



Hyper-V

Enterprise applications

NetApp
February 10, 2026

Table of Contents

Hyper-V	1
Deployment guidelines and storage best practices	1
Overview	1
NetApp storage and Windows Server environment	2
Provisioning in SAN environments	6
Provisioning in SMB environments	14
Hyper-V storage infrastructure on NetApp	17
Storage efficiency	26
Security	28
Deploy Nano server	28
Deploy Hyper-V cluster	31
Deploy Hyper-V Live Migration in a clustered environment	33
Deploy Hyper-V Live Migration outside a clustered environment	33
Deploy Hyper-V Storage Live Migration	35
Deploy Hyper-V Replica outside a clustered environment	35
Deploy Hyper-V replica in a clustered environment	37
Where to find additional information	38

Hyper-V

Deployment guidelines and storage best practices

Overview

Microsoft Windows Server is an enterprise-class operating system (OS) that covers networking, security, virtualization, private cloud, hybrid cloud, virtual desktop infrastructure, access protection, information protection, web services, application platform infrastructure, and much more.



This documentation replaces previously published technical reports *TR-4568: NetApp Deployment Guidelines and Storage Best Practices for Windows Server*

ONTAP runs on NetApp storage controllers. It is available in multiple formats.

- A unified architecture supporting file, object, and block protocols. This enables the storage controllers to act as both NAS and SAN devices as well as object stores
- An All SAN Array (ASA) that focuses only on block protocols and optimizes I/O resume times (IORT) by adding symmetric active-active multipathing for connect hosts
- A software defined unified architecture
 - ONTAP Select running on VMware vSphere or KVM
- First party offerings from hyperscale cloud providers
 - Amazon FSx for NetApp ONTAP
 - Azure NetApp Files
 - Google Cloud NetApp Volumes

ONTAP provides NetApp storage efficiency features such as NetApp snapshot technology, cloning, deduplication, thin provisioning, thin replication, compression, virtual storage tiering, and much more with enhanced performance and efficiency.

Together, Windows Server and ONTAP can operate in large environments and bring immense value to data center consolidation and private or hybrid cloud deployments. This combination also provides nondisruptive workloads efficiently and supports seamless scalability.

Intended audience

This document is intended for system and storage architects who design NetApp storage solutions for the Windows Server.

We make the following assumptions in this document:

- The reader has general knowledge of NetApp hardware and software solutions. See the [System Administration Guide for Cluster Administrators](#) for details.
- The reader has general knowledge of block-access protocols, such as iSCSI, FC and the file-access protocol SMB/CIFS. See the [Clustered Data ONTAP SAN management](#) for SAN-related information. See the [NAS management](#) for CIFS/SMB-related information.
- The reader has general knowledge of the Windows Server OS and Hyper-V.

For a complete, regularly updated matrix of tested and supported SAN and NAS configurations, see the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support Site. With the IMT, you can determine the exact product and feature versions that are supported for your specific environment. The NetApp IMT defines the product components and versions that are compatible with NetApp supported configurations. Specific results depend on each customer's installation in accordance with published specifications.

NetApp storage and Windows Server environment

As mentioned in the [Overview](#), NetApp storage controllers provide a truly unified architecture that supports file, block, and object protocols. This includes SMB/CIFS, NFS, NVMe/TCP, NVMe/FC, iSCSI, FC(FCP) and S3, and they create unified client and host access. The same storage controller can concurrently deliver block storage service in the form of SAN LUNs and file service as NFS and SMB/CIFS. ONTAP is also available as an All SAN Array (ASA) that optimizes host access through symmetric active-active multipathing with iSCSI and FCP, whereas the unified ONTAP systems use asymmetric active-active multipathing. In both modes, ONTAP uses ANA for NVMe over Fabrics (NVMe-oF) multipath management.

A NetApp storage controller running ONTAP can support the following workloads in a Windows Server environment:

- VMs hosted on continuously available SMB 3.0 shares
- VMs hosted on Cluster Shared Volume (CSV) LUNs running on iSCSI or FC
- SQL Server databases on SMB 3.0 shares
- SQL Server databases on NVMe-oF, iSCSI or FC
- Other application workloads

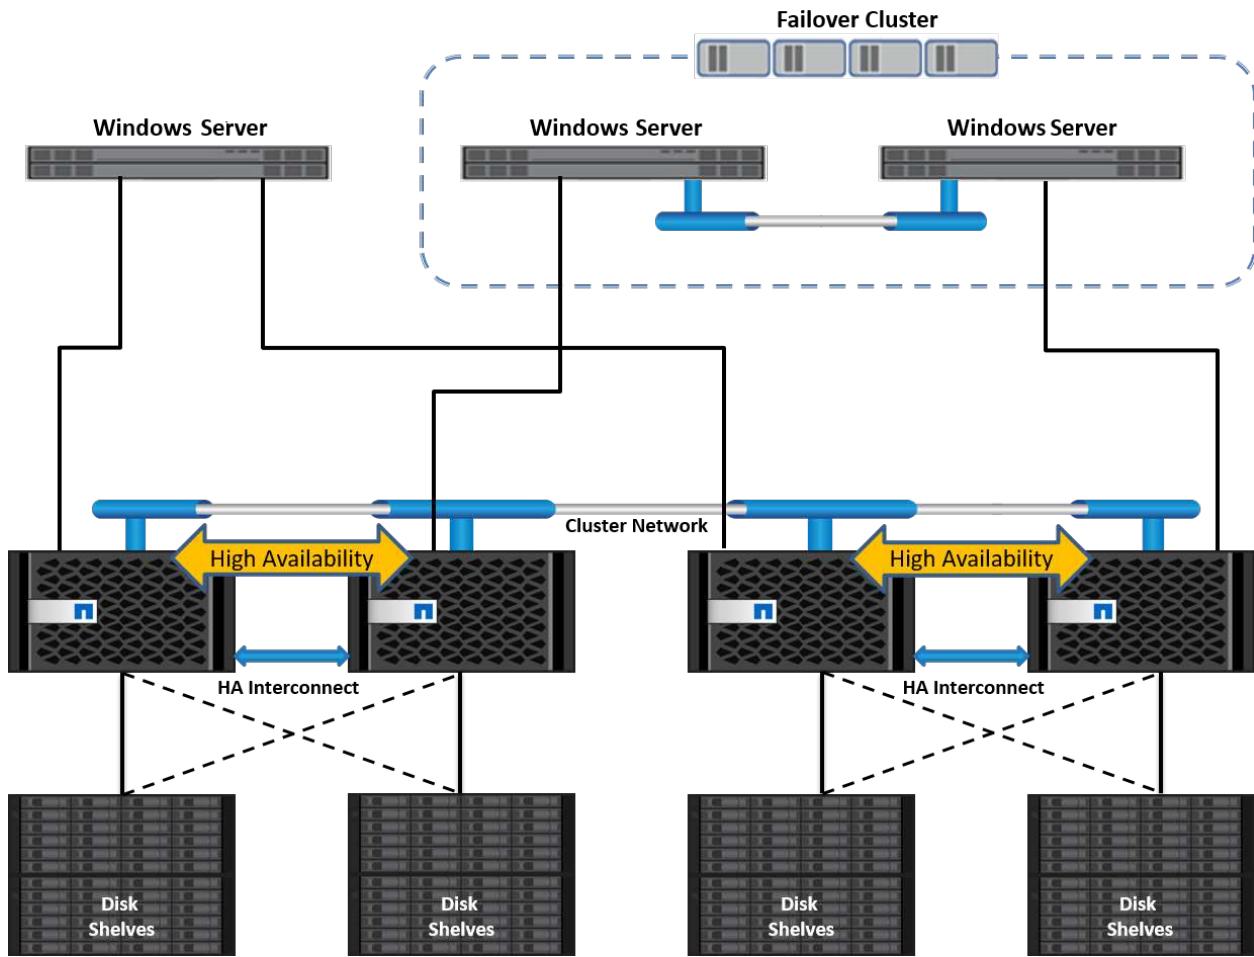
In addition, NetApp storage efficiency features such as deduplication, NetApp FlexClone® copies, NetApp snapshot technology, thin provisioning, compression, and storage tiering provide significant value for workloads running on Windows Server.

ONTAP data management

ONTAP is management software that runs on a NetApp storage controller. Referred to as a node, a NetApp storage controller is a hardware device with a processor, RAM, and NVRAM. The node can be connected to SATA, SAS, or SSD disk drives or a combination of those drives.

Multiple nodes are aggregated into a clustered system. The nodes in the cluster communicate with each other continuously to coordinate cluster activities. The nodes can also move data transparently from node to node by using redundant paths to a dedicated cluster interconnect consisting of two 10Gb Ethernet switches. The nodes in the cluster can take over one another to provide high availability during any failover scenarios. Clusters are administered on a whole-cluster rather than a per-node basis, and data is served from one or more storage virtual machines (SVMs). A cluster must have at least one SVM to serve data.

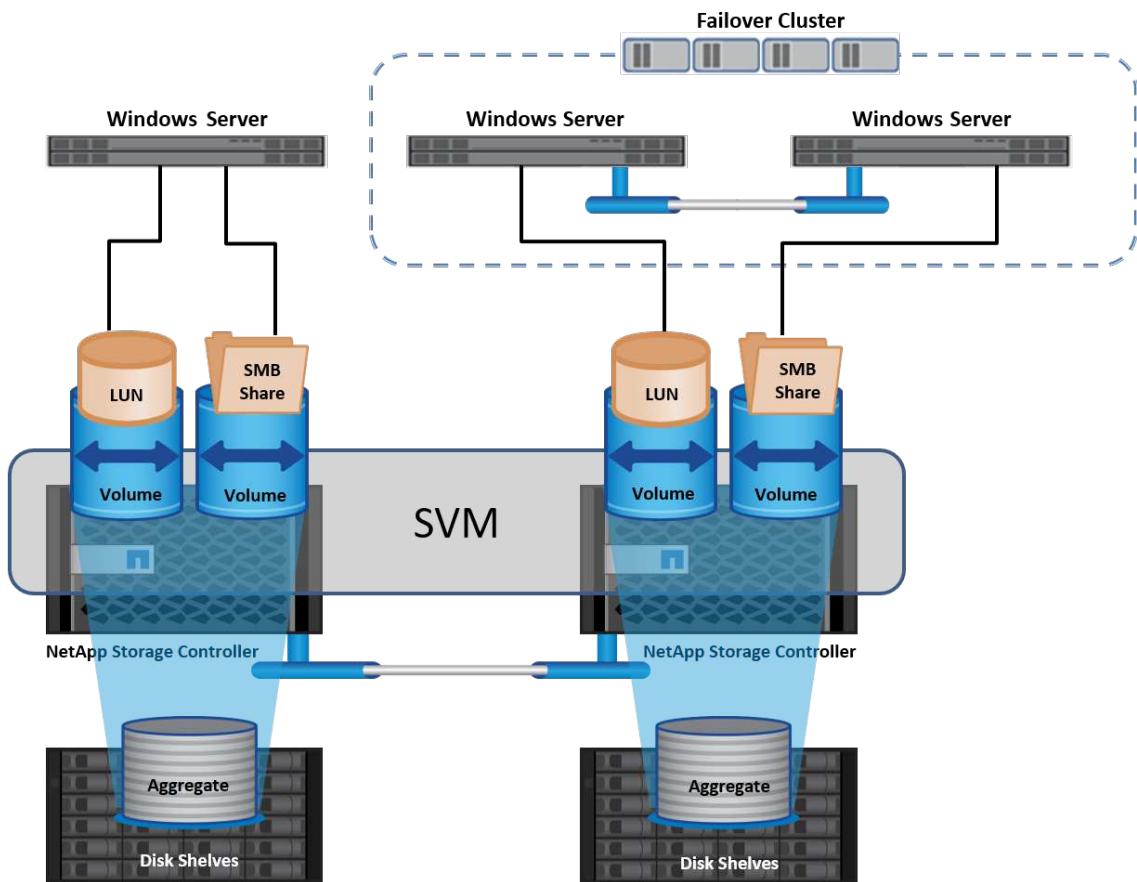
The basic unit of a cluster is the node, and nodes are added to the cluster as part of a high-availability (HA) pair. HA pairs enable high availability by communicating with each other over an HA interconnect (separate from the dedicated cluster interconnect) and by maintaining redundant connections to the HA pair's disks. Disks are not shared between HA pairs, although shelves might contain disks that belong to either member of an HA pair. The following figure depicts a NetApp storage deployment in a Windows Server environment.



Storage Virtual Machines

An ONTAP SVM is a logical storage server that provides data access to LUNs and/or a NAS namespace from one or more logical interfaces (LIFs). The SVM is thus the basic unit of storage segmentation that enables secure multitenancy in ONTAP. Each SVM is configured to own storage volumes provisioned from a physical aggregate and logical interfaces (LIFs) assigned either to a physical Ethernet network or to FC target ports.

Logical disks (LUNs) or CIFS shares are created inside an SVM's volumes and are mapped to Windows hosts and clusters to provide them with storage space, as shown in the following figure. SVMs are node independent and cluster based; they can use physical resources such as volumes or network ports anywhere in the cluster.



Provisioning NetApp storage for Windows Server

Storage can be provisioned to Windows Server in both SAN and NAS environments. In a SAN environment, the storage is provided as disks from LUNs on NetApp volume as block storage. In a NAS environment, the storage is provided as CIFS/SMB shares on NetApp volumes as file storage. These disks and shares can be applied in Windows Server as follows:

- Storage for Windows Server hosts for application workloads
- Storage for Nano Server and containers
- Storage for individual Hyper-V hosts to store VMs
- Shared storage for Hyper-V clusters in the form of CSVs to store VMs
- Storage for SQL Server databases

Managing NetApp storage

To connect, configure, and manage NetApp storage from Windows Server 2016, use one of the following methods:

- **Secure Shell (SSH).** Use any SSH client on Windows Server to run NetApp CLI commands.
- **System Manager.** This is NetApp's GUI-based manageability product.
- **NetApp PowerShell Toolkit.** This is the NetApp PowerShell Toolkit for automating and implementing custom scripts and workflows.

NetApp PowerShell Toolkit

NetApp PowerShell Toolkit (PSTK) is a PowerShell module that provides end-to-end automation and enables storage administration of ONTAP. The ONTAP module contains over 2,000 cmdlets and helps with the administration of NetApp AFF, FAS, commodity hardware, and cloud resources.

Things to remember

- NetApp does not support Windows Server storage spaces. Storage spaces are used only for JBOD (just a bunch of disks) and does not work with any type of RAID (direct-attached storage [DAS] or SAN).
- Clustered storage pools in Windows Server are not supported by ONTAP.
- NetApp supports the shared virtual hard disk format (VHDX) for guest clustering in Windows SAN environments.
- Windows Server does not support creating storage pools using iSCSI or FC LUNs.

Further reading

- For more information about the NetApp PowerShell Toolkit, visit the [NetApp Support Site](#).
- For information about NetApp PowerShell Toolkit best practices, see [TR-4475: NetApp PowerShell Toolkit Best Practices Guide](#).

Networking best practices

Ethernet networks can be broadly segregated into the following groups:

- A client network for the VMs
- One or more storage networks (iSCSI or SMB connecting to the storage systems)
- A cluster communication network (heartbeat and other communication between the nodes of the cluster)
- A management network (to monitor and troubleshoot the system)
- A migration network (for host live migration)
- VM replication (a Hyper-V Replica)

Best practices

- NetApp recommends having dedicated physical ports for each of the preceding functionalities for network isolation and performance.
- For each of the preceding network requirements (except for the storage requirements), multiple physical network ports can be aggregated to distribute load or provide fault tolerance.
- NetApp recommends having a dedicated virtual switch created on the Hyper-V host for guest storage connection within the VM.
- Make sure that the Hyper-V host and guest iSCSI data paths use different physical ports and virtual switches for secure isolation between the guest and the host.
- NetApp recommends avoiding NIC teaming for iSCSI NICs.
- NetApp recommends using ONTAP multipath input/output (MPIO) configured on the host for storage purposes..
- NetApp recommends using MPIO within a guest VM if using guest iSCSI initiators. MPIO usage must be avoided within the guest if you use pass-through disks. In this case, installing MPIO on the host should suffice.

- NetApp recommends not applying QoS policies to the virtual switch assigned for the storage network.
- NetApp recommends not using automatic private IP addressing (APIPA) on physical NICs because APIPA is nonroutable and not registered in the DNS.
- NetApp recommends turning on jumbo frames for CSV, iSCSI, and live migration networks to increase the throughput and reduce CPU cycles.
- NetApp recommends unchecking the option Allow Management Operating System to Share This Network Adapter for the Hyper-V virtual switch to create a dedicated network for the VMs.
- NetApp recommends creating redundant network paths (multiple switches) for live migration and the iSCSI network to provide resiliency and QoS.

Provisioning in SAN environments

ONTAP SVMs support the block protocols iSCSI and FC. When an SVM is created with block protocol iSCSI or FC, the SVM gets either an iSCSI Qualified Name (IQN) or an FC worldwide name (WWN), respectively. This identifier presents a SCSI target to hosts that access NetApp block storage.

Provisioning NetApp LUN on Windows Server

Prerequisites

Using NetApp storage in SAN environments in Windows Server has the following requirements:

- A NetApp cluster is configured with one or more NetApp storage controllers.
- The NetApp cluster or storage controllers have a valid iSCSI license.
- iSCSI and/or FC configured ports are available.
- FC zoning is performed on an FC switch for FC.
- At least one aggregate is created.
- An SVM should have one LIF per Ethernet network or Fibre Channel fabric on every storage controller that is going to serve data using iSCSI or Fibre Channel.

Deployment

1. Create a new SVM with block protocol iSCSI and/or FC enabled. A new SVM can be created with any of the following methods:
 - CLI commands on NetApp storage
 - ONTAP System Manager
 - NetApp PowerShell Toolkit
2. Configure the iSCSI and/or FC protocol.
3. Assign the SVM with LIFs on each cluster node.
4. Start the iSCSI and/or FC service on the SVM.
5. Create iSCSI and/or FC port sets using the SVM LIFs.

5. Create an iSCSI and/or FC initiator group for Windows using the port set created.
6. Add an initiator to the initiator group. The initiator is the IQN for iSCSI and WWPN for FC. They can be queried from Windows Server by running the PowerShell cmdlet Get-InitiatorPort.

```
# Get the IQN for iSCSI
Get-InitiatorPort | Where \{$_.ConnectionType -eq 'iSCSI'\} | Select-Object -Property NodeAddress
```

```
# Get the WWPN for FC
Get-InitiatorPort | Where \{$_.ConnectionType -eq 'Fibre Channel'\} | Select-Object -Property PortAddress
```

```
# While adding initiator to the initiator group in case of FC, make sure to provide the initiator(PortAddress) in the standard WWPN format
```

The IQN for iSCSI on Windows Server can also be checked in the configuration of the iSCSI initiator properties.

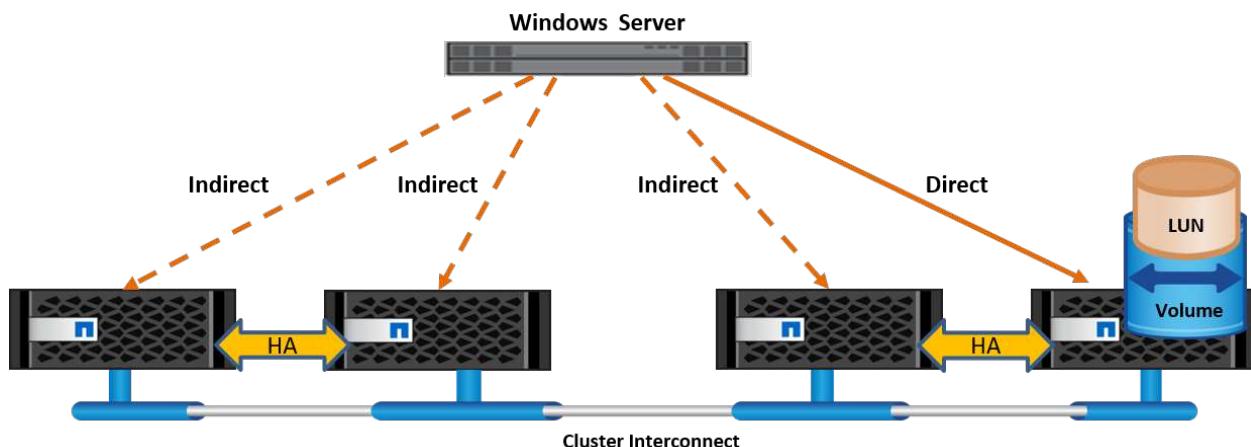
- Create a LUN using Create LUN wizard and associate it with the initiator group created.

Host integration

Windows Server uses Asymmetrical Logical Unit Access (ALUA) extension MPIO to determine direct and indirect paths to LUNs. Even though every LIF owned by an SVM accepts read/write requests for its LUNs, only one of the cluster nodes actually owns the disks backing that LUN at any given moment. This divides available paths to a LUN into two types, direct or indirect, as shown in the following figure.

A direct path for a LUN is a path on which an SVM's LIFs and the LUN being accessed reside on the same node. To go from a physical target port to disk, it is not necessary to traverse the cluster interconnect.

Indirect paths are data paths on which an SVM's LIFs and the LUN being accessed reside on different nodes. Data must traverse the cluster interconnect to go from a physical target port to disk.



MPIO

ONTAP provides highly available storage in which multiple paths from the storage controller to the Windows Server can exist. Multipathing is the ability to have multiple data paths from a server to a storage array. Multipathing protects against hardware failures (cable cuts, switch and host bus adapter [HBA] failure, and so on), and it can provide higher performance limits by using the aggregate performance of multiple connections. When one path or connection becomes unavailable, the multipathing software automatically shifts the load to one of the other available paths. The MPIO feature combines the multiple physical paths to the storage as a single logical path that is used for data access to provide storage resiliency and load balancing. To use this feature, the MPIO feature must be enabled on Windows Server.

Enable MPIO

To enable MPIO on Windows Server, complete the following steps:

1. Log in to Windows Server as a member of the administrator group.
7. Start Server Manager.
8. In the Manage section, click Add Roles and Features.
9. In the Select Features page, select Multipath I/O.

Configure MPIO

When using the iSCSI protocol, you must tell Windows Server to apply multipath support to iSCSI devices in the MPIO properties.

To configure MPIO on Windows Server, complete the following steps:

1. Log on to Windows Server as a member of the administrator group.
10. Start Server Manager.
11. In the Tools section, click MPIO.
12. In MPIO Properties on Discover Multi-Paths, select Add Support for iSCSI Devices and click Add. A prompt then asks you to restart the computer.
13. Reboot Windows Server to see the MPIO device listed in the MPIO Devices section of MPIO Properties.

Configure iSCSI

To detect iSCSI block storage on Windows Server, complete the following steps:

1. Log on to Windows Server as a member of the administrator group.
14. Start Server Manager.
15. In the Tools section, click iSCSI Initiator.
16. Under the Discovery tab, click Discover Portal.
17. Provide the IP address of the LIFs associated with the SVM created for the NetApp storage for SAN protocol. Click Advanced, configure the information in the General tab, and click OK.
18. The iSCSI initiator automatically detects the iSCSI target and lists it in the Targets tab.
19. Select the iSCSI target in Discovered Targets. Click Connect to open the Connect To Target window.
20. You must create multiple sessions from the Windows Server host to the target iSCSI LIFs on the NetApp

storage cluster. To do so, complete the following steps:

- a. In the Connect to Target window, select Enable MPIO and click Advanced.
- b. In Advanced Settings under the General tab, select the local adapter as the Microsoft iSCSI initiator and select the Initiator IP and Target Portal IP.
- c. You must also connect using the second path. Therefore, repeat step 5 through step 8, but this time select the Initiator IP and Target Portal IP for the second path.
- d. Select the iSCSI target in Discovered Targets on the iSCSI Properties main window and click Properties.
- e. The Properties window shows that multiple sessions have been detected. Select the session, click Devices, and then click the MPIO to configure the load balancing policy. All the paths configured for the device are displayed and all load balancing policies are supported. NetApp generally recommends round robin with subset, and this setting is the default for arrays with ALUA enabled. Round robin is the default for active-active arrays that do not support ALUA.

Detect block storage

To detect iSCSI or FC block storage on Windows Server, complete the following steps:

1. Click Computer Management in the Tools section of the Server Manager.
2. In Computer Management, click the Disk Management in Storage section and then click More Actions and Rescan Disks. Doing so displays the raw iSCSI LUNs.
3. Click the discovered LUN and make it online. Then select Initialize Disk using the MBR or GPT partition. Create a new simple volume by providing the volume size and drive letter and format it using FAT, FAT32, NTFS, or the Resilient File System (ReFS).

Best practices

- NetApp recommends enabling thin provisioning on the volumes hosting the LUNs.
- To avoid multipathing problems, NetApp recommends using either all 10Gb sessions or all 1Gb sessions to a given LUN.
- NetApp recommends that you confirm that ALUA is enabled on the storage system. ALUA is enabled by default on ONTAP.
- On the Windows Server host to where the NetApp LUN is mapped, enable iSCSI Service (TCP-In) for Inbound and iSCSI Service (TCP-Out) for Outbound in the firewall settings. These settings allow iSCSI traffic to pass to and from the Hyper-V host and NetApp controller.

Provisioning NetApp LUNs on Nano Server

Prerequisites

In addition to the prerequisites mentioned in the previous section, the storage role must be enabled from the Nano Server side. For example, Nano Server must be deployed using the -Storage option. To deploy Nano Server, see the section "[Deploy Nano Server](#)."

Deployment

To provision NetApp LUNs on a Nano Server, complete the following steps:

1. Connect to the Nano Server remotely using instructions in the section "[Connect to Nano Server](#)."
2. To configure iSCSI, run the following PowerShell cmdlets on the Nano Server:

```
# Start iSCSI service, if it is not already running
Start-Service msiscsi
```

```
# Create a new iSCSI target portal
New-IscsiTargetPortal -TargetPortalAddress <SVM LIF>
```

```
# View the available iSCSI targets and their node address
Get-IscsiTarget
```

```
# Connect to iSCSI target
Connect-IscsiTarget -NodeAddress <NodeAddress>
```

```
# NodeAddress is retrieved in above cmdlet Get-IscsiTarget
# OR
Get-IscsiTarget | Connect-IscsiTarget
```

```
# View the established iSCSI session
Get-IscsiSession
```

```
# Note the InitiatorNodeAddress retrieved in the above cmdlet Get-IscsiSession. This is the IQN for Nano server and this needs to be added in the Initiator group on NetApp Storage
```

```
# Rescan the disks
Update-HostStorageCache
```

3. Add an initiator to the initiator group.

```
Add the InitiatorNodeAddress retrieved from the cmdlet Get-IscsiSession to the Initiator Group on NetApp Controller
```

4. Configure MPIO.

```
# Enable MPIO Feature
Enable-WindowsOptionalFeature -Online -FeatureName MultipathIo
```

```
# Get the Network adapters and their IPs
Get-NetIPAddress -AddressFamily IPv4 -PrefixOrigin <Dhcp or Manual>
```

```
# Create one MPIO-enabled iSCSI connection per network adapter
Connect-IscsiTarget -NodeAddress <NodeAddress> -IsPersistent $True
-IsMultipathEnabled $True -InitiatorPortalAddress <IP Address of
ethernet adapter>
```

```
# NodeAddress is retrieved from the cmdlet Get-IscsiTarget
# IPs are retrieved in above cmdlet Get-NetIPAddress
```

```
# View the connections
Get-IscsiConnection
```

5. Detect block storage.

```
# Rescan disks
Update-HostStorageCache
```

```
# Get details of disks
Get-Disk
```

```
# Initialize disk
Initialize-Disk -Number <DiskNumber> -PartitionStyle <GPT or MBR>
```

```
# DiskNumber is retrieved in the above cmdlet Get-Disk
# Bring the disk online
Set-Disk -Number <DiskNumber> -IsOffline $false
```

```
# Create a volume with maximum size and default drive letter
New-Partition -DiskNumber <DiskNumber> -UseMaximumSize
-AssignDriveLetter
```

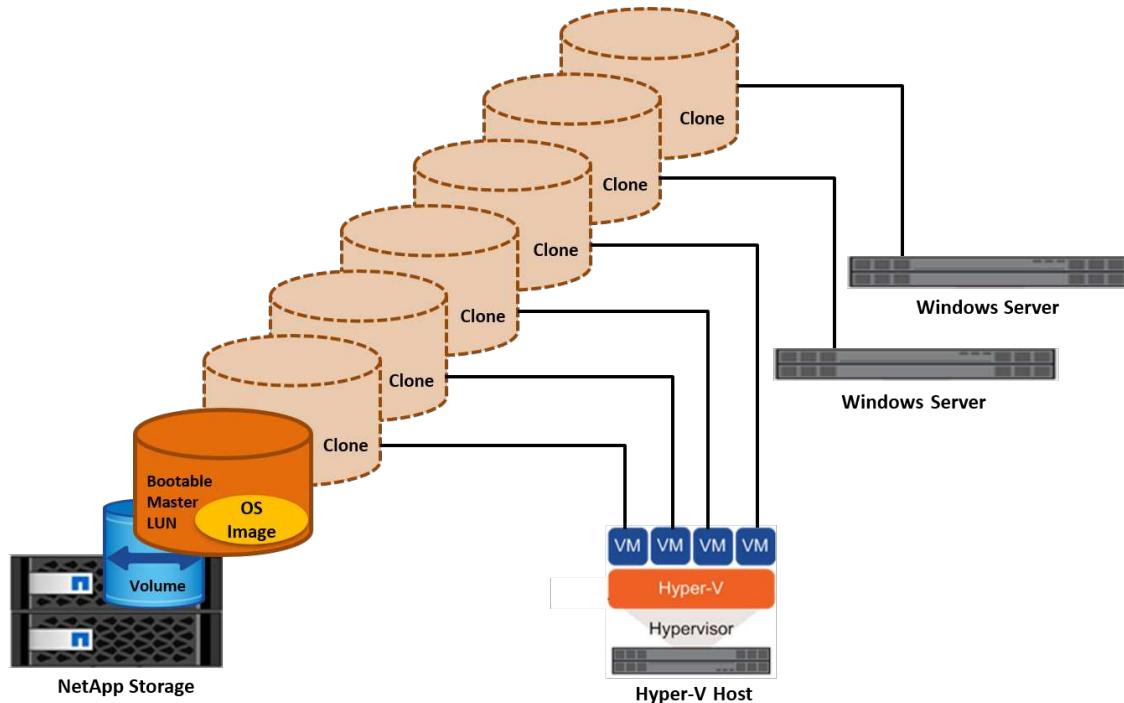
```
# To choose the size and drive letter use -Size and -DriveLetter
parameters
# Format the volume
Format-Volume -DriveLetter <DriveLetter> -FileSystem <FAT32 or NTFS or
REFS>
```

Boot from SAN

A physical host (server) or a Hyper-V VM can boot the Windows Server OS directly from a NetApp LUN instead of its internal hard disk. In the boot-from-SAN approach, the OS image to boot from resides on a NetApp LUN that is attached to a physical host or VM. For a physical host, the HBA of the physical host is configured to use the NetApp LUN for booting. For a VM, the NetApp LUN is attached as a pass-through disk for booting.

NetApp FlexClone approach

Using NetApp FlexClone technology, boot LUNs with an OS image can be cloned instantly and attached to the servers and VMs to rapidly provide clean OS images, as shown in the following figure.



Boot from SAN for physical host

Prerequisites

- The physical host (server) has a proper iSCSI or FC HBA.
- You have downloaded a suitable HBA device driver for the server supporting Windows Server.
- The server has a suitable CD/DVD drive or virtual media to insert the Windows Server ISO image and the HBA device driver has been downloaded.
- A NetApp iSCSI or FC LUN is provisioned on the NetApp storage controller.

Deployment

To configure booting from SAN for a physical host, complete the following steps:

1. Enable BootBIOS on the server HBA.
2. For iSCSI HBAs, configure the Initiator IP, iSCSI node name, and adapter boot mode in the boot BIOS settings.
3. When creating an initiator group for iSCSI and/or FC on a NetApp storage controller, add the server HBA initiator to the group. The HBA initiator of the server is the WWPN for the FC HBA or iSCSI node name for iSCSI HBA.
4. Create a LUN on the NetApp storage controller with a LUN ID of 0 and associate it with the initiator group created in the previous step. This LUN serves as a boot LUN.
5. Restrict the HBA to a single path to the boot LUN. Additional paths can be added after Windows Server is installed on the boot LUN to exploit the multipathing feature.
6. Use the HBA's BootBIOS utility to configure the LUN as a boot device.
7. Reboot the host and enter the host BIOS utility.
8. Configure the host BIOS to make the boot LUN the first device in the boot order.
9. From the Windows Server ISO, launch the installation setup.
10. When the installation asks, "Where Do You Want to Install Windows?", click Load Driver at the bottom of the installation screen to launch the Select Driver to Install page. Provide the path of the HBA device driver downloaded earlier and finish the installation of the driver.
11. Now the boot LUN created previously must be visible on the Windows installation page. Select the boot LUN for installation of Windows Server on the boot LUN and finish the installation.

Boot from SAN for virtual machine

To configure booting from SAN for a VM, complete the following steps:

Deployment

1. When creating an initiator group for iSCSI or FC on a NetApp storage controller, add the IQN for iSCSI or the WWN for FC of the Hyper-V server to the controller.
2. Create LUNs or LUN clones on the NetApp storage controller and associate them with the initiator group created in the previous step. These LUNs serve as boot LUNs for the VMs.
3. Detect the LUNs on the Hyper-V server, bring them online, and initialize them.
4. Bring the LUNs offline.
5. Create VMs with the option Attach a Virtual Hard Disk Later on the Connect Virtual Hard Disk page.
6. Add a LUN as a pass-through disk to a VM.

- a. Open the VM settings.
- b. Click IDE Controller 0, select Hard Drive, and click Add. Selecting IDE Controller 0 makes this disk the first boot device for the VM.
- c. Select Physical Hard Disk in the Hard Disk options and select a disk from the list as a pass-through disk. The disks are the LUNs configured in the previous steps.

7. Install Windows Server on the pass-through disk.

Best practices

- Make sure that the LUNs are offline. Otherwise, the disk cannot be added as a pass-through disk to a VM.
- When multiple LUNs exist, be sure to note the disk number of the LUN in disk management. Doing so is necessary because disks listed for the VM are listed with the disk number. Also, the selection of the disk as a pass-through disk for the VM is based on this disk number.
- NetApp recommends avoiding NIC teaming for iSCSI NICs.
- NetApp recommends using ONTAP MPIO configured on the host for storage purposes.

Provisioning in SMB environments

ONTAP provides resilient and high performance NAS storage for Hyper-V virtual machines using the SMB3 protocol.

When an SVM is created with the CIFS protocol, a CIFS server runs on top of the SVM that is part of the Windows Active Directory Domain. SMB shares can be used for a home directory and to host Hyper-V and SQL Server workloads. The following SMB 3.0 features are supported in ONTAP:

- Persistent handles (continuously available file shares)
- Witness protocol
- Clustered client failover
- Scale-out awareness
- ODX
- Remote VSS

Provisioning SMB shares on Windows Server

Prerequisites

Using NetApp storage in NAS environments in Windows Server has the following requirements:

- ONTAP cluster have a valid CIFS license.
- At least one aggregate is created.
- One data logical interface (LIF) is created and the data LIF must be configured for CIFS.
- A DNS-configured Windows Active Directory domain server and domain administrator credentials are present.
- Each node in the NetApp cluster is time synchronized with the Windows domain controller.

Active Directory Domain Controller

A NetApp storage controller can be joined to and operate within an Active Directory similar to a Windows Server. During the creation of the SVM, you can configure the DNS by providing the domain name and name server details. The SVM attempts to search for an Active Directory domain controller by querying the DNS for an Active Directory/Lightweight Directory Access Protocol (LDAP) server in a manner similar to Windows Server.

For the CIFS setup to work properly, the NetApp storage controllers must be time synchronized with the Windows domain controller. NetApp recommends having a time skew between the Windows domain controller and the NetApp storage controller of not more than five minutes. It is a best practice to configure the Network Time Protocol (NTP) server for the ONTAP cluster to synchronize with an external time source. To configure the Windows domain controller as the NTP server, run the following command on your ONTAP cluster:

```
$domainControllerIP = "<input IP Address of windows domain controller>"  
cluster::> system services ntp server create -server $domainControllerIP
```

Deployment

1. Create a new SVM with the NAS protocol CIFS enabled. A new SVM can be created with any of the following methods:
 - CLI commands on ONTAP
 - System Manager
 - The NetApp PowerShell Toolkit
2. Configure the CIFS protocol
 - a. Provide the CIFS server name.
 - b. Provide the Active Directory to which the CIFS server must be joined. You must have the domain administrator credentials to join the CIFS server to the Active Directory.
3. Assign the SVM with LIFs on each cluster node.
4. Start the CIFS service on the SVM.
5. Create a volume with the NTFS security style from the aggregate.
6. Create a qtree on the volume (optional).
7. Create shares that correspond to the volume or qtree directory so that they can be accessed from Windows Server. Select Enable Continuous Availability for Hyper-V during the creation of the share if the share is used for Hyper-V storage. Doing so enables high availability for file shares.
8. Edit the share created and modify the permissions as required for accessing the share. The permissions for the SMB share must be configured to grant access for the computer accounts of all the servers accessing this share.

Host integration

The NAS protocol CIFS is natively integrated into ONTAP. Therefore, Windows Server does not require any additional client software to access data on ONTAP. A NetApp storage controller appears on the network as a native file server and supports Microsoft Active Directory authentication.

To detect the CIFS share created previously with Windows Server, complete the following steps:

1. Log in to Windows Server as a member of the administrator group.
2. Go to run.exe and type the complete path of the CIFS share created to access the share.
3. To permanently map the share onto the Windows Server, right-click This PC, click Map Network Drive, and provide the path of the CIFS share.
4. Certain CIFS management tasks can be performed using Microsoft Management Console (MMC). Before performing these tasks, you must connect the MMC to the ONTAP storage using the MMC menu commands.
 - a. To open the MMC in Windows Server, click Computer Management in the Tools section of Server Manager.
 - b. Click More Actions and Connect to Another Computer, which opens the Select Computer dialog.
 - c. Enter the name of the CIFS server or the IP address of the SVM LIF to connect to the CIFS server.
 - d. Expand System Tools and Shared Folders to view and manage open files, sessions, and shares.

Best practices

- To confirm that there is no downtime when a volume is moved from one node to another or in the case of a node failure, NetApp recommends that you enable the continuous availability option on the file share.
- When provisioning VMs for a Hyper-V-over-SMB environment, NetApp recommends that you enable copy offload on the storage system. Doing so reduces the VMs' provisioning time.
- If the storage cluster hosts multiple SMB workloads such as SQL Server, Hyper-V, and CIFS servers, NetApp recommends hosting different SMB workloads on separate SVMs on separate aggregates. This configuration is beneficial because each of these workloads warrants unique storage networking and volume layouts.
- NetApp recommends connecting Hyper-V hosts and the ONTAP storage with a 10GB network if one is available. In the case of 1GB network connectivity, NetApp recommends creating an interface group consisting of multiple 1GB ports.
- When migrating VMs from one SMB 3.0 share to another, NetApp recommends enabling the CIFS copy offload functionality on the storage system so that migration is faster.

Things to remember

- When you provision volumes for SMB environments, the volumes must be created with the NTFS security style.
- Time settings on nodes in the cluster should be set up accordingly. Use the NTP if the NetApp CIFS server must participate in the Windows Active Directory domain.
- Persistent handles work only between nodes in an HA pair.
- The witness protocol works only between nodes in an HA pair.
- Continuously available file shares are supported only for Hyper-V and SQL Server workloads.
- The SMB multichannel is supported from ONTAP 9.4 onwards.
- RDMA is not supported.
- ReFS is not supported.

Provisioning SMB shares on Nano Server

Nano Server does not require additional client software to access data on the CIFS share on a NetApp storage controller.

To copy files from Nano Server to a CIFS share, run the following cmdlets on the remote server:

```
$ip = "<input IP Address of the Nano Server>"
```

```
# Create a New PS Session to the Nano Server
$session = New-PSSession -ComputerName $ip -Credential ~\Administrator
```

```
Copy-Item -FromSession $s -Path C:\Windows\Logs\DISM\dism.log -Destination
\\cifsshare
```

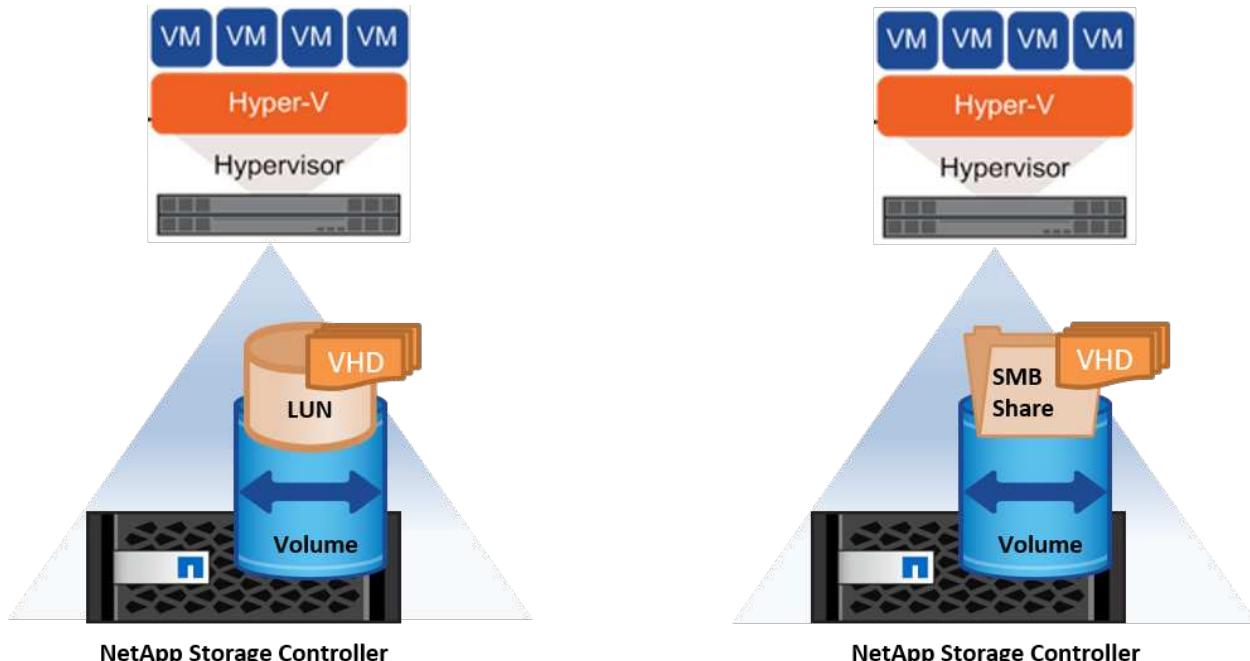
- `cifsshare` is the CIFS share on the NetApp storage controller.
- To copy files to Nano Server, run the following cmdlet:

```
Copy-Item -ToSession $s -Path \\cifsshare\<file> -Destination C:\
```

To copy the entire contents of a folder, specify the folder name and use the `-Recurse` parameter at the end of the cmdlet.

Hyper-V storage infrastructure on NetApp

A Hyper-V storage infrastructure can be hosted on ONTAP storage systems. Storage for Hyper-V to store the VM files and its disks can be provided using NetApp LUNs or NetApp CIFS shares, as shown in the following figure.



Hyper-V Storage on NetApp LUNs

- Provision a NetApp LUN on the Hyper-V server machine. For more information, see the section "[Provisioning in SAN Environments](#)."
- Open Hyper-V Manager from the Tools section of Server Manager.
- Select the Hyper-V server and click Hyper-V Settings.
- Specify the default folder to store the VM and its disk as the LUN. Doing so sets the default path as the LUN for the Hyper-V storage. If you want to specify the path explicitly for a VM, then you can do so during VM creation.

Hyper-V Storage on NetApp CIFS

Before beginning the steps listed in this section, review the section "[Provisioning in SMB Environments](#)." To configure Hyper-V storage on the NetApp CIFS share, complete the following steps:

1. Open Hyper-V Manager from the Tools section of Server Manager.
2. Select the Hyper-V server and click Hyper-V Settings.
3. Specify the default folder to store the VM and its disk as the CIFS share. Doing so sets the default path as the CIFS share for the Hyper-V storage. If you want to specify the path explicitly for a VM, then you can do so during VM creation.

Each VM in Hyper-V can in turn be provided with the NetApp LUNs and CIFS shares that were provided to the physical host. This procedure is the same as for any physical host. The following methods can be used to provision storage to a VM:

- Adding a storage LUN by using the FC initiator within the VM
- Adding a storage LUN by using the iSCSI initiator within the VM
- Adding a pass-through physical disk to a VM
- Adding VHD/VHDX to a VM from the host

Best practices

- When a VM and its data are stored on NetApp storage, NetApp recommends running NetApp deduplication at the volume level at regular intervals. This practice results in significant space savings when identical VMs are hosted on a CSV or SMB share. Deduplication runs on the storage controller and it does not affect the host system and VM performance.
- When using iSCSI LUNs for Hyper-V, make sure to enable 'iSCSI Service (TCP-In) for Inbound' and 'iSCSI Service (TCP-Out) for Outbound' in the firewall settings on the Hyper-V host. Doing so allows iSCSI traffic to pass to and from the Hyper-V host and the NetApp controller.
- NetApp recommends unchecking the option Allow Management Operating System to Share This Network Adapter for the Hyper-V virtual switch. Doing so creates a dedicated network for the VMs.

Things to remember

- Provisioning a VM by using virtual Fibre Channel requires an N_Port ID Virtualization enabled FC HBA. A maximum of four FC ports is supported.
- If the host system is configured with multiple FC ports and presented to the VM, then MPIO must be installed in the VM to enable multipathing.
- Pass-through disks cannot be provisioned to the host if MPIO is being used on that host, because pass-

through disks do not support MPIO.

- Disk used for VHD/VHDx files should use 64K formatting for allocation.

Further reading

- For information about FC HBAs, see the [Interoperability Matrix Tool \(IMT\)](#).
- For more information about virtual Fibre Channel, see the [Microsoft Hyper-V Virtual Fibre Channel Overview](#) page.

Offloaded data transfer

Microsoft ODX, also known as copy offload, enables direct data transfers within a storage device or between compatible storage devices without transferring the data through the host computer. ONTAP supports the ODX feature for both CIFS and SAN protocols. ODX can potentially improve performance if copies are within same volume, reduce utilization of CPU and memory on the client, and reduce network I/O bandwidth utilization.

With ODX, it is faster and efficient to copy files within the SMB shares, within the LUNs, and between the SMB shares and LUNs if it's in same volume. This approach is more helpful in a scenario for which multiple copies of the golden image of an OS (VHD/VHDx) are required within same volume. Several copies of the same golden image can be made in significantly less time if copies are within same volume. ODX is also applied in Hyper-V storage live migration for moving VM storage.

If copy is across volumes, there may not be significant performance gains compared to host-based copies.

To enable the ODX feature on CIFS, run the following CLI commands on the NetApp storage controller:

1. Enable ODX for CIFS.

```
#set the privilege level to diagnostic  
cluster::> set -privilege diagnostic
```

```
#enable the odx feature  
cluster::> vserver cifs options modify -vserver <vserver_name> -copy  
-offload-enabled true
```

```
#return to admin privilege level  
cluster::> set privilege admin
```

2. To enable the ODX feature on SAN, run the following CLI commands on the NetApp storage controller:

```
#set the privilege level to diagnostic  
cluster::> set -privilege diagnostic
```

```
#enable the odx feature  
cluster::> copy-offload modify -vserver <vserver_name> -scsi enabled
```

```
#return to admin privilege level  
cluster::> set privilege admin
```

Things to remember

- For CIFS, ODX is available only when both the client and the storage server support SMB 3.0 and the ODX feature.
- For SAN environments, ODX is available only when both the client and the storage server support the ODX feature.

Further reading

For information about ODX, see [Improving Microsoft Remote Copy Performance](#) and [Microsoft Offloaded Data Transfers](#).

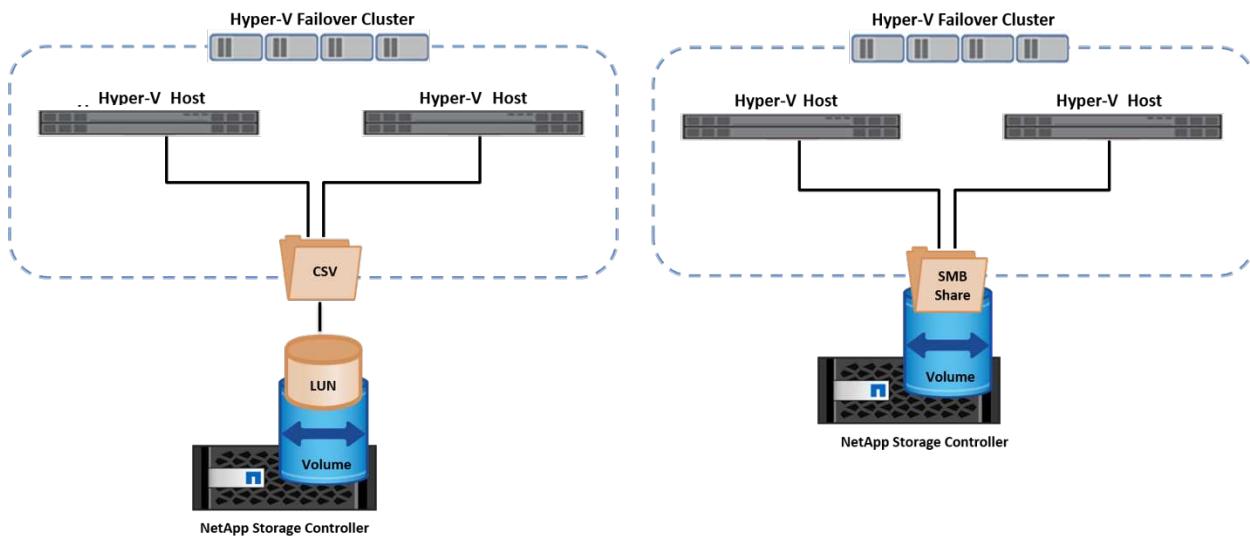
Hyper-V clustering: High availability and scalability for virtual machines

Failover clusters provide high availability and scalability to Hyper-V servers. A failover cluster is a group of independent Hyper-V servers that work together to increase availability and scalability for the VMs.

Hyper-V clustered servers (called nodes) are connected by the physical network and by cluster software. These nodes use shared storage to store the VM files, which include configuration, virtual hard disk (VHD) files, and snapshots. The shared storage can be a NetApp SMB/CIFS share or a CSV on top of a NetApp LUN, as shown below. This shared storage provides a consistent and distributed namespace that can be accessed simultaneously by all the nodes in the cluster. Therefore, if one node fails in the cluster, the other node provides service by a process called failover. Failover clusters can be managed by using the Failover Cluster Manager snap-in and the failover clustering Windows PowerShell cmdlets.

Cluster Shared Volumes

CSVs enable multiple nodes in a failover cluster to simultaneously have read/write access to the same NetApp LUN that is provisioned as an NTFS or ReFS volume. With CSVs, clustered roles can fail over quickly from one node to another without requiring a change in drive ownership or dismounting and remounting a volume. CSVs also simplify the management of a potentially large number of LUNs in a failover cluster. CSVs provide a general-purpose clustered file system that is layered above NTFS or ReFS.



Best practices

- NetApp recommends turning off cluster communication on the iSCSI network to prevent internal cluster communication and CSV traffic from flowing over the same network.

- NetApp recommends having redundant network paths (multiple switches) to provide resiliency and QoS.

Things to remember

- Disks used for CSV must be partitioned with NTFS or ReFS. Disks formatted with FAT or FAT32 cannot be used for a CSV.
- Disks used for CSVs should use 64K formatting for allocation.

Further reading

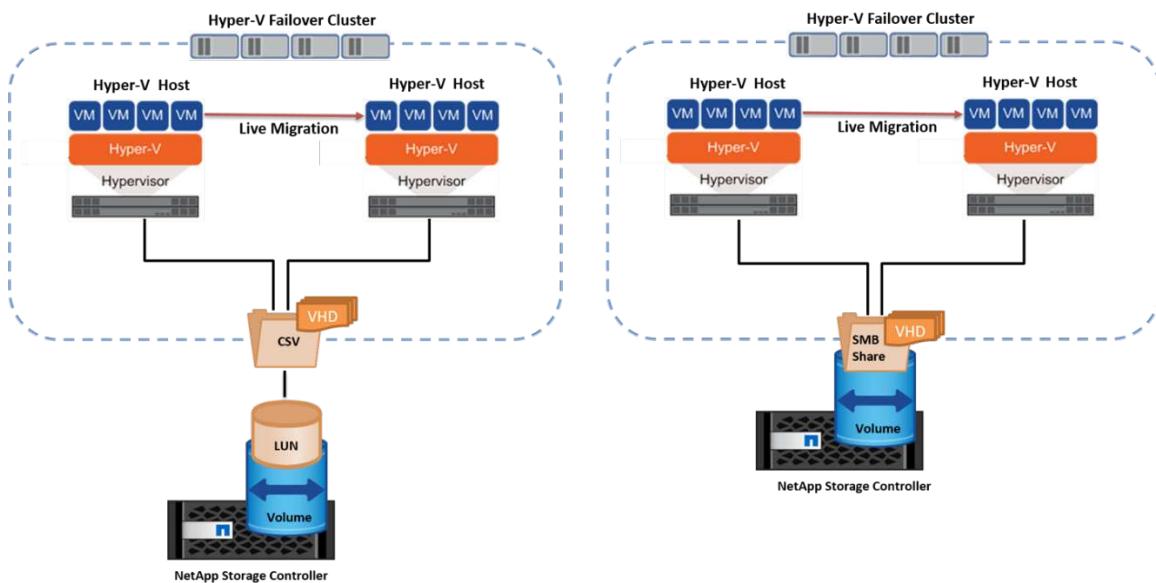
For information about deploying a Hyper-V cluster, see Appendix B: [Deploy Hyper-V Cluster](#).

Hyper-V Live Migration: Migration of VMs

It is sometimes necessary during the lifetime of VMs to move them to a different host on the Windows cluster. Doing so might be required if the host is running out of system resources or if the host is required to reboot for maintenance reasons. Similarly, it might be necessary to move a VM to a different LUN or SMB share. This might be required if the present LUN or share is running out of space or yielding lower than expected performance. Hyper-V live migration moves running VMs from one physical Hyper-V server to another with no effect on VM availability to users. You can live migrate VMs between Hyper-V servers that are part of a failover cluster or between independent Hyper-V servers that are not part of any cluster.

Live Migration in a clustered environment

VMs can be moved seamlessly between the nodes of a cluster. VM migration is instantaneous because all the nodes in the cluster share the same storage and have access to the VM and its disk. The following figure depicts live migration in a clustered environment.



Best practice

- Have a dedicated port for live migration traffic.
- Have a dedicated host live migration network to avoid network-related issues during migration.

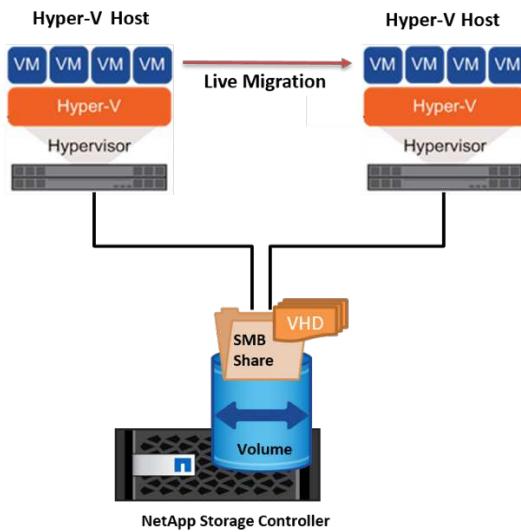
Further reading

For information about deploying live migration in a clustered environment, see [Appendix C: Deploy Hyper-V Live Migration in a Clustered Environment](#).

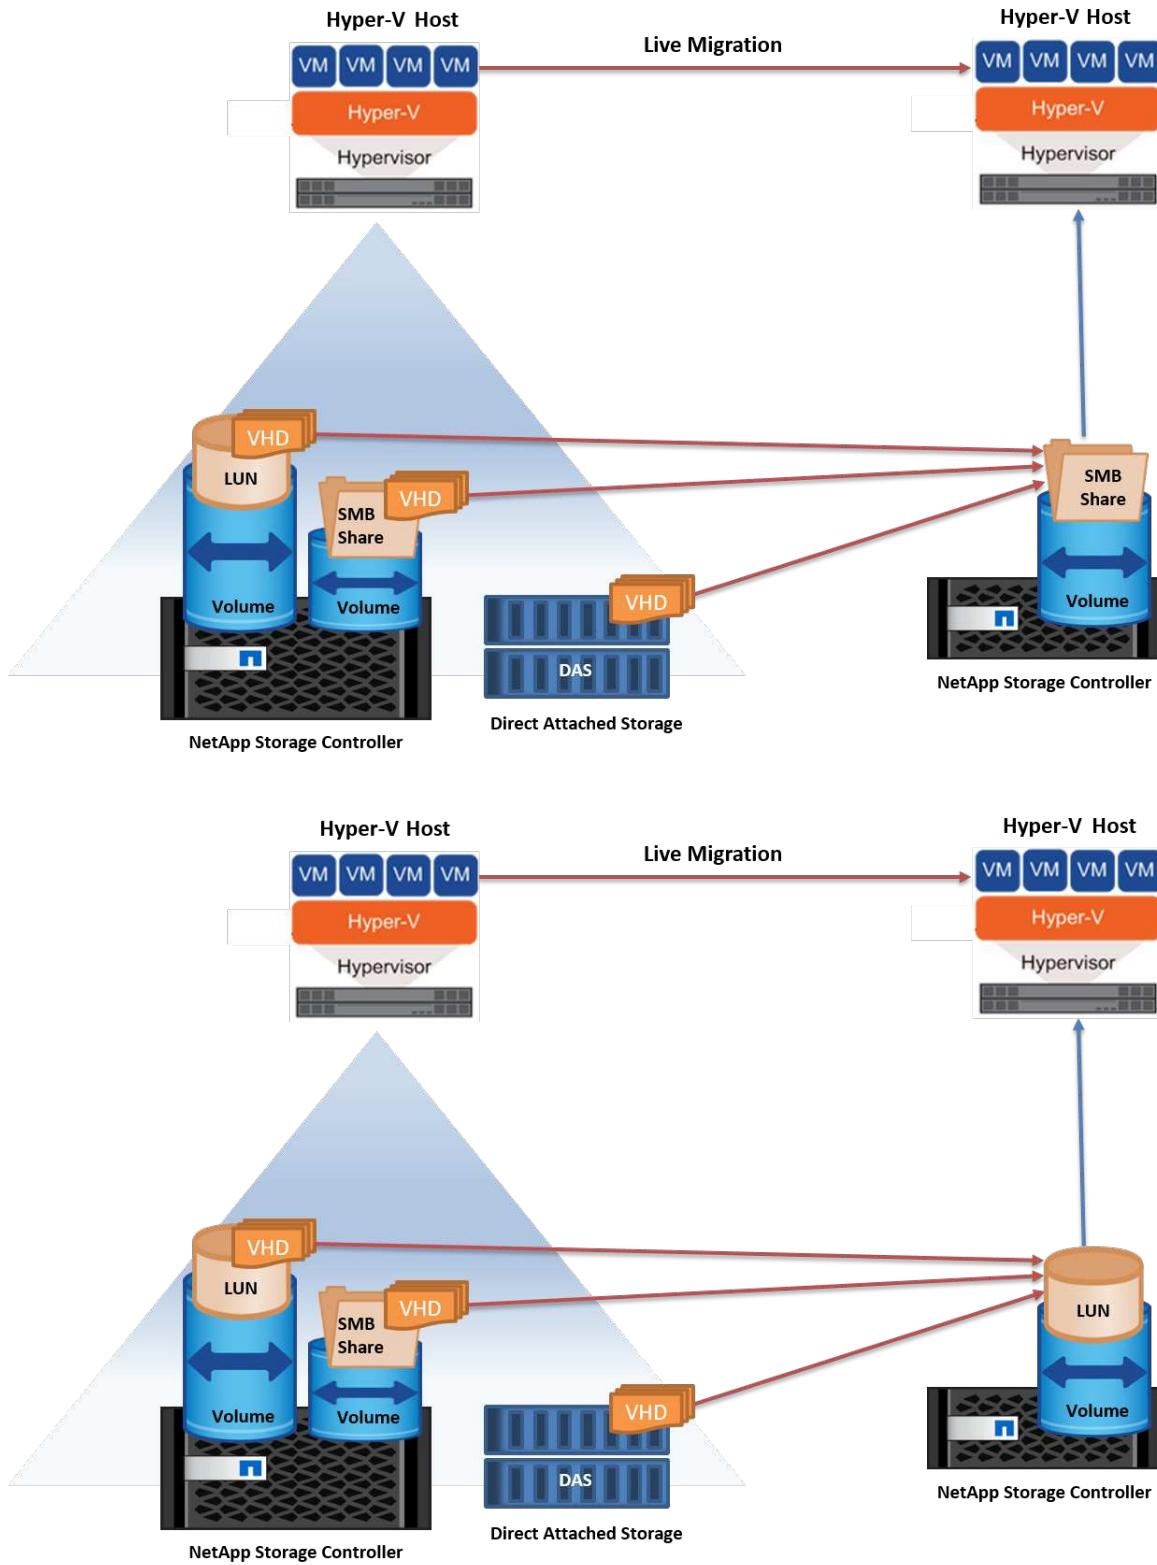
Live Migration outside a clustered environment

You can live migrate a VM between two nonclustered, independent Hyper-V servers. This process can use either shared or shared nothing live migration.

- In shared live migration, the VM is stored on an SMB share. Therefore, when you live migrate a VM, the VM's storage remains on the central SMB share for instant access by the other node, as shown below.



- In shared nothing live migration, each Hyper-V server has its own local storage (it can be an SMB share, a LUN, or DAS), and the VM's storage is local to its Hyper-V server. When a VM is live migrated, the VM's storage is mirrored to the destination server over the client network and then the VM is migrated. The VM stored on DAS, a LUN, or an SMB/CIFS share can be moved to an SMB/CIFS share on the other Hyper-V server, as shown in the following figure. It can also be moved to a LUN, as shown in the second figure.



Further reading

For information about deploying live migration outside a clustered environment, see [Appendix D: Deploy Hyper-V Live Migration Outside of a Clustered Environment](#).

Hyper-V Storage Live Migration

During the lifetime of a VM, you might need to move the VM storage (VHD/VHDX) to a different LUN or SMB share. This might be required if the present LUN or share is running out of space or yielding lower than expected performance.

The LUN or the share that currently hosts the VM can run out of space, be repurposed, or provide reduced performance. Under these circumstances, the VM can be moved without downtime to another LUN or share on a different volume, aggregate, or cluster. This process is faster if the storage system has copy-offload capabilities. NetApp storage systems are copy-offload enabled by default for CIFS and SAN environments.

The ODX feature performs full-file or sub-file copies between two directories residing on remote servers. A copy is created by copying data between the servers (or the same server if both the source and the destination files are on the same server). The copy is created without the client reading the data from the source or writing to the destination. This process reduces processor and memory use for the client or server and minimizes network I/O bandwidth. The copy is faster if its within same volume. If copy is across volumes, there may not be significant performance gains compared to host-based copies. Before proceeding with a copy operation on the host, confirm that the copy offload settings are configured on the storage system.

When VM storage live migration is initiated from a host, the source and the destination are identified, and the copy activity is offloaded to the storage system. Because the activity is performed by the storage system, there is negligible use of the host CPU, memory, or network.

NetApp storage controllers support the following different ODX scenarios:

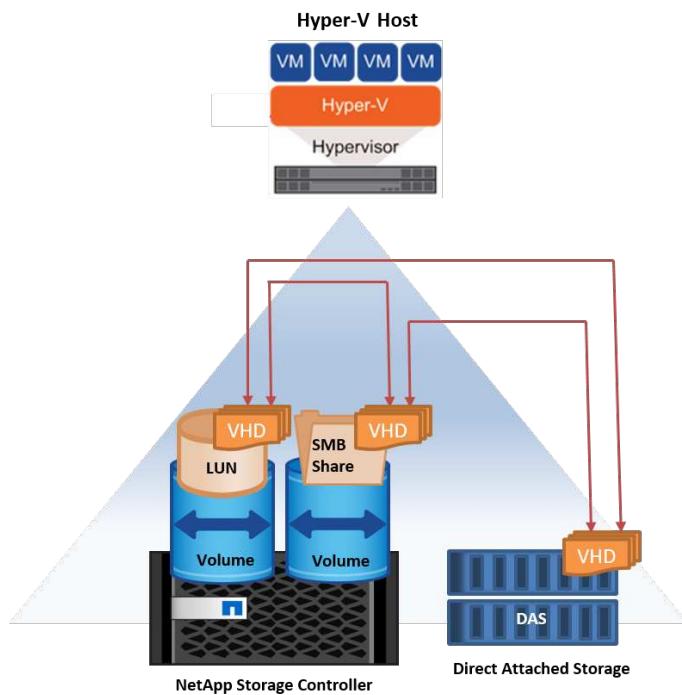
- **IntraSVM.** The data is owned by the same SVM:
- **Intervolume, intranode.** The source and destination files or LUNs reside within the same volume. The copy is performed with FlexClone file technology, which provides additional remote copy performance benefits.
- **Intervolume, intranode.** The source and destination files or LUNs are on different volumes that are on the same node.
- **Intervolume, internodes.** The source and destination files or LUNs are on different volumes that are located on different nodes.
- **InterSVM.** The data is owned by different SVMs.
- **Intervolume, intranode.** The source and destination files or LUNs are on different volumes that are on the same node.
- **Intervolume, internodes.** The source and destination files or LUNs are on different volumes that are on different nodes.
- **Intercluster.** Beginning with ONTAP 9.0, ODX is also supported for intercluster LUN transfers in SAN environments. Intercluster ODX is supported for SAN protocols only, not for SMB.

After the migration is complete, the backup and replication policies must be reconfigured to reflect the new volume holding the VMs. Any previous backups that were taken cannot be used.

VM storage (VHD/VHDX) can be migrated between the following storage types:

- DAS and the SMB share
- DAS and LUN
- An SMB share and a LUN
- Between LUNs

- Between SMB shares

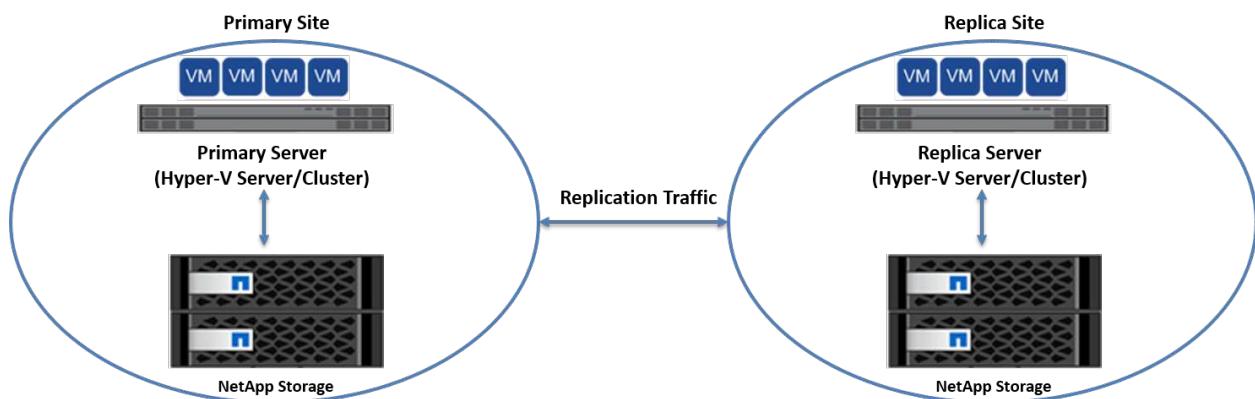


Further reading

For information about deploying storage live migration, see [Appendix E: Deploy Hyper-V Storage Live Migration](#).

Hyper-V Replica: Disaster recovery for virtual machines

Hyper-V Replica replicates the Hyper-V VMs from a primary site to replica VMs on a secondary site, asynchronously providing disaster recovery for the VMs. The Hyper-V server at the primary site hosting the VMs is known as the primary server; the Hyper-V server at the secondary site that receives replicated VMs is known as the replica server. A Hyper-V Replica example scenario is shown in the following figure. You can use Hyper-V Replica for VMs between Hyper-V servers that are part of a failover cluster or between independent Hyper-V servers that are not part of any cluster.



Replication

After Hyper-V Replica is enabled for a VM on the primary server, initial replication creates an identical VM on the replica server. After the initial replication, Hyper-V Replica maintains a log file for the VHDs of the VM. The

log file is replayed in reverse order to the replica VHD in accordance with the replication frequency. This log and the use of reverse order make sure that the latest changes are stored and replicated asynchronously. If replication does not occur in line with the expected frequency, an alert is issued.

Extended replication

Hyper-V Replica supports extended replication in which a secondary replica server can be configured for disaster recovery. A secondary replica server can be configured for the replica server to receive the changes on the replica VMs. In an extended replication scenario, the changes on the primary VMs on the primary server are replicated to the replica server. Then the changes are replicated to the extended replica server. The VMs can be failed over to the extended replica server only when both primary and replica servers go down.

Failover

Failover is not automatic; the process must be manually triggered. There are three types of failover:

- **Test failover.** This type is used to verify that a replica VM can start successfully on the replica server and is initiated on the replica VM. This process creates a duplicate test VM during failover and does not affect regular production replication.
- **Planned failover.** This type is used to fail over VMs during planned downtime or expected outages. This process is initiated on the primary VM, which must be turned off on the primary server before a planned failover is run. After the machine fails over, Hyper-V Replica starts the replica VM on the replica server.
- **Unplanned failover.** This type is used when unexpected outages occur. This process is initiated on the replica VM and should be used only if the primary machine fails.

Recovery

When you configure replication for a VM, you can specify the number of recovery points. Recovery points represent points in time from which data can be recovered from a replicated machine.

Further reading

- For information about deploying Hyper-V Replica outside a clustered environment, see the section "[Deploy Hyper-V Replica Outside of a Clustered Environment](#)."
- For information about deploying Hyper-V Replica in a clustered environment, see the section "[Deploy Hyper-V Replica in a Clustered Environment](#)."

Storage efficiency

ONTAP provides industry leading storage efficiency for virtualized environments including Microsoft Hyper-V. NetApp also offers storage efficiency guarantee programs.

NetApp deduplication

NetApp deduplication works by removing duplicate blocks at the storage volume level, storing only one physical copy, regardless of how many logical copies are present. Therefore, deduplication creates the illusion that there are numerous copies of that block. Deduplication automatically removes duplicate data blocks on a 4KB block level across an entire volume. This process reclaims storage to achieve space and potential performance savings by reducing the number of physical writes to the disk. Deduplication can provide more than 70% space savings in Hyper-V environments.

Thin provisioning

Thin provisioning is an efficient way to provision storage because the storage is not preallocated up front. In other words, when a volume or LUN is created using thin provisioning, the space on the storage system is unused. The space remains unused until the data is written to the LUN or volume and only the necessary space to store the data is used. NetApp recommends enabling thin provisioning on the volume and disabling LUN reservation.

Quality of Service

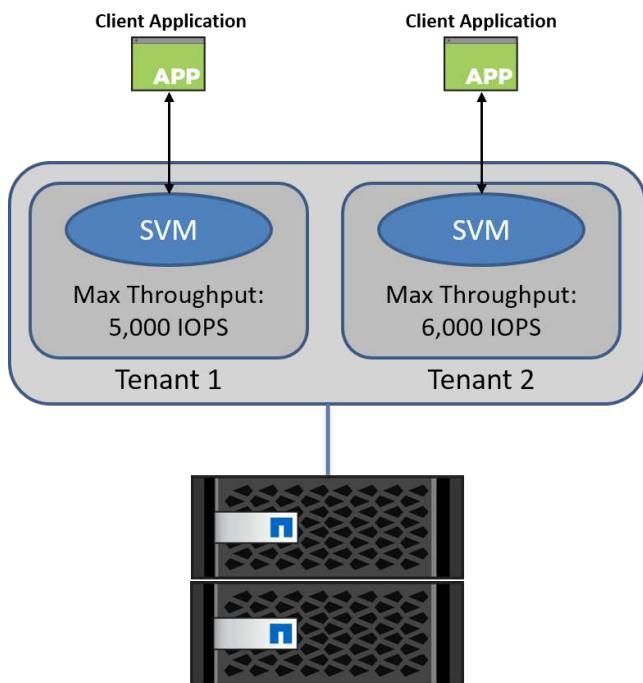
Storage QoS in NetApp ONTAP enables you to group storage objects and set throughput limits on the group. Storage QoS can be used to limit the throughput to workloads and to monitor workload performance. With this ability, a storage administrator can separate workloads by organization, application, business unit, or production or development environments.

In enterprise environments, storage QoS helps to achieve the following:

- Prevents user workloads from affecting each other.
- Protects critical applications that have specific response times that must be met in IT-as-a-service (ITaaS) environments.
- Prevents tenants from affecting each other.
- Avoids performance degradation with the addition of each new tenant.

QoS allows you to limit the amount of I/O sent to an SVM, a flexible volume, a LUN, or a file. I/O can be limited by the number of operations or the raw throughput.

The following figure illustrates SVM with its own QoS policy enforcing a maximum throughput limit.



To configure an SVM with its own QoS policy and monitor policy group, run the following commands on your ONTAP cluster:

```
# create a new policy group pg1 with a maximum throughput of 5,000 IOPS
cluster::> qos policy-group create pg1 -vserver vs1 -max-throughput
5000iops
```

```
# create a new policy group pg2 without a maximum throughput
cluster::> qos policy-group create pg2 -vserver vs2
```

```
# monitor policy group performance
cluster::> qos statistics performance show
```

```
# monitor workload performance
cluster::> qos statistics workload performance show
```

Security

ONTAP provides a secure storage system for the Windows operating system.

Windows Defender Antivirus

Windows Defender is antimalware software installed and enabled on Windows Server by default. This software actively protects Windows Server against known malware and can regularly update antimalware definitions through Windows Update. NetApp LUNs and SMB shares can be scanned using Windows Defender.

Further reading

For further information, see the [Windows Defender Overview](#).

BitLocker

BitLocker drive encryption is a data protection feature continued from Windows Server 2012. This encryption protects physical disks, LUNs, and CSVs.

Best practice

Before enabling BitLocker, the CSV must be put into maintenance mode. Therefore, NetApp recommends that decisions pertaining to BitLocker-based security be made before creating VMs on the CSV to avoid downtime.

Deploy Nano server

Learn about deploying Microsoft Windows Nano Server.

Deployment

To deploy a Nano Server as a Hyper-V host, complete the following steps:

1. Log in to Windows Server as a member of the administrator group.
2. Copy the NanoServerImageGenerator folder from the \NanoServer folder in the Windows Server ISO to the local hard disk.
3. To create a Nano Server VHD/VHDX, complete the following steps:
 - a. Start Windows PowerShell as an administrator, navigate to the copied NanoServerImageGenerator folder on the local hard disk, and run the following cmdlet:

```
Set-ExecutionPolicy RemoteSigned  
Import-Module .\NanoServerImageGenerator -Verbose
```

- b. Create a VHD for the Nano Server as a Hyper-V host by running the following PowerShell cmdlet. This command prompts you for an administrator password for the new VHD.

```
New-NanoServerImage -Edition Standard -DeploymentType Guest  
-MediaPath <"input the path to the root of the contents of Windows  
Server 2016 ISO"> -TargetPath <"input the path, including the  
filename and extension where the resulting VHD/VHDX will be created">  
-ComputerName <"input the name of the nano server computer you are  
about to create"> -Compute
```

- c. In the following example, we create a Nano Server VHD with the feature Hyper-V host with failover clustering enabled. This example creates a Nano Server VHD from an ISO mounted at f:\. The newly created VHD is placed in a folder named NanoServer in the folder from where the cmdlet is run. The computer name is NanoServer and the resulting VHD contains the standard edition of Windows Server.

```
New-NanoServerImage -Edition Standard -DeploymentType Guest  
-MediaPath f:\ -TargetPath .\NanoServer.vhd -ComputerName NanoServer  
-Compute -Clustering
```

- d. With the cmdlet New-NanoServerImage, configure parameters that set the IP address, the subnet mask, the default gateway, the DNS server, the domain name, and so on.
4. Use the VHD in a VM or physical host to deploy Nano Server as a Hyper-V host:
 - a. For deployment on a VM, create a new VM in Hyper-V Manager and use the VHD created in Step 3.
 - b. For deployment on a physical host, copy the VHD to the physical computer and configure it to boot from this new VHD. First, mount the VHD, run bcdboot e:\windows (where the VHD is mounted under E:\), unmount the VHD, restart the physical computer, and boot to the Nano Server.
5. Join the Nano Server to a domain (optional):
 - a. Log in to any computer in the domain and create a data blob by running the following PowerShell cmdlet:

```
$domain = "<input the domain to which the Nano Server is to be joined>"  
$nanoserver = "<input name of the Nano Server>"
```

```
djoin.exe /provision /domain $domain /machine $nanoserver /savefile C:\temp\odblob /reuse
```

b. Copy the odblob file to the Nano Server by running the following PowerShell cmdlets on a remote machine:

```
$nanoserver = "<input name of the Nano Server>"  
$nanouname = ""<input username of the Nano Server>"  
$nanopwd = ""<input password of the Nano Server>"
```

```
$filePath = 'c:\temp\odblob'  
$fileContents = Get-Content -Path $filePath -Encoding Unicode
```

```
$securenanopwd = ConvertTo-SecureString -AsPlainText -Force $nanopwd  
$nanosecurecred = new-object management.automation.pscredential  
$nanouname, $securenanopwd
```

```
Invoke-Command -VMName $nanoserver -Credential $nanosecurecred  
-ArgumentList @($filePath,$fileContents) -ScriptBlock \  
param($filePath,$data)  
New-Item -ItemType directory -Path c:\temp  
Set-Content -Path $filePath -Value $data -Encoding Unicode  
cd C:\temp  
djoin /requestodj /loadfile c:\temp\odblob /windowspath  
c:\windows /localos  
}
```

c. Reboot the Nano Server.

Connect to Nano Server

To connect to the Nano Server remotely using PowerShell, complete the following steps:

1. Add the Nano Server as a trusted host on the remote computer by running the following cmdlet on the remote server:

```
Set-Item WSMan:\localhost\Client\TrustedHosts "<input IP Address of the Nano Server>"
```

2. If the environment is safe and if you want to set all the hosts to be added as trusted hosts on a server, run the following command:

```
Set-Item WSMan:\localhost\Client\TrustedHosts *
```

3. Start the remote session by running the following cmdlet on the remote server. Provide the password for the Nano Server when prompted.

```
Enter-PSSession -ComputerName "<input IP Address of the Nano Server>"  
-Credential ~\Administrator
```

To connect to the Nano Server remotely using GUI management tools from a remote Windows Server, complete the following commands:

4. Log in to the Windows Server as a member of the administrator group.
5. Start Server Manager.
6. To manage a Nano Server remotely from Server Manager, right-click All Servers, click Add Servers, provide the Nano Server's information, and add it. You can now see the Nano Server listed in the server list. Select the Nano Server, right-click it, and start managing it with the various options provided.
7. To manage services on a Nano Server remotely, complete the following steps:
 - a. Open Services from the Tools section of Server Manager.
 - b. Right-click Services (Local).
 - c. Click Connect to Server.
 - d. Provide the Nano Server details to view and manage the services on the Nano Server.
8. If the Hyper-V role is enabled on the Nano Server, complete the following steps to manage it remotely from Hyper-V Manager:
 - a. Open Hyper-V Manager from the Tools section of Server Manager.
 - b. Right-click Hyper-V Manager.
 - c. Click Connect to Server and provide the Nano Server details. Now the Nano Server can be managed as a Hyper-V server to create and manage VMs on top of it.
9. If the failover clustering role is enabled on the Nano Server, complete the following steps to manage it remotely from the failover cluster manager:
 - a. Open Failover Cluster Manager from the Tools section of Server Manager.
 - b. Perform clustering-related operations with the Nano Server.

Deploy Hyper-V cluster

This appendix describes deploying a Hyper-V cluster.

Prerequisites

- At least two Hyper-V servers exist connected to each other.
- At least one virtual switch is configured on each Hyper-V server.
- The failover cluster feature is enabled on each Hyper-V server.
- SMB shares or CSVs are used as shared storage to store VMs and their disks for Hyper-V clustering.
- Storage should not be shared between different clusters. You should have only one CSV/CIFS share per cluster.
- If the SMB share is used as shared storage, then permissions on the SMB share must be configured to grant access to the computer accounts of all the Hyper-V servers in the cluster.

Deployment

1. Log in to one of the Windows Hyper-V servers as a member of the administrator group.
2. Start Server Manager.
3. In the Tools section, click Failover Cluster Manager.
4. Click the Create Cluster from Actions menu.
5. Provide details for the Hyper-V server that is part of this cluster.
6. Validate the cluster configuration. Select Yes when prompted for cluster configuration validation and select the tests required to validate whether the Hyper-V servers pass the prerequisites to be part of the cluster.
7. After validation succeeds, the Create Cluster wizard is started. In the wizard, provide the cluster name and the cluster IP address for the new cluster. A new failover cluster is then created for the Hyper-V server.
8. Click the newly created cluster in Failover Cluster Manager and manage it.
9. Define shared storage for the cluster to use. It can be either an SMB share or a CSV.
10. Using an SMB share as shared storage requires no special steps.
 - Configure a CIFS share on a NetApp storage controller. To do so, see the section "[Provisioning in SMB Environments](#)".
11. To use a CSV as shared storage, complete the following steps:
 - a. Configure LUNs on a NetApp storage controller. To do so, see the section "[Provisioning in SAN Environments](#)".
 - b. Make sure that all the Hyper-V servers in the failover cluster can see the NetApp LUNs. To do this for all the Hyper-V servers that are part of the failover cluster, make sure that their initiators are added to the initiator group on NetApp storage. Also be sure that their LUNs are discovered and make sure that MPIO is enabled.
 - c. On any one of the Hyper-V servers in the cluster, complete the following steps:
 - i. Take the LUN online, initialize the disk, create a new simple volume, and format it using NTFS or ReFS.
 - ii. In Failover Cluster Manager, expand the cluster, expand Storage, right-click Disks, and then click Add Disks. Doing so opens the Add Disks to a Cluster wizard showing the LUN as a disk. Click OK to add the LUN as a disk.
 - iii. Now the LUN is named 'Clustered Disk' and is shown as Available Storage under Disks.
 - d. Right-click the LUN ('Clustered Disk') and click Add to Cluster Shared Volumes. Now the LUN is shown as a CSV.

- e. The CSV is simultaneously visible and accessible from all the Hyper-V servers of the failover cluster at its local location C:\ClusterStorage\.

12. Create a highly available VM:

- a. In Failover Cluster Manager, select and expand the cluster created previously.
- b. Click Roles and then click Virtual Machines in Actions. Click New Virtual Machine.
- c. Select the node from the cluster where the VM should reside.
- d. In the Virtual Machine Creation wizard, provide the shared storage (SMB share or CSV) as the path to store the VM and its disks.
- e. Use Hyper-V Manager to set the shared storage (SMB share or CSV) as the default path to store the VM and its disks for a Hyper-V server.

13. Test planned failover. Move VMs to another node using live migration, quick migration, or storage migration (move). Review [Live Migration in a Clustered Environment](#) for more details.

14. Test unplanned failover. Stop cluster service on the server owning the VM.

Deploy Hyper-V Live Migration in a clustered environment

This appendix describes deploying live migration in a clustered environment.

Prerequisites

To deploy live migration, you need to have Hyper-V servers configured in a failover cluster with shared storage. Review [Deploy Hyper-V Cluster](#) for more details.

Deployment

To use live migration in a clustered environment, complete the following steps:

1. In Failover Cluster Manager, select and expand the cluster. If the cluster is not visible, then click Failover Cluster Manager, click Connect to Cluster, and provide the cluster name.
2. Click Roles, which lists all the VMs available in a cluster.
3. Right-click on the VM and click Move. Doing so provides you with three options:
 - **Live migration.** You can select a node manually or allow the cluster to select the best node. In live migration, the cluster copies the memory used by the VM from the current node to another node. Therefore, when the VM is migrated to another node, the memory and state information needed by the VM are already in place for the VM. This migration method is nearly instantaneous, but only one VM can be live migrated at a time.
 - **Quick migration.** You can select a node manually or allow the cluster to select the best node. In quick migration, the cluster copies the memory used by a VM to a disk in storage. Therefore, when the VM is migrated to another node, the memory and state information needed by the VM can be quickly read from the disk by the other node. With quick migration, multiple VMs can be migrated simultaneously.
 - **Virtual machine storage migration.** This method uses the Move Virtual Machine Storage wizard. With this wizard, you can select the VM disk along with other files to be moved to another location, which can be a CSV or an SMB share.

Deploy Hyper-V Live Migration outside a clustered environment

This section describes the deployment of Hyper-V live migration outside a clustered

environment.

Prerequisites

- Standalone Hyper-V servers with independent storage or shared SMB storage.
- The Hyper-V role installed on both the source and destination servers.
- Both Hyper-V servers belong to the same domain or to domains that trust each other.

Deployment

To perform live migration in a non-clustered environment, configure source and destination Hyper-V servers so that they can send and receive live migration operations. On both Hyper-V servers, complete the following steps:

1. Open Hyper-V Manager from the Tools section of Server Manager.
2. In Actions, click Hyper-V Settings.
3. Click Live Migrations and select Enable Incoming and Outgoing Live Migrations.
4. Choose whether to allow live migration traffic on any available network or only on specific networks.
5. Optionally, you can configure the authentication protocol and performance options from the Advanced section of Live Migrations.
6. If CredSSP is used as the authentication protocol, then make sure to log in to the source Hyper-V server from the destination Hyper-V server before moving the VM.
7. If Kerberos is used as the authentication protocol, then configure the constrained delegation. Doing so requires Active Directory domain controller access. To configure the delegation, complete the following steps:
 - a. Log in to the Active Directory domain controller as the administrator.
 - b. Start Server Manager.
 - c. In the Tools section, click Active Directory Users and Computers.
 - d. Expand the domain and click Computers.
 - e. Select the source Hyper-V server from the list, right-click it, and click Properties.
 - f. In the Delegation tab, select Trust This Computer for Delegation to Specified Services Only.
 - g. Select Use Kerberos Only.
 - h. Click Add, which opens the Add Services wizard.
 - i. In Add Services, click Users and Computers, which opens Select Users or Computers.
 - j. Provide the destination Hyper-V server name and click OK.
 - To move VM storage, select CIFS.
 - To move VMs, select the Microsoft Virtual System Migration service.
 - k. In the Delegation tab, click OK.
 - l. From the Computers folder, select the destination Hyper-V server from the list and repeat the process. In Select Users or Computers, provide the source Hyper-V server name.
8. Move the VM.
 - a. Open Hyper-V Manager.

- b. Right-click a VM and click Move.
- c. Choose Move the Virtual Machine.
- d. Specify the destination Hyper-V server for the VM.
- e. Choose the move options. For Shared Live Migration, choose Move Only the Virtual Machine. For Shared Nothing Live Migration, choose any of the other two options based on your preferences.
- f. Provide the location for the VM on the destination Hyper-V server based on your preferences.
- g. Review the summary and click OK to move the VM.

Deploy Hyper-V Storage Live Migration

Learn how to configure Hyper-V storage live migration

Prerequisites

- You must have a standalone Hyper-V server with independent storage (DAS or LUN) or SMB storage (local or shared among other Hyper-V servers).
- The Hyper-V server must be configured for live migration. Review the section on deployment in [Live Migration Outside of a Clustered Environment](#).

Deployment

1. Open Hyper-V Manager.
2. Right-click a VM and click Move.
3. Select Move the Virtual Machine's Storage.
4. Select options for moving the storage based on your preferences.
5. Provide the new location for the VM's items.
6. Review the summary and click OK to move the VM's storage.

Deploy Hyper-V Replica outside a clustered environment

This appendix describes deploying Hyper-V Replica outside a clustered environment.

Prerequisites

- You need standalone Hyper-V servers located in the same or separate geographical locations serving as primary and replica servers.
- If separate sites are used, then the firewall at each site must be configured to allow communication between the primary and replica servers.
- The replica server must have enough space to store the replicated workloads.

Deployment

1. Configure the replica server.
 - a. So that the inbound firewall rules allow incoming replication traffic, run the following PowerShell cmdlet:

```
Enable-Netfirewallrule -displayname "Hyper-V Replica HTTP Listener (TCP-In)"
```

- b. Open Hyper-V Manager from the Tools section of Server Manager.
- c. Click Hyper-V Settings from Actions.
- d. Click Replication Configuration and select Enable This Computer as a Replica Server.
- e. In the Authentication and Ports section, select the authentication method and port.
- f. In the Authorization and Storage section, specify the location to store the replicated VMs and files.

2. Enable VM replication for VMs on the primary server. VM replication is enabled on a per-VM basis and not for the entire Hyper-V server.
 - a. In Hyper-V Manager, right-click a VM and click Enable Replication to open the Enable Replication wizard.
 - b. Provide the name of the replica server where the VM must be replicated.
 - c. Provide the authentication type and the replica server port that was configured to receive replication traffic on the replica server.
 - d. Select the VHDs to be replicated.
 - e. Choose the frequency (duration) at which the changes are sent to the replica server.
 - f. Configure recovery points to specify the number of recovery points to maintain on the replica server.
 - g. Choose Initial Replication Method to specify the method to transfer the initial copy of the VM data to the replica server.
 - h. Review the summary and click Finish.
 - i. This process creates a VM replica on the replica server.

Replication

1. Run a test failover to make sure that the replica VM functions properly on the replica server. The test creates a temporary VM on the replica server.
 - a. Log in to the replica server.
 - b. In Hyper-V Manager, right-click a replica VM, click Replication, and click Test Failover.
 - c. Choose the recovery point to use.
 - d. This process creates a VM of the same name appended with -Test.
 - e. Verify the VM to make sure that everything works well.
 - f. After failover, the replica test VM is deleted if you select Stop Test Failover for it.
2. Run a planned failover to replicate the latest changes on the primary VM to the replica VM.
 - a. Log in to the primary server.
 - b. Turn off the VM to be failed over.
 - c. In Hyper-V Manager, right-click the turned-off VM, click Replication, and click Planned Failover.
 - d. Click Failover to transfer the latest VM changes to the replica server.
3. Run an unplanned failover in the case of primary VM failure.
 - a. Log in to the replica server.

- b. In Hyper-V Manager, right-click a replica VM, click Replication, and click Failover.
- c. Choose the recovery point to use.
- d. Click Failover to fail over the VM.

Deploy Hyper-V replica in a clustered environment

Learn how to deploy and configure Hyper-V replica with Windows Server Failover Cluster.

Prerequisites

- You need to have Hyper-V clusters located in the same or in separate geographical locations serving as primary and replica clusters. Review [Deploy Hyper-V Cluster](#) for more details.
- If separate sites are used, then the firewall at each site must be configured to allow communication between the primary and replica clusters.
- The replica cluster must have enough space to store the replicated workloads.

Deployment

1. Enable firewall rules on all the nodes of a cluster. Run the following PowerShell cmdlet with admin privileges on all the nodes in both the primary and replica clusters.

```
# For Kerberos authentication
get-clusternode | ForEach-Object \{Invoke-command -computername $_.name
-scriptblock \{Enable-Netfirewallrule -displayname "Hyper-V Replica HTTP
Listener (TCP-In)"\}\}
```

```
# For Certificate authentication
get-clusternode | ForEach-Object \{Invoke-command -computername $_.name
-scriptblock \{Enable-Netfirewallrule -displayname "Hyper-V Replica
HTTPS Listener (TCP-In)"\}\}
```

2. Configure the replica cluster.
 - a. Configure the Hyper-V Replica broker with a NetBIOS name and IP address to use as the connection point to the cluster that is used as the replica cluster.
 - i. Open Failover Cluster Manager.
 - ii. Expand the cluster, click Roles, and click the Configure Role from Actions pane.
 - iii. Select Hyper-V Replica Broker in the Select Role page.
 - iv. Provide the NetBIOS name and IP address to be used as the connection point to the cluster (client access point).
 - v. This process creates a Hyper-V Replica broker role. Verify that it comes online successfully.
 - b. Configure replication settings.
 - i. Right-click the replica broker created in the previous steps and click Replication Settings.
 - ii. Select Enable This Cluster as a Replica Server.

- iii. In the Authentication and Ports section, select the authentication method and port.
- iv. In the authorization and storage section, select the servers allowed to replicate VMs to this cluster. Also, specify the default location where the replicated VMs are stored.

Replication

Replication is similar to the process described in the section [Replica Outside a Clustered Environment](#).

Where to find additional information

Additional resources for Microsoft Windows and Hyper-V

- ONTAP concepts
<https://docs.netapp.com/us-en/ontap/concepts/introducing-data-management-software-concept.html>
- Best practices for modern SAN
<https://www.netapp.com/media/10680-tr4080.pdf>
- NetApp All-SAN Array Data Availability and Integrity with the NetApp ASA
<https://www.netapp.com/pdf.html?item=/media/85671-tr-4968.pdf>
- SMB documentation
<https://docs.netapp.com/us-en/ontap/smb-admin/index.html>
- Getting Started with Nano Server
<https://technet.microsoft.com/library/mt126167.aspx>
- What's New in Hyper-V on Windows Server
<https://technet.microsoft.com/windows-server-docs/compute/hyper-v/what-s-new-in-hyper-v-on-windows>

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.