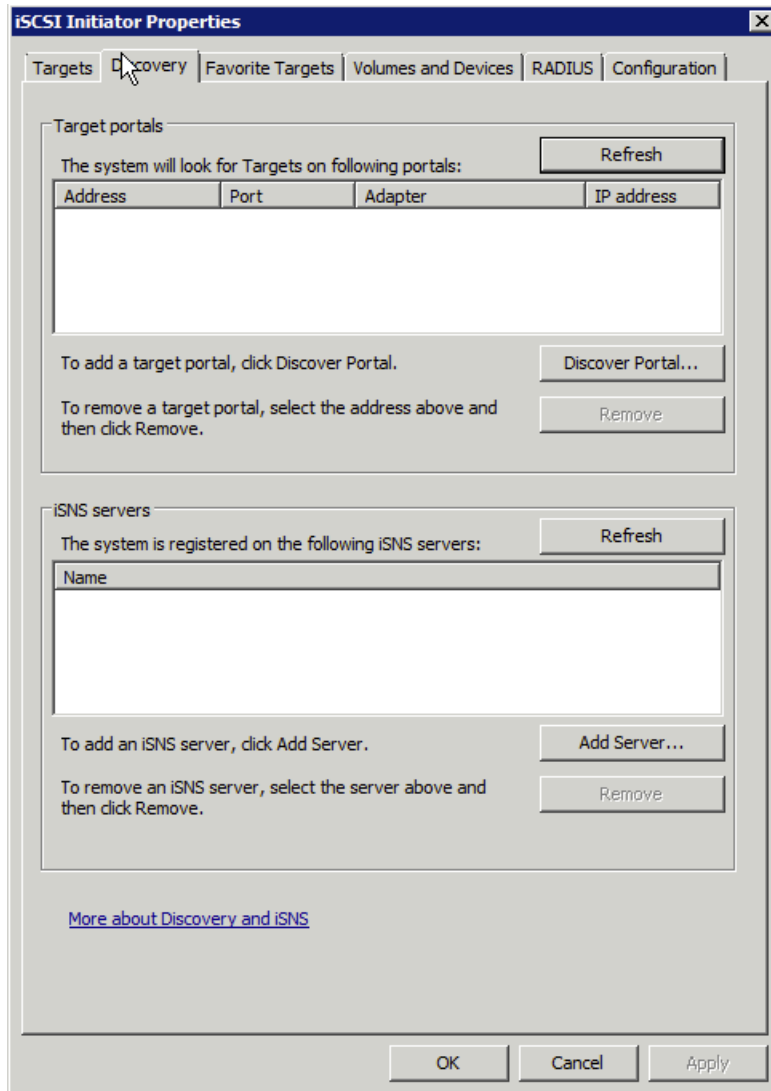


Mounting Storage Gateway Volumes

To access the AWS storage volumes on a Windows client, the iSCSI initiator must be configured to connect to the gateway VM.

On the start menu of the client, type “iscsicpl.exe” in the start menu to launch the iSCSI initiator. Alternatively, the iSCSI initiator can be found in the control panel.

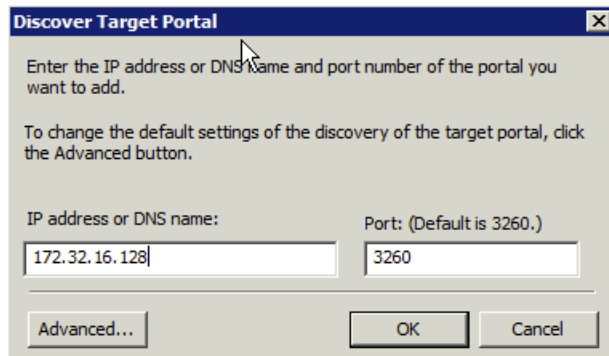
STEP 1



Once the iSCSI initiator loads, select the **Discovery** tab and click **Discover Portal** to add a target portal.



STEP 2



Discover Target Portal

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.

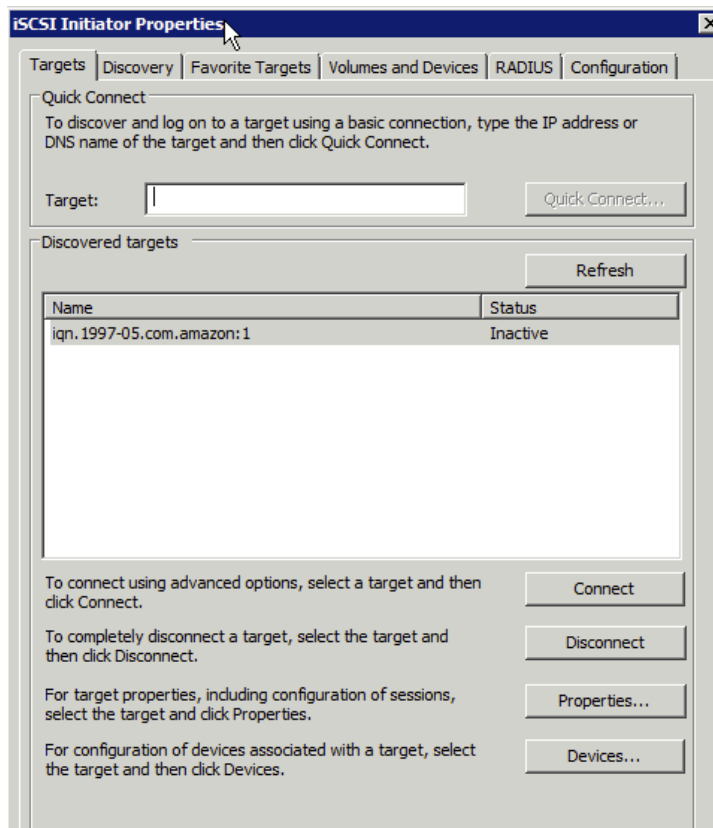
IP address or DNS name: Port: (Default is 3260.)

Advanced... OK Cancel

In the **IP address or DNS name** field, type the IP address of the gateway VM (the IP address displayed in the virtual machine summary tab of the vSphere management console).

Verify that the port is set 3260, and click **OK**.

STEP 3



iSCSI Initiator Properties

Targets | Discovery | Favorite Targets | Volumes and Devices | RADIUS | Configuration

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: Quick Connect...

Discovered targets

Refresh

Name	Status
iqn.1997-05.com.amazon:1	Inactive

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.

Connect Disconnect Properties... Devices...

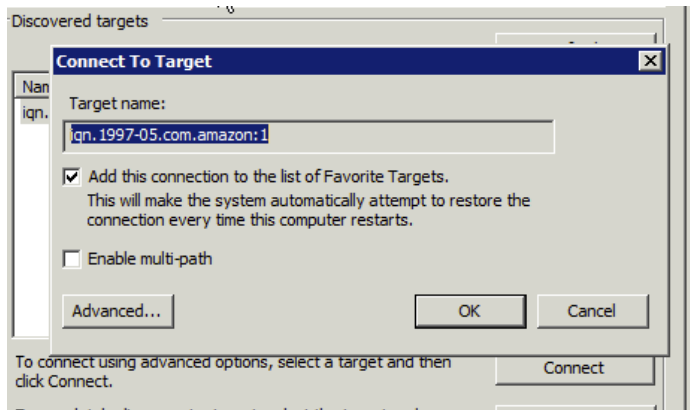
Select the **Targets** tab and look under **Discovered targets** for a list of available storage targets.

If no targets are listed, click the **Refresh** button. It may take several minutes before the initiator detects the newly created target.

Select the target (targets available for mapping are listed as Inactive), and click **Connect** to map the storage to the server.



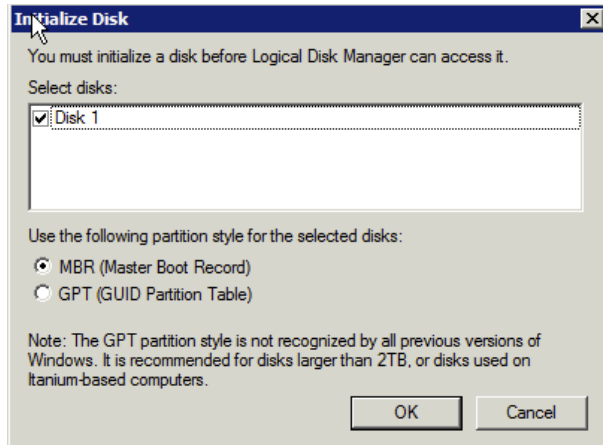
STEP 4



Check the box for **Add this connection** to the list of Favorite Targets and then click **OK** to map the storage.

The target should now be listed as Active and you can close the iSCSI initiator.

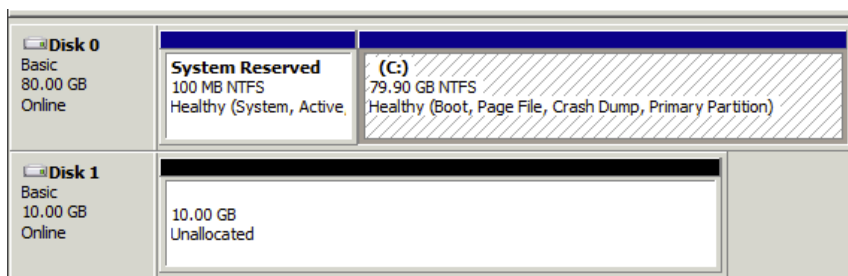
STEP 5



On the start menu type **diskmgmt.msc** to open the Disk Management console. Once the console displays, you will see the option to **Initialize Disk** for the newly created disk.


Select **MBR (Master Boot Record)** for the partition style and click **OK** to continue.

STEP 6



Verify that the new disk is displayed as **Online**. If the disk is offline, right-click on the drive and select **Online**.

STEP 7




To create a partition, right-click on the disk and select **New Simple Volume**.



CAUTION: The disk you are actively working with displays parallel diagonal lines. Verify that the correct disk has been selected before proceeding. Failure to do so can result in data loss.

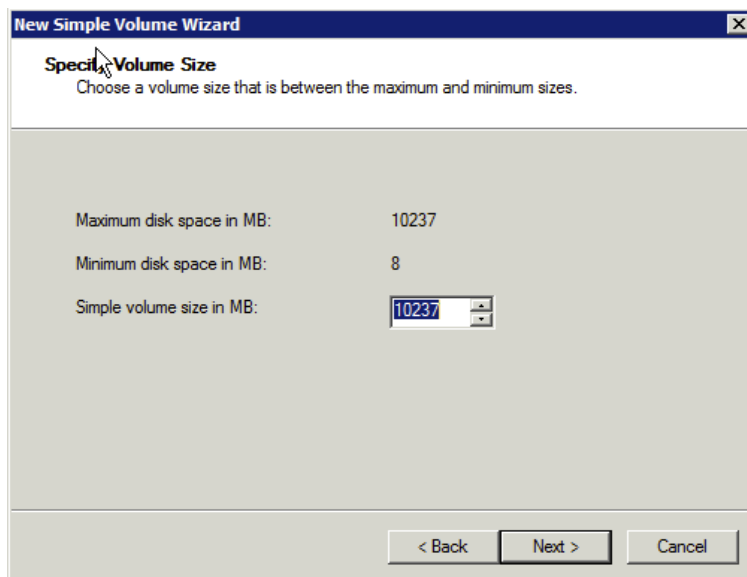
STEP 8



The volume wizard takes you through the process of partitioning and formatting the disk.

Click **Next** to continue.

STEP 9

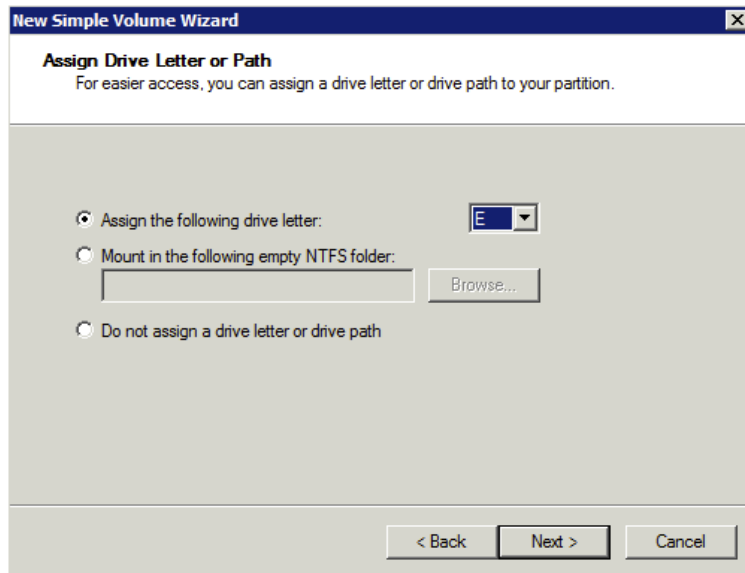


The new disk must be a single partition using all available space.

Leave all values set to the defaults and click **Next**.



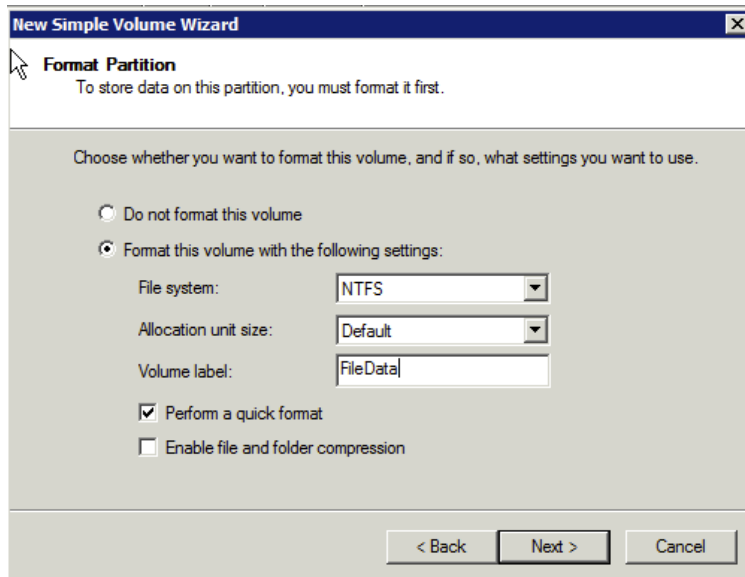
STEP 10



The new disk can be assigned to any available drive letter. Accept the default drive letter or change it if you wish.

Click **Next** to continue.

STEP 11

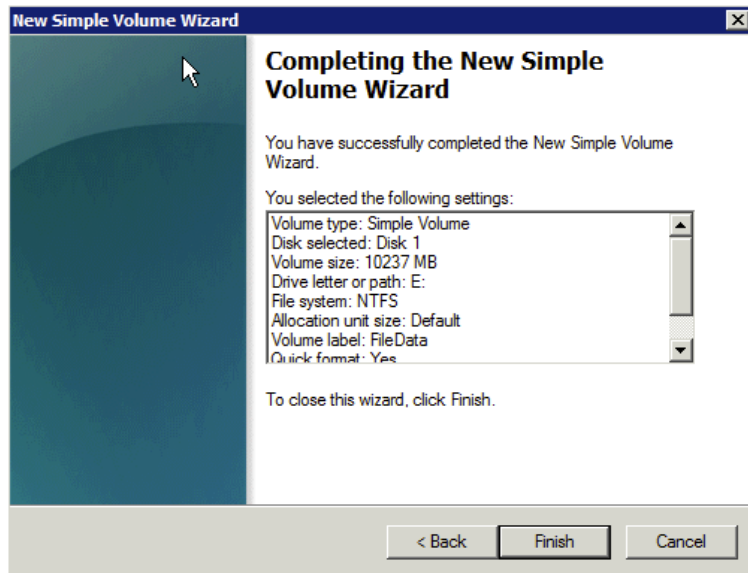


Select **Format this volume with the following settings** leaving:

- File system set to NTFS
- Allocation unit size to Default

Check the option to **Perform a quick format**.

NOTE: Performing a standard format of an iSCSI disk can take several hours to complete. There is no advantage to performing a standard (full) format in this scenario.



The drive is now mapped to the server and is available to store data.

Any data stored on this disk will be replicated to AWS cloud storage.

Restoring Storage Volume Snapshots

A snapshot is a point-in time-backup of the storage volume. When a storage gateway volume is created, a task is configured in the AWS Management Console to create a snapshot of the volume every 24 hours.

To modify the snapshot schedule, edit the task from the **Snapshot Schedules** tab of the storage gateway.

Automatic expiration of volume snapshots is not currently a feature of the AWS Management Console. It is your responsibility to manually delete snapshots that are no longer required to support data retention policies. Each snapshot is billed as the total provisioned size of the storage volume and additional storage fees can add up quickly if you don't monitor them.

NOTE: Snapshots can be scheduled to occur more frequently, but the maximum recurrence interval is 24 hours. Automatic snapshotting cannot be disabled at this time.



Before a snapshot can be restored a new disk must be added to the storage gateway virtual appliance.



STEP 1

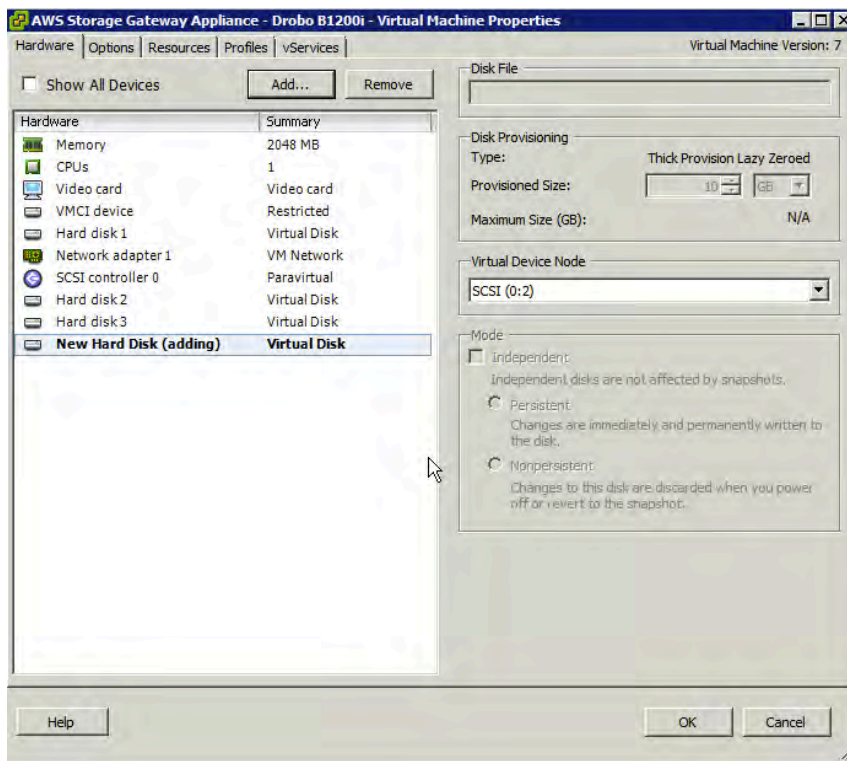


Select **Snapshots** from the Navigation pane to view all available snapshots. Select a snapshot with a Volume ID corresponding to the volume to be restored to the storage gateway.

Based on the snapshot schedule, there may be a large number of snapshots available. Before proceeding, record the **Snapshot ID** and **Capacity** as it will be needed in the following steps.

Before a snapshot can be restored, an additional hard disk must be added to the Storage Gateway virtual machine. Follow Steps 3 – 7 in the “Deploying the AWS Storage Gateway Appliance” for instructions on creating a disk.

STEP 2




In this example a third 10 gigabyte disk is added to the virtual machine to hold the restored snapshot. This disk must be the same size as the original data disk.

The additional hard disk can be added while the virtual machine is running and no reboot is required.



STEP 3



Select the storage gateway from the Navigation pane. In the Volumes tab, click **Create New Volume** to start the volume creation wizard.

STEP 4



Select an available hard drive from the Disk drop-down menu.

The iSCSI target name can be changed to better describe the storage volume or left as the default.

Type the Snapshot ID that was recorded in Step 1, and click **Create Volume**.

STEP 5



Return to the **Volumes** tab to view the status the the new volume.

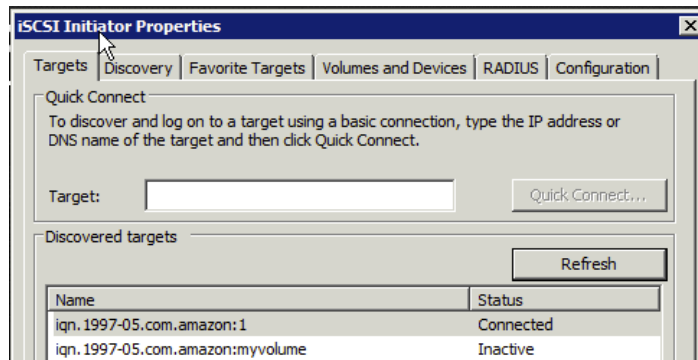
Once the volume creation process is complete, the restoration of the volume begins.

While the full restoration can take anywhere from a few minutes to several hours, data on the volume can accessed immediately.

NOTE: When your application requests data that has not yet been restored, the gateway will download the requested information from AWS and resume the restoration process.



STEP 6

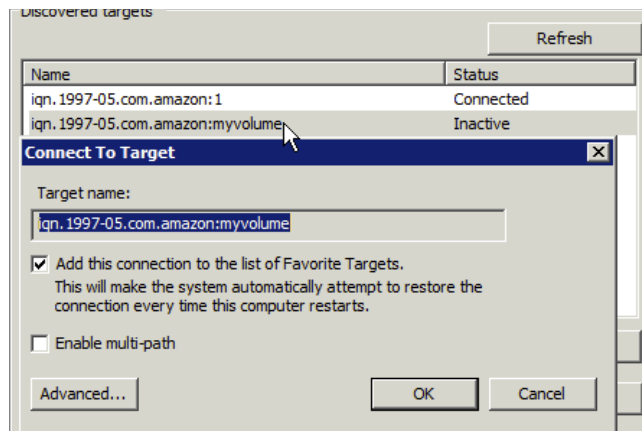


On the start menu of the server, type "iscsicpl.exe" in the start menu to launch the iSCSI initiator. You can also find the iSCSI initiator in the control panel.

In the Targets tab, click **Refresh** to update the list of available targets. The volume created from the snapshot is listed as Inactive.

Select the new volume, and click **Connect**.

STEP 7



Check the box for **Add this connection to the list of Favorite Targets**, and click **OK** to map the storage.

The target should now be listed as Active and you can close the iSCSI initiator.