



Deploying Microsoft Hyper-V on NetApp Storage

NetApp Solutions

NetApp
May 14, 2024

Table of Contents

- Deploying Microsoft Hyper-V on NetApp Storage 1
 - Deploying Microsoft Hyper-V on NetApp Storage 1
 - Deploying Microsoft Hyper-V on NetApp Storage: Pre-Requisites 1
 - Deploying Microsoft Hyper-V on NetApp Storage: Considerations 7
 - Deploying Microsoft Hyper-V on NetApp Storage: Data Protection 13
 - Deploying Microsoft Hyper-V on NetApp Storage: Conclusion 19
 - Deploying Microsoft Hyper-V on NetApp Storage: Migration Script 20

Deploying Microsoft Hyper-V on NetApp Storage

Deploying Microsoft Hyper-V on NetApp Storage

The Windows Server platform uses the Hyper-V role to provide virtualization technology. Hyper-V is one of many optional roles that are offered with Windows Server.

Overview

The Hyper-V role enables us to create and manage a virtualized computing environment by using virtualization technology built into Windows Server. The Hyper-V technology virtualizes hardware to provide an environment in which you can run multiple operating systems at the same time on one physical computer. Hyper-V enables you to create and manage virtual machines and their resources. Each virtual machine is an isolated, virtualized computer system that can run its own operating system. Hyper-V provides infrastructure to virtualize applications and workloads that supports a variety of business goals aimed at improving efficiency and reducing costs which is a perfect alternative to VMware® vSphere, especially when organizations are looking for co-existence of multiple hypervisors during the current market conditions.

Audience

This document describes the architecture and deployment procedures for the Hyper-V Cluster configuration with the NetApp ONTAP systems. The intended audience for this document includes sales engineers, field consultants, professional services, IT managers, partner engineers, and customers who want to deploy Hyper-V as the primary or as an alternate hypervisor.

Architecture

The architecture described in this document specifically includes Microsoft® Windows Server® 2022 and Hyper-V® virtualization. NetApp strongly recommends virtualization software and infrastructure management software as part of every deployment. The configuration uses the best practices for each component to enable a reliable, enterprise-class infrastructure.

Use Case Summary

This document describes the deployment procedures and best practices to set up Hyper-V cluster to optimally perform as a workload on Microsoft Windows Server 2022 using NetApp All-flash FAS and ASA arrays models. The server operating system/hypervisor is Microsoft Windows Server 2022. The guidance covers NetApp storage systems that serve data over storage area network (SAN) and network-attached storage (NAS) protocols.

Deploying Microsoft Hyper-V on NetApp Storage: Pre-Requisites

This topic provides steps to configure and deploy a two-node failover cluster and clustered Hyper-V virtual machines leveraging ONTAP storage system.

Pre-requisites for Deployment Procedure

- All hardware must be certified for the version of Windows Server that you are running, and the complete failover cluster solution must pass all tests in the Validate a Configuration Wizard

- Hyper-V nodes joined to the domain controller (recommended) and appropriate connectivity between each other.
- Every Hyper-V node should be configured identically.
- Network adapters and designated virtual switches configured on each Hyper-V server for segregated traffic for mgmt, iSCSI, SMB, live migrate.
- 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. Plan for one or multiple 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 nodes in the cluster.

For more information, see:

- [System Requirements for Hyper-V on Windows Server](#)
- [Validate Hardware for a Failover Cluster](#)
- [Deploy a Hyper-V Cluster](#)

Installing Windows Features

The following steps describe how to install the required Windows Server 2022 features.

All Hosts

1. Prepare the windows OS 2022 with necessary updates and device drivers on all the designated nodes.
2. Log into each Hyper-V node using the administrator password entered during installation.
3. Launch a PowerShell prompt by right clicking the PowerShell icon in the taskbar and selecting `Run as Administrator`.
4. Add the Hyper-V, MPIO, and clustering features.

```
Add-WindowsFeature Hyper-V, Failover-Clustering, Multipath-IO `-
IncludeManagementTools -Restart
```

Configuring Networks

Proper network planning is key to achieving fault tolerant deployment. Setting up distinct physical network adapters for each type of traffic was the standard suggestion for a failover cluster. With the ability to add virtual network adapters, switch embedded teaming (SET) and features like Hyper-V QoS introduced, condense network traffic on fewer physical adapters. Design the network configuration with quality of service, redundancy, and traffic isolation in mind. Configuring network isolation techniques like VLANs in conjunction with traffic isolation techniques provides redundancy for the traffic and quality of service which would improve and add consistency to storage traffic performance.

It is advised to separate and isolate specific workloads using multiple logical and/or physical networks. Typical network traffic examples that are typically divided into segments are as follows:

- iSCSI Storage network.

- CSV (Cluster Shared Volume) or Heartbeat network.
- Live Migration
- VM network
- Management network

Note: When iSCSI is used with dedicated NICs, then using any teaming solution is not recommended and MPIO/DSPM should be used.

Note: Hyper-V networking best practices also do not recommend using NIC teaming for SMB 3.0 storage networks in Hyper-V environment.

For additional information, refer to [Plan for Hyper-V networking in Windows Server](#)

Deciding on Storage Design for Hyper-V

Hyper-V supports NAS (SMB3.0) and Block storage (iSCSI/FC) as the backing storage for virtual machines. NetApp supports SMB3.0, iSCSI and FC protocol which can be used as native storage for VMs - Cluster Shared Volumes (CSV) using iSCSI/FC and SMB3. Customers can also use SMB3 and iSCSI as guest connected storage options for workloads that require direct access to the storage. ONTAP provides flexible options with unified storage (All Flash Array) - for workload that requires mixed protocol access and SAN optimized storage (All SAN Array) for SAN only configurations.

The decision to use SMB3 vs iSCSI/FC is driven by the existing infrastructure in place today, SMB3/iSCSI allow customers to use existing network infrastructure. For customers that have existing FC infrastructure can leverage that infrastructure and present storage as FC based Clustered Shared Volumes.

Note: A NetApp storage controller running ONTAP software can support the following workloads in a Hyper-V environment:

- VMs hosted on continuously available SMB 3.0 shares
- VMs hosted on Cluster Shared Volume (CSV) LUNs running on iSCSI or FC
- In-Guest storage and pass through disks to guest virtual machines

Note: Core ONTAP features such as thin provisioning, deduplication, compression, data compaction, flex clones, snapshots, and replication work seamlessly in the background regardless of the platform or operating system and provide significant value for the Hyper-V workloads. The default settings for these features are optimal for Windows Server and Hyper-V.

Note: MPIO is supported on the guest VM using in-guest initiators if multiple paths are available to the VM, and the multipath I/O feature is installed and configured.

Note: ONTAP supports all major industry-standard client protocols: NFS, SMB, FC, FCoE, iSCSI, NVMe/FC, and S3. However, NVMe/FC and NVMe/TCP are not supported by Microsoft.

Installing NetApp Windows iSCSI Host Utilities

The following section describes how to perform an unattended installation of the NetApp Windows iSCSI Host Utilities. For detailed information regarding the installation see the [Install Windows Unified Host Utilities 7.2 \(or the latest supported version\)](#)

All Hosts

1. Download [Windows iSCSI Host Utilities](#)
2. Unblock the downloaded file.

```
Unblock-file ~\Downloads\netapp_windows_host_utilities_7.2_x64.msi
```

3. Install the Host Utilities.

```
~\Downloads\netapp_windows_host_utilities_7.2_x64.msi /qn  
"MULTIPATHING=1"
```

Note: The system will reboot during this process.

Configuring Windows Host iSCSI initiator

The following steps describe how to configure the built in Microsoft iSCSI initiator.

All Hosts

1. Launch a PowerShell prompt by right clicking the PowerShell icon in the taskbar and selecting Run as Administrator.
2. Configure the iSCSI service to start automatically.

```
Set-Service -Name MSiSCSI -StartupType Automatic
```

3. Start the iSCSI service.

```
Start-Service -Name MSiSCSI
```

4. Configure MPIO to claim any iSCSI device.

```
Enable-MSDSMAutomaticClaim -BusType iSCSI
```

5. Set the default load balance policy of all newly claimed devices to round robin.

```
Set-MSDSMGlobalDefaultLoadBalancePolicy -Policy RR
```

6. Configure an iSCSI target for each controller.

```

New-IscsiTargetPortal -TargetPortalAddress <<iscsia_lif01_ip>>
-InitiatorPortalAddress <iscsia_ipaddress>

New-IscsiTargetPortal -TargetPortalAddress <<iscsib_lif01_ip>>
-InitiatorPortalAddress <iscsib_ipaddress>

New-IscsiTargetPortal -TargetPortalAddress <<iscsia_lif02_ip>>
-InitiatorPortalAddress <iscsia_ipaddress>

New-IscsiTargetPortal -TargetPortalAddress <<iscsib_lif02_ip>>
-InitiatorPortalAddress <iscsib_ipaddress>

```

7. Connect a session for each iSCSI network to each target.

```

Get-IscsiTarget | Connect-IscsiTarget -IsPersistent $true
-IsMultipathEnabled $true -InitiatorPortalAddress <iscsia_ipaddress>

Get-IscsiTarget | Connect-IscsiTarget -IsPersistent $true
-IsMultipathEnabled $true -InitiatorPortalAddress <iscsib_ipaddress>

```

Note: Add multiple sessions (min of 5-8) for increased performance and utilizing the bandwidth.

Creating a Cluster


One Server Only

1. Launch a PowerShell prompt with administrative permissions, by right clicking the PowerShell icon and selecting Run as Administrator`.
2. Create a new cluster.

```

New-Cluster -Name <cluster_name> -Node <hostnames> -NoStorage
-StaticAddress <cluster_ip_address>

```




3. Select the appropriate cluster network for Live migration.
4. Designate the CSV network.

```
(Get-ClusterNetwork -Name Cluster).Metric = 900
```

5. Change the cluster to use a quorum disk.
 - a. Launch a PowerShell prompt with administrative permissions by right clicking the PowerShell icon and selecting 'Run as Administrator'.

```
start-ClusterGroup "Available Storage" | Move-ClusterGroup -Node $env:COMPUTERNAME
```

- b. In Failover Cluster Manager, select Configure Cluster Quorum Settings.



- c. Click Next through the Welcome page.
- d. Select the quorum witness and click Next.
- e. Select Configure a disk witness` and click Next.
- f. Select Disk W: from the available storage and click Next.
- g. Click Next through the confirmation page and Finish on the summary page.

For more detailed information about quorum and witness, see [Configuring and manage quorum](#)

6. Run the Cluster Validation wizard from Failover Cluster Manager to validate deployment.
7. Create CSV LUN to store virtual machine data and create highly available virtual machines via Roles within Failover Cluster Manager.

Deploying Microsoft Hyper-V on NetApp Storage: Considerations

This step is vital to ascertain that the applications, services, and workloads can operate effectively in the Hyper-V environment. Compatibility checks must encompass operating system versions, Windows server versions, application dependencies, database systems, and any specific configurations or customisations that exist in the existing environment.

Right sizing the storage

Before deploying the workload or migrating from existing hypervisor, ensure the workload is sized to meet the required performance. This can be easily done by collecting performance data for each individual VM that collects statistics for CPU (used/provisioned), Memory (used/provisioned), Storage (provisioned/utilized), Network throughput and latency along with aggregation of the Read/Write IOPs, throughput and block size. These parameters are mandatory for have a successful deployment and to correctly size the storage array and workload hosts.

Note: Plan for IOPS and capacity when sizing storage for Hyper-V and associated workloads.

Note: For higher-I/O intensive VMs or those that require lots of resources and capacity, segregate the OS and data disks. Operating system and application binaries change infrequently, and volume crash consistency is acceptable.

Note: Use Guest connected storage (aka in-guest) for high performance data disks than using VHDs. This helps with easier cloning process as well.

Enhance Virtual Machine performance

Choose the right amount of RAM and vCPUs for optimal performance along with attaching multiple disks to a single virtual SCSI controller. Using fixed VHDX is still recommended as the primary choice for virtual disks for deployments and there are no restrictions for using any type of VHDX virtual disks.

Note: Avoid installing unnecessary roles on Windows Server that will not be utilized.

Note: Choose Gen2 as the generation for virtual machines able to load VMs from the SCSI controller and is based on the VMBUS and VSP / VSC architecture for the boot level, which significantly increases the overall VM performance.

Note: Avoid making frequent checkpoints because it has a negative impact on the performance of the VM.

SMB3.0 Design and Consideration

SMB 3.0 file shares can be used as shared storage for Hyper-V. ONTAP supports nondisruptive operations over SMB shares for Hyper-V. Hyper-V can use SMB file shares to store virtual machine files, such as configuration, snapshots, and virtual hard disk (VHD) files. Use dedicated ONTAP CIFS SVM for SMB3.0 based shares for Hyper-V. The volumes used to store virtual machine files must be created with NTFS security-style volumes. Connectivity between Hyper-V hosts and the NetApp array is recommended on a 10GB network if one is available. In case of 1GB network connectivity, NetApp recommends creating an interface group consisting of multiple 1GB ports. Connect each NIC serving SMB multichannel to its dedicated IP subnet so that each subnet provides a single path between the client and server.

Key Points

- Enable SMB multi-channel on ONTAP SVM
- ONTAP CIFS SVMs should have at least one data LIF on each node in a cluster.
- Shares used must be configured with the continuously available property set.
- ONTAP One is now included on every AFF (A-Series and C-Series), All-SAN Array (ASA), and FAS system. Hence there is no separate licenses needed.
- For Shared VHDX, use guest connected iSCSI LUN

Note: ODX is supported and works across protocols. Copying data between a file share and iSCSI or an FCP-attached LUN also utilizes ODX.

Note: Time settings on nodes in the cluster should be set up accordingly. Network Time Protocol (NTP) should be used if the NetApp CIFS server must participate in the Windows Active Directory (AD) domain.

Note: Large MTU values must be enabled through the CIFS server. Small packet sizes might result in performance degradation.

Provisioning SMB volume

1. Verify that the required CIFS server options are enabled on the storage virtual machine (SVM)
2. The following options should be set to true: smb2-enabled smb3-enabled copy-offload-enabled shadowcopy-enabled is-multichannel-enabled is-large-mtu-enabled

```
HV_NestedCluster::> vserver cifs options show -vserver NestedHvsvm01 -fields copy-offload-enabled, is-multichannel-enabled, is-large-mtu-enabled, smb2-enabled, smb3-enabled, copy-offload-enabled, shadowcopy-enabled
vserver      smb2-enabled smb3-enabled copy-offload-enabled shadowcopy-enabled is-multichannel-enabled is-large-mtu-enabled
-----
NestedHvsvm01 true         true         true         true         true         true
```

3. Create NTFS data volumes on the storage virtual machine (SVM) and then configure continuously available shares for use with Hyper-V


```

HV_NestedCluster:~> volume create -vsrver NestedHvVm01 -volume hvdemo0mb -aggregate HV_NestedCluster_01_VM_DISK_1 -size 500GB -security-style ntfs -type smbpath/activate Windows.
[Job 169] Job succeeded: Successful

```

Note: Nondisruptive operations for Hyper-V over SMB do not work correctly unless the volumes used in the configuration are created as NTFS security-style volumes.

4. Enable continuously available and configure NTFS permissions on the share to include Hyper-V nodes with full control.



For detailed best practices guidance, see [Deployment Guidelines and best practices for Hyper-V](#).

For additional information, refer to [SMB server and volume requirements for Hyper-V over SMB](#).

Block Protocol Design and Consideration

Key Points

- Use multipathing (MPIO) on hosts to manage the multiple paths. Create more paths as needed, either to facilitate data mobility operations or to leverage additional I/O resources, but do not exceed the maximum number of paths a host OS can support.
- Install the Host Utilities Kit on hosts accessing the LUNs.
- Create a minimum of 8 volumes.

Note: Use one LUN per volume, thus having 1:1 mapping for LUN to CSV ratio.

- 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.

- SVMs serving data with FCP, or iSCSI need an SVM management interface.

Provisioning iSCSI volume

To provision iSCSI volume, ensure the following pre-requisites are met.


- The storage virtual machine (SVM) should have the iSCSI protocol enabled and the appropriate logical interfaces (LIFs) created.
- The designated aggregate must have enough free space to contain the LUN.

Note: By default, ONTAP uses Selective LUN Map (SLM) to make the LUN accessible only through paths on the node owning the LUN and its high-availability (HA) partner.

- Configure all the iSCSI LIFs on every node for LUN mobility in case the LUN is moved to another node in the cluster.

Steps

1. Use System Manager and navigate to the LUNs window (ONTAP CLI can be used for the same operation).
2. Click Create.
3. Browse and select the designated SVM in which the LUNs to be created and the Create LUN Wizard is displayed.
4. On the General Properties page, select Hyper-V for LUNs containing virtual hard disks (VHDs) for Hyper-V virtual machines.




5. <click on More options> On the LUN Container page, select an existing FlexVol volume otherwise a new volume will be created.
6. <click on More options> On the Initiators Mapping page, click Add Initiator Group, enter the required information on the General tab, and then on the Initiators tab, enter the iSCSI initiator node name of the hosts.
7. Confirm the details, and then click Finish to complete the wizard.

Once the LUN is created, go to the Failover Cluster Manager. To add a disk to CSV, the disk must be added to the Available Storage group of the cluster (if it is not already added), and then add the disk to CSV on the cluster.

Note: The CSV feature is enabled by default in Failover Clustering.

Adding a disk to Available Storage:

1. In Failover Cluster Manager, in the console tree, expand the name of the cluster, and then expand Storage.
2. Right-click Disks, and then select Add Disk. A list appears showing the disks that can be added for use in a failover cluster.
3. Select the disk or disks you want to add, and then select OK.
4. The disks are now assigned to the Available Storage group.
5. Once done, select the disk that was just assigned to Available Storage, right-click the selection, and then select Add to Cluster Shared Volumes.



6. The disks are now assigned to the Cluster Shared Volume group in the cluster. The disks are exposed to each cluster node as numbered volumes (mount points) under the %SystemDrive%\ClusterStorage folder. The volumes appear in the CSVFS file system.

For additional information, refer to [Use Cluster Shared Volumes in a failover cluster](#).

Create highly available virtual machines:

To create a highly available virtual machine, follow the below steps:

1. In Failover Cluster Manager, select or specify the cluster that you want. Ensure that the console tree under the cluster is expanded.
2. Click Roles.
3. In the Actions pane, click Virtual Machines, and then click New Virtual Machine. The New Virtual Machine Wizard appears. Click Next.

4. On the Specify Name and Location page, specify a name for the virtual machine, such as nimdemo. Click Store the virtual machine in a different location, and then type the full path or click Browse and navigate to the shared storage.
5. Assign Memory and configure network adapter to the virtual switch that is associated with the physical network adapter.
6. On the Connect Virtual Hard Disk page, click Create a virtual hard disk.
7. On the Installation Options page, click Install an operating system from a boot CD/DVD-ROM. Under Media, specify the location of the media, and then click Finish.
8. The virtual machine is created. The High Availability Wizard in Failover Cluster Manager then automatically configures the virtual machine for high availability.

Fast Provisioning of Virtual Disks Using ODX Feature

The ODX feature in ONTAP allows making copies of master VHDXs by simply copying a master VHDX file hosted by ONTAP storage system. Because an ODX-enabled copy does not put any data on the network wire, the copy process happens on the NetApp storage side and as a result can be up to six to eight times faster. General considerations for fast provisioning include master sysprepped images stored on file shares and regular copy processes initiated by the Hyper-V host machines.

Note: ONTAP supports ODX for both the SMB and SAN protocols.

Note: To take advantage of the use cases for ODX copy offload pass-through with Hyper-V, the guest operating system must support ODX, and the guest operating system's disks must be SCSI disks backed by storage (either SMB or SAN) that supports ODX. IDE disks on the guest operating system do not support ODX pass-through.

Performance optimization

Although the recommended number of VMs per CSV is subjective, numerous factors determine the optimum number of VMs that can be placed on each CSV or SMB volume. Although most administrators only consider capacity, the amount of concurrent I/O being sent to the VHDx is one of the most key factors for overall performance. The easiest way to control performance is by regulating the number of virtual machines that are placed on each CSV or share. If the concurrent virtual machine I/O patterns are sending too much traffic to the CSV or share, the disk queues fill, and higher latency are generated.

SMB Volume and CSV sizing

Ensure the solution is adequately sized end-to-end to avoid bottlenecks and when a volume is created for Hyper-V VM storage purposes, the best practice is to create a volume no larger than required. Right sizing volumes prevent accidentally placing too many virtual machines on the CSV and decreases the probability of resource contention. Each cluster shared volume (CSV) supports one VM or multiple VMs. The number of VMs to place on a CSV is determined by the workload and business preferences, and how ONTAP storage features such as snapshots and replication will be used. Placing multiple VMs on a CSV is a good starting point in most deployment scenarios. Adjust this approach for specific use cases to meet performance and data protection requirements.

Since volumes and VHDx sizes can be easily increased, if a VM needs extra capacity, it is not necessary to size CSVs larger than required. Diskpart can be used for extending the CSV size or an easier approach is to create a new CSV and migrate the required VMs to the new CSV. For optimal performance, the best practice is to increase the number of CSVs rather than increase their size as an interim measure.

Migration

One of the most common use cases in the current market condition is migration. Customers can use VMM Fabric or other third-party migration tools to migrate VMs. These tools use host level copy to move data from the source platform to the destination platform, which can be time consuming depending on number of virtual machines that are in scope of migration.

Using ONTAP in such scenario's enable quicker migration than using host based migration process. ONTAP also enables swift migration of VMs from one hypervisor to another (ESXi in this case to Hyper-V). VMDK of any size can be converted to VHDx in seconds on NetApp Storage. That is our PowerShell way - It leverages NetApp FlexClone® technology for the rapid conversion of VM hard disks. It also handles the creation and configuration of target and destination VMs.

This process helps minimize downtime and enhances business productivity. It also offers choice and flexibility by reducing licensing costs, lock-in, and commitments to a single vendor. This is also beneficial for organizations looking to optimize VM licensing costs and extend IT budgets.

The following video demonstrates the process to migrate virtual machines from VMware ESX to Hyper-V.

[Zero touch migration from ESX to Hyper-V](#)

For additional information about migration using Flexclone and PowerShell, see the [PowerShell script for migration](#).

Deploying Microsoft Hyper-V on NetApp Storage: Data Protection

Data protection is a key tenant for any production workload. This section describes how to backup and restore Hyper-V virtual machines.

Restore using NetApp Storage snapshot

Backing up VMs and quickly recovering or cloning them are among the great strengths of ONTAP volumes. Use Snapshot copies to make quick FlexClone copies of the VMs or even the whole CSV volume without affecting performance. This enables working with production data without the risk of data corruption when cloning production data volumes and mounting them on QA, staging and development environments. FlexClone volumes are useful for making test copies of production data, without having to double the amount of space required to copy the data.

Keep in mind, Hyper-V nodes assign each disk a unique ID and taking a snapshot of the volume that has respective partition (MBR or GPT) will carry the same unique identification. MBR uses disk signatures and GPT uses GUIDs (Global Unique Identifiers). In case of standalone Hyper-V host, the FlexClone volume can be easily mounted without any conflicts. This is because stand-alone Hyper-V servers can automatically detect duplicate disk IDs and change them dynamically without user intervention. This approach can be used to recover the VM(s) by copying the VHDs as the scenario demands.

While it is straightforward with standalone Hyper-V hosts, the procedure is different for Hyper-V clusters. The recovery process involves mapping the FlexClone volume to a standalone Hyper-V host or using diskpart to manually change the signature by mapping FlexClone volume to a standalone Hyper-V host (it is important because a disk ID conflict results in inability to bring the disk online) and once done, map the FlexClone volume to the cluster.

Backup and Restore using Third party solution

Note: This section uses Commvault, however this is applicable to other third-party solutions.

Leveraging ONTAP snapshots, CommVault IntelliSnap® creates hardware-based snapshots of Hyper-V. Backups can be automated based on the configuration for a Hyper-V hypervisor or VM group, or manually for a VM group or a specific VM. IntelliSnap enables fast protection of Hyper-V environments placing minimal load on the production Virtualization Farm. The integration of IntelliSnap technology with the Virtual Server Agent (VSA) enables NetApp ONTAP Array to complete backups with a large number of virtual machines and data stores in a matter of minutes. Granular access provides individual file and folder recovery from the secondary tier of storage along with the full guest .vhd files.

Prior to configuring the virtualization environment, deploy the proper agents requiring snapshot integration with the Array. Microsoft Hyper-V virtualization environments require the following agents:


- MediaAgent
- Virtual Server Agent (VSA)
- VSS Hardware Provider (Windows Server 2012 and newer operating systems)

Configure NetApp Array using Array Management

The following steps show how to configure IntelliSnap virtual machine backups in an environment utilizing an ONTAP array and Hyper-V.

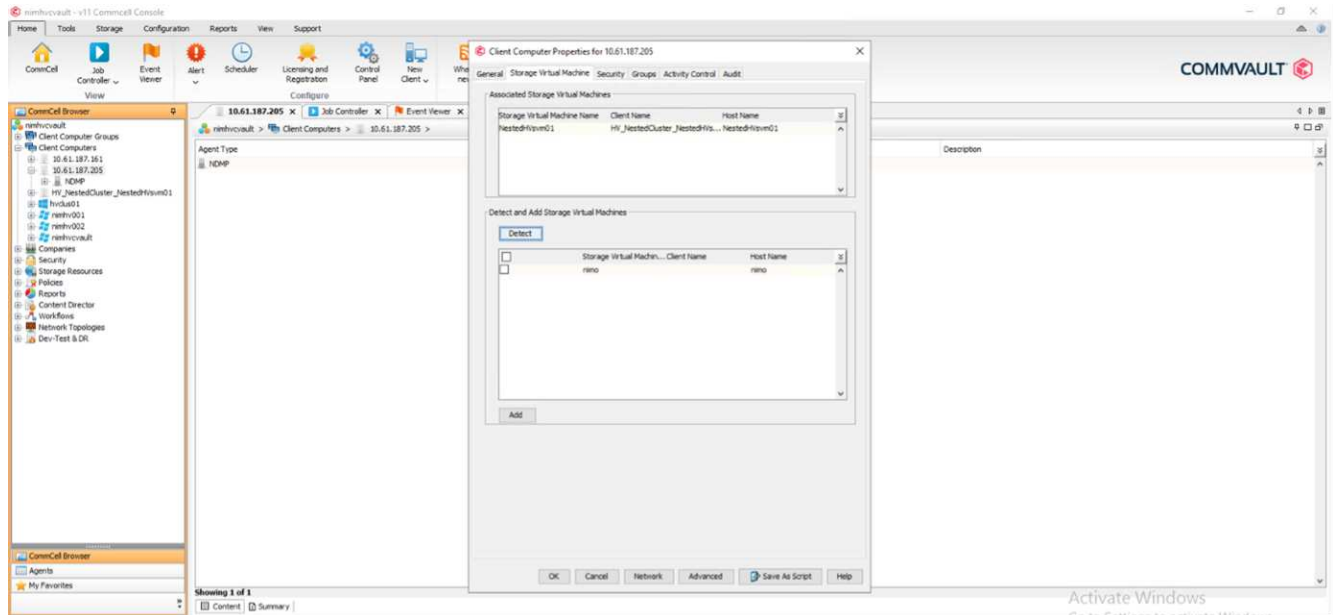
1. On the ribbon in the CommCell Console, click the Storage tab, and then click Array Management.
2. The Array Management dialog box appears.
3. Click Add.

The Array Properties dialog box appears.

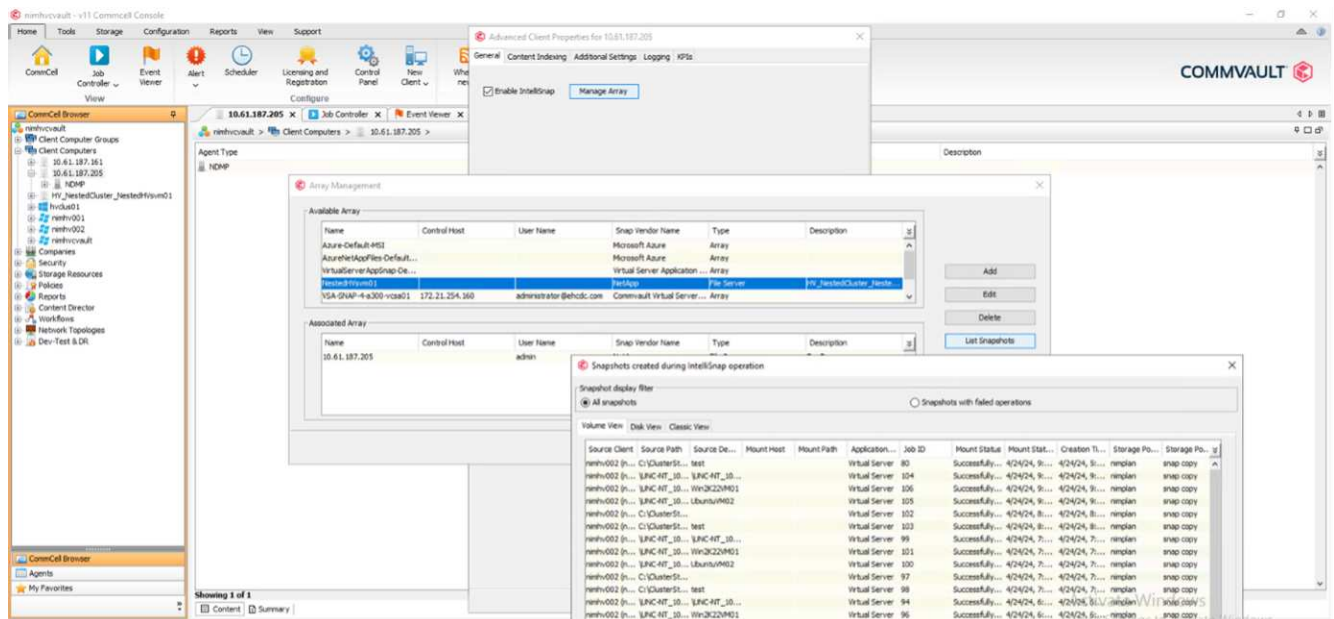


4. On the General tab, specify the following information:
5. From the Snap Vendor list, select NetApp.
6. In the Name box, enter the host name, the fully qualified domain name (FQDN), or the TCP/IP address of the primary file server.

7. On the Array Access Nodes tab, select available media agents.
8. On the Snap Configuration tab, configure Snapshot Configuration Properties according to your needs.
9. Click OK.
10. <Mandatory step> Once done, also configure SVM on the NetApp storage array by using the detect option to automatically detect storage virtual machines (SVM), then choose an SVM, and with the add option, add the SVM in the CommServe database, as an array management entry.



11. Click on Advanced (as shown in the below graphics) and select “Enable IntelliSnap” checkbox.



For detailed steps about configuring the array, see [Configuring NetApp Array](#) and [Configuring Storage Virtual machines on NetApp Arrays](#)

Add Hyper-V as the Hypervisor

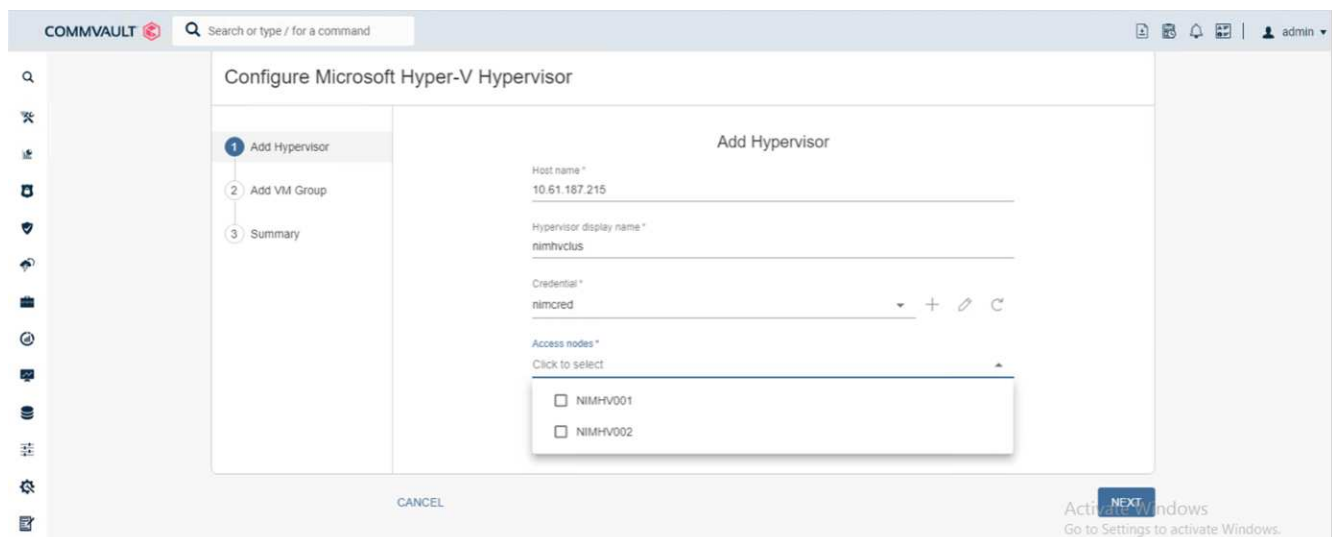
Next step is to add Hyper-V hypervisor and adding a VM group.

Pre-requisites

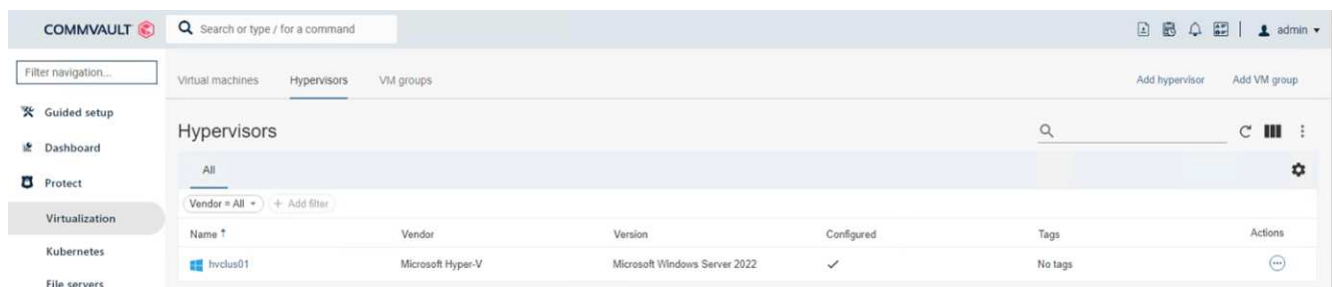
- The hypervisor can be a Hyper-V cluster, a Hyper-V server in a cluster, or a standalone Hyper-V server.
- The user must belong to the Hyper-V administrators' group for Hyper-V Server 2012 and later. For a Hyper-V cluster, the user account must have full cluster permissions (Read and Full Control).
- Identify one or more nodes on which you will install the Virtual Server Agent (VSA) to create access nodes (VSA proxies) for backup and restore operations. To discover Hyper-V servers, the CommServe system must have the VSA installed.
- To use Changed Block Tracking for Hyper-V 2012 R2, select all nodes in the Hyper-V cluster.

The following steps show how to add Hyper-V as a hypervisor.


1. After the core setup is complete, on the Protect tab, click the Virtualization tile.
2. On the Create server backup plan page, type a name for the plan, then provide information about storage, retention, and backup schedules.
3. Now the Add hypervisor page appears > Select vendor: Select Hyper-V (Enter the IP address or FQDN and user credentials)
4. For a Hyper-V server, click Discover nodes. When the Nodes field is populated, select one or more nodes on which to install the Virtual Server Agent.



5. Click Next and the Save.



6. On the Add VM group page, select the virtual machines to be protected (Demogrpf is the VM group created in this case) and enable IntelliSnap option as shown below.



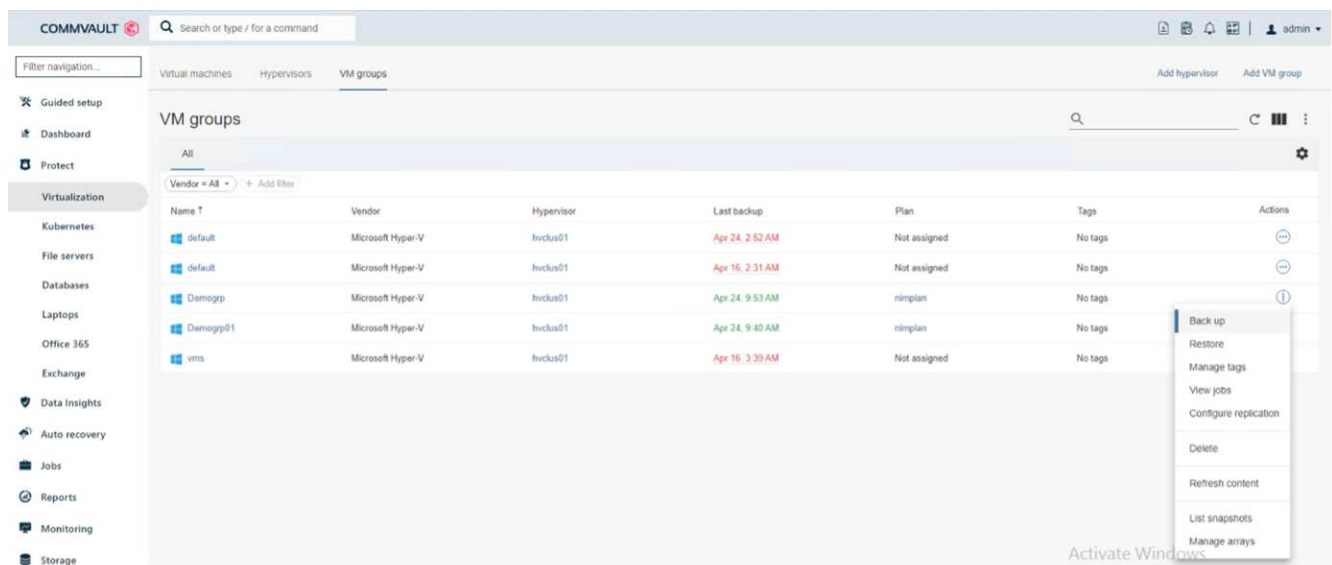
Note: When IntelliSnap is enabled on a VM group, Commvault automatically creates schedule policies for the primary (snap) and backup copies.

7. Click Save.


For detailed steps about configuring the array, see [Adding a Hypervisor](#).

Performing a backup:

1. From the navigation pane, go to Protect > Virtualization. The Virtual machines page appears.
2. Back up the VM or the VM group. In this demo, VM group is selected. In the row for the VM group, click the action button `action_button`, and then select Back up. In this case, `nimplan` is the plan associated against `Demogr` and `Demogr01`.



3. Once the backup is successful, restore points are available as shown in the screen capture. From the snap copy, restore of full VM and restore of guest files and folders can be performed.

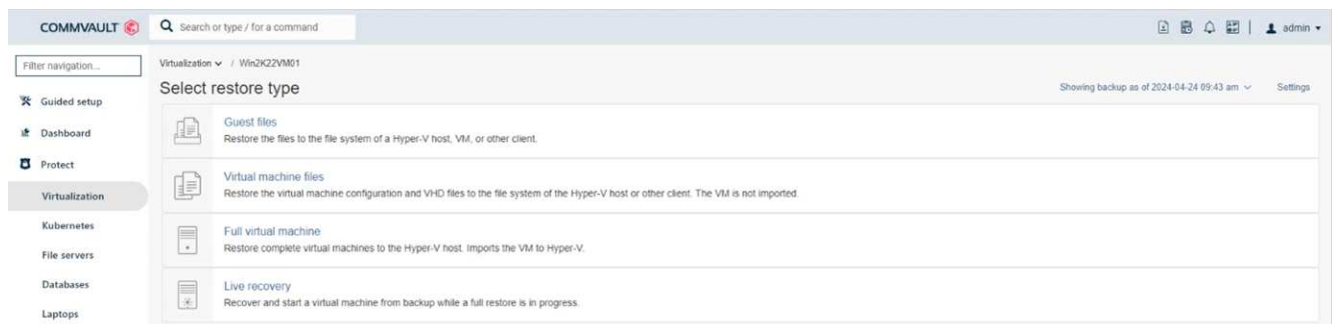


Note: For critical and heavily utilized virtual machines, keep fewer virtual machines per CSV

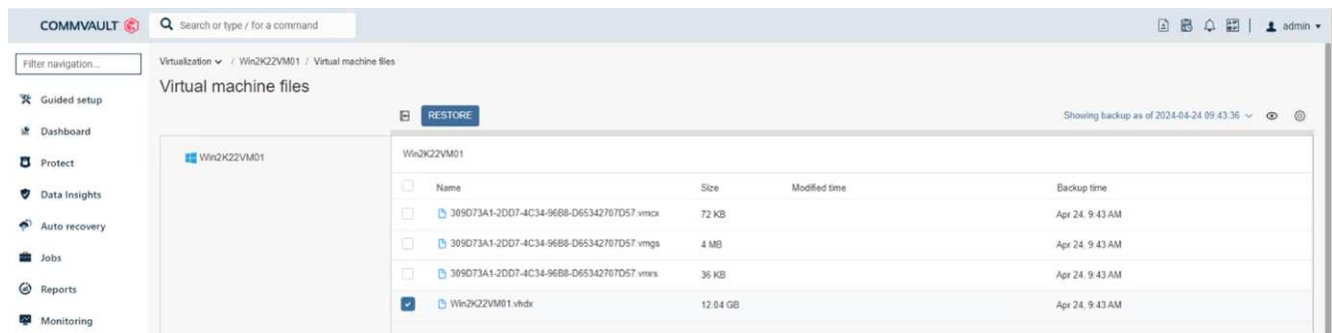
Performing a restore operation:

Restore full VMs, guest files and folders, or virtual disk files via the restore points.

1. From the navigation pane, go to Protect > Virtualization, the Virtual machines page appears.
2. Click the VM groups tab.
3. The VM group page appears.
4. In the VM groups area, click Restore for the VM group that contains the virtual machine.
5. The Select restore type page appears.



6. Select Guest files or Full virtual machine depending on the selection and trigger the restore.



For detailed steps for all supported restore options, see [Restores for Hyper-V](#).

Advanced NetApp ONTAP options

NetApp SnapMirror enables efficient site-to-site storage replication, making disaster recovery rapid, reliable, and manageable to suit today's global enterprises. Replicating data at high speeds over LANs and WANs, SnapMirror provides high data availability and fast recovery for mission-critical applications, as well as outstanding storage deduplication and network compression capabilities. With NetApp SnapMirror technology, disaster recovery can protect the entire data center. Volumes can back up to an off-site location incrementally. SnapMirror performs incremental, block-based replication as frequently as the required RPO. The block-level updates reduce bandwidth and time requirements, and data consistency is maintained at the DR site.

An important step is to create a one-time baseline transfer of the entire dataset. This is required before incremental updates can be performed. This operation includes the creation of a Snapshot copy at the source and the transfer of all the data blocks referenced by it to the destination file system. After the initialization is complete, scheduled or manually triggered updates can occur. Each update transfers only the new and changed blocks from the source to the destination file system. This operation includes creating a Snapshot copy at the source volume, comparing it with the baseline copy, and transferring only the changed blocks to the destination volume. The new copy becomes the baseline copy for the next update. Because the replication is periodic, SnapMirror can consolidate the changed blocks and conserve network bandwidth. The impact on write throughput and write latency is minimal.

Recovery is performed by completing the following steps:

1. Connect to the storage system on the secondary site.
2. Break the SnapMirror relationship.
3. Map the LUNs in the SnapMirror volume to the initiator group (igroup) for the Hyper-V servers on the secondary site.
4. Once the LUNs are mapped to the Hyper-V cluster, make these disks online.
5. Using the failover-cluster PowerShell cmdlets, add the disks to available storage and convert them to CSVs.
6. Import the virtual machines in the CSV to the Hyper-V manager, make them highly available, and then add them to the cluster.
7. Turn on the VMs.

Deploying Microsoft Hyper-V on NetApp Storage: Conclusion

ONTAP is the optimal shared storage foundation to deploy a variety of IT workloads. ONTAP AFF or ASA platforms are both flexible and scalable for multiple use cases and applications. Windows Server 2022 and Hyper-V enabled on it is one common use case as the virtualization solution, which is described in this document. The flexibility and scalability of ONTAP storage and associated features enable customers to start out with a right-sized storage layer that can grow with and adapt to their evolving business requirements. In current market conditions, Hyper-V offers a perfect alternate hypervisor option which provides most of the functionalities that was provided VMware.

Deploying Microsoft Hyper-V on NetApp Storage: Migration Script

This section contains a PowerShell script that can be used for migration using Flexclone.

Powershell script

```
param (
    [Parameter(Mandatory=$True, HelpMessage="VCenter DNS name or IP Address")]
    [String]$VCENTER,
    [Parameter(Mandatory=$True, HelpMessage="NetApp ONTAP NFS Datastore name")]
    [String]$DATASTORE,
    [Parameter(Mandatory=$True, HelpMessage="VCenter credentials")]
    [System.Management.Automation.PSCredential]$VCENTER_CREDS,
    [Parameter(Mandatory=$True, HelpMessage="The IP Address of the ONTAP Cluster")]
    [String]$ONTAP_CLUSTER,
    [Parameter(Mandatory=$True, HelpMessage="NetApp ONTAP VServer/SVM name")]
    [String]$VSERVER,
    [Parameter(Mandatory=$True, HelpMessage="NetApp ONTAP NFS,SMB Volume name")]
    [String]$ONTAP_VOLUME_NAME,
    [Parameter(Mandatory=$True, HelpMessage="ONTAP NFS/CIFS Volume mount Drive on Hyper-V host")]
    [String]$ONTAP_NETWORK_SHARE_ADDRESS,
    [Parameter(Mandatory=$True, HelpMessage="NetApp ONTAP Volume QTree folder name")]
    [String]$VHDX_QTREE_NAME,
    [Parameter(Mandatory=$True, HelpMessage="The Credential to connect to the ONTAP Cluster")]
    [System.Management.Automation.PSCredential]$ONTAP_CREDS,
    [Parameter(Mandatory=$True, HelpMessage="Hyper-V VM switch name")]
    [String]$HYPERV_VM_SWITCH
)

function main {

    ConnectVCenter

    ConnectONTAP

    GetVMList
```

```

GetVMInfo

#PowerOffVMs

CreateOntapVolumeSnapshot

Shift

ConfigureVMsOnHyperV
}

function ConnectVCenter {
    Write-Host
    "-----"
    -----" -ForegroundColor Cyan
    Write-Host "Connecting to vCenter $VCENTER" -ForegroundColor Magenta
    Write-Host
    "-----"
    -----"`n" -ForegroundColor Cyan

    [string]$vmwareModuleName = "VMware.VimAutomation.Core"

    Write-Host "Importing VMware $vmwareModuleName Powershell module"
    if ((Get-Module|Select-Object -ExpandProperty Name) -notcontains
$vmwareModuleName) {
        Try {
            Import-Module $vmwareModuleName -ErrorAction Stop
            Write-Host "$vmwareModuleName imported successfully"
-ForegroundColor Green
        } Catch {
            Write-Error "Error: $vmwareMdouleName PowerShell module not
found"
            break;
        }
    }
    else {
        Write-Host "$vmwareModuleName Powershell module already imported"
-ForegroundColor Green
    }

    Write-Host "`nConnecting to vCenter $VCENTER"
    Try {
        $connect = Connect-VIServer -Server $VCENTER -Protocol https
-Credential $VCENTER_CREDS -ErrorAction Stop
        Write-Host "Connected to vCenter $VCENTER" -ForegroundColor Green
    } Catch {

```

```

        Write-Error "Failed to connect to vCenter $VCENTER. Error : $($_.Exception.Message)"
        break;
    }
}

function ConnectONTAP {
    Write-Host "`n
-----" -ForegroundColor Cyan
    Write-Host "Connecting to VSerevr $VSERVER at ONTAP Cluster
$ONTAP_CLUSTER" -ForegroundColor Magenta
    Write-Host
    "-----"
    ----`n" -ForegroundColor Cyan

    [string]$ontapModuleName = "NetApp.ONTAP"

    Write-Host "Importing NetApp ONTAP $ontapModuleName Powershell module"
    if ((Get-Module|Select-Object -ExpandProperty Name) -notcontains
$ontapModuleName) {
        Try {
            Import-Module $ontapModuleName -ErrorAction Stop
            Write-Host "$ontapModuleName imported successfully"
            -ForegroundColor Green
        } Catch {
            Write-Error "Error: $vmwareModule Name PowerShell module not
found"
            break;
        }
    }
    else {
        Write-Host "$ontapModuleName Powershell module already imported"
        -ForegroundColor Green
    }

    Write-Host "`nConnecting to ONTAP Cluster $ONTAP_CLUSTER"
    Try {
        $connect = Connect-NcController -Name $ONTAP_CLUSTER -Credential
$ONTAP_CREDS -Vserver $VSERVER
        Write-Host "Connected to ONTAP Cluster $ONTAP_CLUSTER"
        -ForegroundColor Green
    } Catch {
        Write-Error "Failed to connect to ONTAP Cluster $ONTAP_CLUSTER.
Error : $($_.Exception.Message)"
        break;
    }
}

```



```

    }
}

function GetVMList {
    Write-Host "`n
-----" -ForegroundColor Cyan
    Write-Host "Fetching powered on VMs list with Datastore $DATASTORE"
    -ForegroundColor Magenta
    Write-Host
    "-----"
    ----`n" -ForegroundColor Cyan
    try {
        $vmList = VMware.VimAutomation.Core\Get-VM -Datastore $DATASTORE
        -ErrorAction Stop| Where-Object {$_.PowerState -eq "PoweredOn"} | OUT-
        GridView -OutputMode Multiple
        # $vmList = Get-VM -Datastore $DATASTORE -ErrorAction Stop| Where-
        Object {$_.PowerState -eq "PoweredOn"}

        if($vmList) {
            Write-Host "Selected VMs for Shift" -ForegroundColor Green
            $vmList | Format-Table -Property Name
            $Script:VMList = $vmList
        }
        else {
            Throw "No VMs selected"
        }
    }
    catch {
        Write-Error "Failed to get VM List. Error : $($_.Exception.
Message)"
        Break;
    }
}

function GetVMInfo {
    Write-Host
    "-----"
    ----" -ForegroundColor Cyan
    Write-Host "VM Information" -ForegroundColor Magenta
    Write-Host
    "-----"
    ----" -ForegroundColor Cyan
    $vmObjArray = New-Object System.Collections.ArrayList

    if($VMList) {

```

```

foreach($vm in $VMList) {
    $vmObj = New-Object -TypeName System.Object

    $vmObj | Add-Member -MemberType NoteProperty -Name ID -Value
$vm.Id
    $vmObj | Add-Member -MemberType NoteProperty -Name Name -Value
$vm.Name
    $vmObj | Add-Member -MemberType NoteProperty -Name NumCpu
-Value $vm.NumCpu
    $vmObj | Add-Member -MemberType NoteProperty -Name MemoryGB
-Value $vm.MemoryGB
    $vmObj | Add-Member -MemberType NoteProperty -Name Firmware
-Value $vm.ExtensionData.Config.Firmware

    $vmDiskInfo = $vm | VMware.VimAutomation.Core\Get-HardDisk

    $vmDiskArray = New-Object System.Collections.ArrayList
    foreach($disk in $vmDiskInfo) {
        $diskObj = New-Object -TypeName System.Object

        $diskObj | Add-Member -MemberType NoteProperty -Name Name
-Value $disk.Name

        $fileName = $disk.FileName
        if ($fileName -match '\[(.*?)\]') {
            $dataStoreName = $Matches[1]
        }

        $parts = $fileName -split " "
        $pathParts = $parts[1] -split "/"
        $folderName = $pathParts[0]
        $fileName = $pathParts[1]

        $diskObj | Add-Member -MemberType NoteProperty -Name
DataStore -Value $dataStoreName
        $diskObj | Add-Member -MemberType NoteProperty -Name
Folder -Value $folderName
        $diskObj | Add-Member -MemberType NoteProperty -Name
Filename -Value $fileName
        $diskObj | Add-Member -MemberType NoteProperty -Name
CapacityGB -Value $disk.CapacityGB

        $null = $vmDiskArray.Add($diskObj)
    }

    $vmObj | Add-Member -MemberType NoteProperty -Name

```

```

PrimaryHardDisk -Value "[ $($vmDiskArray[0].DataStore) ] $($vmDiskArray[0]
.Folder)/ $($vmDiskArray[0].Filename) "
    $vmObj | Add-Member -MemberType NoteProperty -Name HardDisks
-Value $vmDiskArray

    $null = $vmObjArray.Add($vmObj)

    $vmNetworkArray = New-Object System.Collections.ArrayList

    $vm |
    ForEach-Object {
        $VM = $_
        $VM | VMware.VimAutomation.Core\Get-VMGuest | Select-Object
-ExpandProperty Nics |
        ForEach-Object {
            $Nic = $_
            foreach ($IP in $Nic.IPAddress)
            {
                if ($IP.Contains('.'))
                {
                    $networkObj = New-Object -TypeName System.Object

                    $vlanId = VMware.VimAutomation.Core\Get-
VirtualPortGroup | Where-Object {$_.Key -eq $Nic.NetworkName}
                    $networkObj | Add-Member -MemberType NoteProperty
-Name VlanID -Value $vlanId
                    $networkObj | Add-Member -MemberType NoteProperty
-Name IPv4Address -Value $IP

                    $null = $vmNetworkArray.Add($networkObj)
                }
            }
        }
    }

    $vmObj | Add-Member -MemberType NoteProperty -Name PrimaryIPv4
-Value $vmNetworkArray[0].IPv4Address
    $vmObj | Add-Member -MemberType NoteProperty -Name
PrimaryVlanID -Value $vmNetworkArray.VlanID
    $vmObj | Add-Member -MemberType NoteProperty -Name Networks
-Value $vmNetworkArray

    $guest = $vm.Guest
    $parts = $guest -split ":"
    $afterColon = $parts[1]

```

```

        $osFullName = $afterColon

        $vmObj | Add-Member -MemberType NoteProperty -Name OSFullName
-Value $osFullName
        $vmObj | Add-Member -MemberType NoteProperty -Name GuestID
-Value $vm.GuestId
    }
}

$vmObjArray | Format-Table -Property ID, Name, NumCpu, MemoryGB,
PrimaryHardDisk, PrimaryIPv4, PrimaryVlanID, GuestID, OSFullName, Firmware

$Script:VMObjList = $vmObjArray
}

function PowerOffVMs {
    Write-Host "`n"

    -----
    ----" -ForegroundColor Cyan
        Write-Host "Power Off VMs" -ForegroundColor Magenta
        Write-Host

    -----
    ----`n" -ForegroundColor Cyan
        foreach($vm in $VMObjList) {
            try {
                Write-Host "Powering Off VM $($vm.Name) in vCenter $($VCENTER
)"

                $null = VMware.VimAutomation.Core\Stop-VM -VM $vm.Name
-Confirm:$false -ErrorAction Stop
                Write-Host "Powered Off VM $($vm.Name)" -ForegroundColor Green
            }
            catch {
                Write-Error "Failed to Power Off VM $($vm.Name). Error :
$. _Exception.Message"
                Break;
            }
            Write-Host "`n"
        }
}

function CreateOntapVolumeSnapshot {
    Write-Host "`n"

    -----
    ----" -ForegroundColor Cyan
        Write-Host "Taking ONTAP Snapshot for Volume $ONTAP_VOLUME_NAME"
-ForegroundColor Magenta

```

```

Write-Host
"-----"
----`n" -ForegroundColor Cyan

Try {
    Write-Host "Taking snapshot for Volume $ONTAP_VOLUME_NAME"
    $timestamp = Get-Date -Format "yyyy-MM-dd_HH:mm:ss"
    $snapshot = New-NcSnapshot -VserverContext $VSERVER -Volume
$ONTAP_VOLUME_NAME -Snapshot "snap.script-$timestamp"

    if($snapshot) {
        Write-Host "Snapshot ""$($snapshot.Name)"" created for Volume
$ONTAP_VOLUME_NAME" -ForegroundColor Green
        $Script:OntapVolumeSnapshot = $snapshot
    }
} Catch {
    Write-Error "Failed to create snapshot for Volume
$ONTAP_VOLUME_NAME. Error : $_.Exception.Message"
    Break;
}

function Shift {
    Write-Host
"-----"
----" -ForegroundColor Cyan
    Write-Host "VM Shift" -ForegroundColor Magenta
    Write-Host
"-----"
----`n" -ForegroundColor Cyan

    $Script:HypervVMList = New-Object System.Collections.ArrayList
    foreach($vmObj in $VMObjList) {

        Write-Host "*****"
        Write-Host "Performing VM conversion for $($vmObj.Name)"
    -ForegroundColor Blue
        Write-Host "*****"

        $hypervVMObj = New-Object -TypeName System.Object

        $directoryName = "/vol/$( $ONTAP_VOLUME_NAME )/$( $VHDX_QTREE_NAME )
/$( $vmObj.HardDisks[0].Folder )"

        try {
            Write-Host "Creating Folder ""$directoryName"" for VM $(

```

```

$vmObj.Name) "
    $dir = New-NcDirectory -VserverContext $VSERVER -Path
$directoryName -Permission 0777 -Type directory -ErrorAction Stop
    if($dir) {
        Write-Host "Created folder ""$directoryName"" for VM
$($vmObj.Name) `n" -ForegroundColor Green
    }
}
catch {
    if($_.Exception.Message -eq "[500]: File exists") {
        Write-Warning "Folder ""$directoryName"" already exists!
`n"
    }
    Else {
        Write-Error "Failed to create folder ""$directoryName""
for VM $($vmObj.Name). Error : $($_.Exception.Message)"
        Break;
    }
}

$vmDiskArray = New-Object System.Collections.ArrayList

foreach($disk in $vmObj.HardDisks) {
    $vmDiskObj = New-Object -TypeName System.Object
    try {
        Write-Host "`nConverting $($disk.Name)"
        Write-Host "-----"

        $vmdkPath = "/vol/$(ONTAP_VOLUME_NAME)/$(disk.Folder)/
$(disk.FileName)"
        $fileName = $disk.FileName -replace '\.vmdk$', ''
        $vhdxPath = "$($directoryName)/$(fileName).vhdx"

        Write-Host "Converting ""$(disk.Name)"" VMDK path ""
$(vmdkPath)"" to VHDX at Path ""$(vhdxPath)"" for VM $($vmObj.Name)"
        $convert = ConvertTo-NcVhdx -SourceVmdk $vmdkPath
-DestinationVhdx $vhdxPath -SnapshotName $OntapVolumeSnapshot
-ErrorAction Stop -WarningAction SilentlyContinue
        if($convert) {
            Write-Host "Successfully converted VM ""$(vmObj.Name)
)"" VMDK path ""$(vmdkPath)"" to VHDX at Path ""$(vhdxPath)""
-ForegroundColor Green

            $vmDiskObj | Add-Member -MemberType NoteProperty -Name
Name -Value $disk.Name
            $vmDiskObj | Add-Member -MemberType NoteProperty -Name

```

```

VHDXPath -Value $vhdxPath

        $null = $vmDiskArray.Add($vmDiskObj)
    }
}
catch {
    Write-Error "Failed to convert ""${$disk.Name}"" VMDK to
VHDX for VM ${$vmObj.Name}. Error : ${$_.Exception.Message}"
    Break;
}
}

$hypervVMObj | Add-Member -MemberType NoteProperty -Name Name
-Value $vmObj.Name
$hypervVMObj | Add-Member -MemberType NoteProperty -Name HardDisks
-Value $vmDiskArray
$hypervVMObj | Add-Member -MemberType NoteProperty -Name MemoryGB
-Value $vmObj.MemoryGB
$hypervVMObj | Add-Member -MemberType NoteProperty -Name Firmware
-Value $vmObj.Firmware
$hypervVMObj | Add-Member -MemberType NoteProperty -Name GuestID
-Value $vmObj.GuestID

$null = $HypervVMList.Add($hypervVMObj)
Write-Host "`n"

}
}

function ConfigureVMsOnHyperV {
    Write-Host
    "-----"
    -----" -ForegroundColor Cyan
    Write-Host "Configuring VMs on Hyper-V" -ForegroundColor Magenta
    Write-Host
    "-----"
    -----`n" -ForegroundColor Cyan

    foreach($vm in $HypervVMList) {
        try {

            # Define the original path
            $originalPath = $vm.HardDisks[0].VHDXPath
            # Replace forward slashes with backslashes

```

```

$windowsPath = $originalPath -replace "/", "\"

# Replace the initial part of the path with the Windows drive
letter
$windowsPath = $windowsPath -replace "^\\vol\\", "\\
$(\$ONTAP_NETWORK_SHARE_ADDRESS)\\"

$vmGeneration = if ($vm.Firmware -eq "bios") {1} else {2};

Write-Host "*****"
Write-Host "Creating VM $($vm.Name)" -ForegroundColor Blue
Write-Host "*****"
Write-Host "Creating VM $($vm.Name) with Memory $($vm.
MemoryGB)GB, vSwitch $($HYPERV_VM_SWITCH), $($vm.HardDisks[0].Name) ""
$( $windowsPath)"" , Generation $($vmGeneration) on Hyper-V"

$createVM = Hyper-V\New-VM -Name $vm.Name -VHDPATH
$windowsPath -SwitchName $HYPERV_VM_SWITCH -MemoryStartupBytes (Invoke-
Expression "$($vm.MemoryGB)GB") -Generation $vmGeneration -ErrorAction
Stop

if($createVM) {
    Write-Host "VM $($createVM.Name) created on Hyper-V host
`n" -ForegroundColor Green

    $index = 0
    foreach($vmDisk in $vm.HardDisks) {
        $index++
        if ($index -eq 1) {
            continue
        }

        Write-Host "`nAttaching $($vmDisk.Name) for VM $($vm
.Name) "

        Write-Host
"-----"

        $originalPath = $vmDisk.VHDXPath

        # Replace forward slashes with backslashes
        $windowsPath = $originalPath -replace "/", "\"

        # Replace the initial part of the path with the
Windows drive letter
        $windowsPath = $windowsPath -replace "^\\vol\\", "\\
$(\$ONTAP_NETWORK_SHARE_ADDRESS)\\"

```



```

        try {
            $attachDisk = Hyper-v\Add-VMHardDiskDrive -VMName
$vm.Name -Path $windowsPath -ErrorAction Stop
            Write-Host "Attached $($vmDisk.Name) ""
$($windowsPath)"" to VM $($vm.Name)" -ForegroundColor Green
        }
        catch {
            Write-Error "Failed to attach $($vmDisk.Name)
$($windowsPath) to VM $($vm.Name): Error : $($_.Exception.Message) "
            Break;
        }
    }

    if($vmGeneration -eq 2 -and $vm.GuestID -like "*rhel*") {
        try {
            Write-Host "`nDisabling secure boot"
            Hyper-V\Set-VMFirmware -VMName $createVM.Name
-EnableSecureBoot Off -ErrorAction Stop
            Write-Host "Secure boot disabled" -ForegroundColor
Green
        }
        catch {
            Write-Error "Failed to disable secure boot for VM
$($createVM.Name). Error : $($_.Exception.Message) "
        }
    }

    try {
        Write-Host "`nStarting VM $($createVM.Name) "
        Hyper-v\Start-VM -Name $createVM.Name -ErrorAction
Stop
        Write-Host "Started VM $($createVM.Name)`n"
-ForegroundColor Green
    }
    catch {
        Write-Error "Failed to start VM $($createVM.Name).
Error : $($_.Exception.Message) "
        Break;
    }
}
}
catch {
    Write-Error "Failed to create VM $($vm.Name) on Hyper-V.
Error : $($_.Exception.Message) "
    Break;
}
}

```

```
}  
}
```

main

Copyright information

Copyright © 2024 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.