

Dell PowerStore: Microsoft Hyper-V Best Practices

White Paper

July 2022

H18069.5

White Paper

Abstract

This document provides best practices for configuring Microsoft Windows Server Hyper-V to perform optimally with Dell PowerStore storage.

Dell Technologies

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2020-2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners.

Published in the USA July 2022 H18069.5.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Executive summary	4
Introduction	5
Storage and transport best practices.....	11
Hyper-V best practices	14
PowerStore management and configuration best practices.....	36
Data protection	37
Disaster recovery with Hyper-V and PowerStore.....	47
Conclusion.....	49
References.....	50

Executive summary

Overview

This document provides best-practice guidance for deploying and optimizing the Microsoft Windows Server Hyper-V hypervisor role with Dell PowerStore storage. This paper focuses on the block (SAN) capabilities of single- or multiappliance PowerStore T model cluster configurations in support of Windows Server Hyper-V.

Hyper-V and PowerStore storage are feature-rich solutions. Hyper-V and PowerStore seamlessly integrate to offer a diverse range of configuration options that solve key business objectives such as storage capacity, performance, and resiliency.

Audience

This document is intended for IT administrators, storage architects, partners, and Dell Technologies employees. This audience also includes individuals who may evaluate, acquire, manage, operate, or design a Dell networked storage environment using PowerStore systems. Readers should have working knowledge of Dell PowerStore storage and Microsoft Hyper-V.

Revisions

Date	Description
April 2020	Initial release PowerStoreOS 1.0
July 2020	Minor updates
October 2020	Minor updates
April 2021	Update for PowerStoreOS 2.0
October 2021	Minor updates
October 2021	Template update
July 2022	Update for PowerStoreOS 3.0

We value your feedback

Dell Technologies and the author of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: Marty Glaser

Note: For links to other documentation for this topic, see the [PowerStore Info Hub](#) and [Dell Technologies Support](#).

Introduction

PowerStore overview

Dell PowerStore is a feature-rich, next-generation, enterprise storage appliance. PowerStore provides an ideal block-based storage platform for Microsoft Windows Server including the Hyper-V role. PowerStore provides powerful storage integrations, optimizations, and management tools that are well suited to support Microsoft environments.

PowerStore achieves new levels of operational simplicity and agility. It uses a container-based microservices architecture, advanced storage technologies, and integrated machine learning to unlock the power of your data. PowerStore is a versatile platform with a performance-centric design that delivers multidimensional scale, always-on data reduction, and support for next-generation media.

PowerStore brings the simplicity of public cloud to on-premises infrastructure, streamlining operations with an integrated machine-learning engine and seamless automation. It also offers predictive analytics to easily monitor, analyze, and troubleshoot the environment. PowerStore is highly adaptable, providing the flexibility to host specialized workloads directly on the appliance and modernize infrastructure without disruption. It also offers investment protection through flexible payment solutions and data-in-place upgrades.

PowerStore appliances

PowerStoreOS 1.x introduced the 1000, 3000, 5000, 7000, and 9000 models (PowerStore T and X). PowerStoreOS 2.x introduced the PowerStore 500 model (PowerStore T only).

PowerStoreOS 3.0 introduces the 1200, 3200, 5200, and 9200 model appliances as PowerStore T models. The 1200 through 9200 models incorporate a hardware refresh that offers improved CPU performance, more memory, and enhanced features. The different CPU, memory, drive, and IOPS capacities available with these models provide performance and price options to satisfy a wide variety of storage needs.



Figure 1. Dell PowerStore 2U storage appliance (shown with and without the front bezel)

PowerStore X models and PowerStore T models

Each PowerStore appliance comes in a 2U rackmount configuration. Depending on the model, PowerStore is configurable from the factory in two base-model configurations:

- **PowerStore X:** VMware ESXi-based hypervisor for storage and guest workloads (AppsON functionality)
- **PowerStore T:** Two configuration options are supported with PowerStore T:
 - Unified storage (block storage and file storage (NAS) on the same appliance)
 - Block-optimized (optimized for block storage only)

PowerStore X and PowerStore T models support VMware vSphere Virtual Volumes (vVols).

This paper focuses on configuring Hyper-V to run optimally with PowerStore T models.

PowerStore management tools

PowerStore Manager is an intuitive, all-inclusive HTML5-based management UI that is integrated on each PowerStore appliance.

- No client installation required
- No Java requirement
- Supported with most popular web browsers

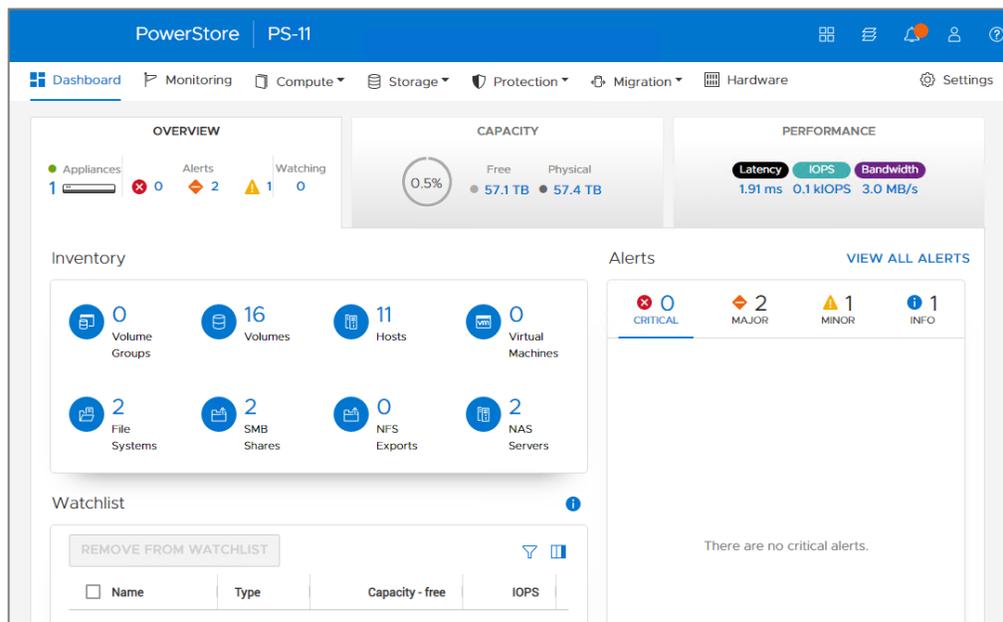


Figure 2. PowerStore Manager UI dashboard

Each PowerStore cluster provides an easy-to-use, web-based REST API interface for cluster management and automation tasks. To access the REST API interface, open a supported web browser, and add **/swaggerui** to the management IP of the cluster.

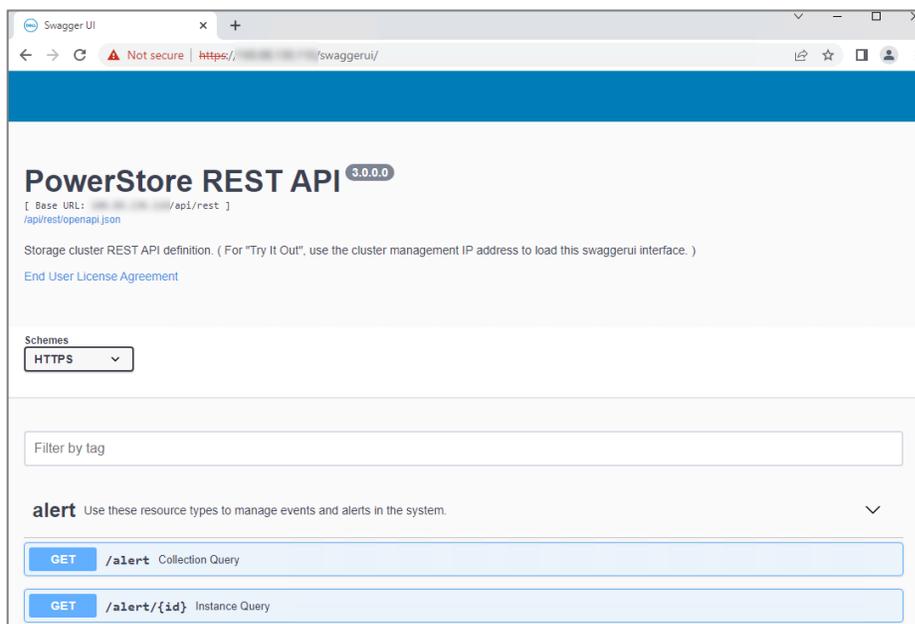


Figure 3. PowerStore REST API

PowerStore also provides a CLI as an extra management tool.

Documentation and support

To learn more about specific PowerStore models and features, see the [Dell Data Storage Portfolio](#).

PowerStore documentation is located at the [PowerStore Info Hub](#) and [Dell Support](#).

PowerStore product documentation and support are also available in PowerStore Manager. Click **Online Help** or **General Support**.

This white paper provides guidance that is supplemental to these resources.

Note: Most PowerStore storage features work seamlessly in the background, regardless of the platform or workload. Usually, the default storage settings for PowerStore are optimal for Hyper-V environments. This document provides configuration strategies and configuration options for PowerStore and Hyper-V that may enhance usability, performance, and resiliency in your environment.

Microsoft Hyper-V overview

Hyper-V is a mature, robust, proven virtualization platform. Hyper-V is a software layer that abstracts physical host server hardware resources. It presents these resources in an optimized and virtualized manner to guest virtual machines (VMs) and their workloads. Hyper-V optimizes the use of physical resources in a host server such as CPUs, memory, NICs, and power. Hyper-V virtualization allows many VMs to share physical host resources concurrently.

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.

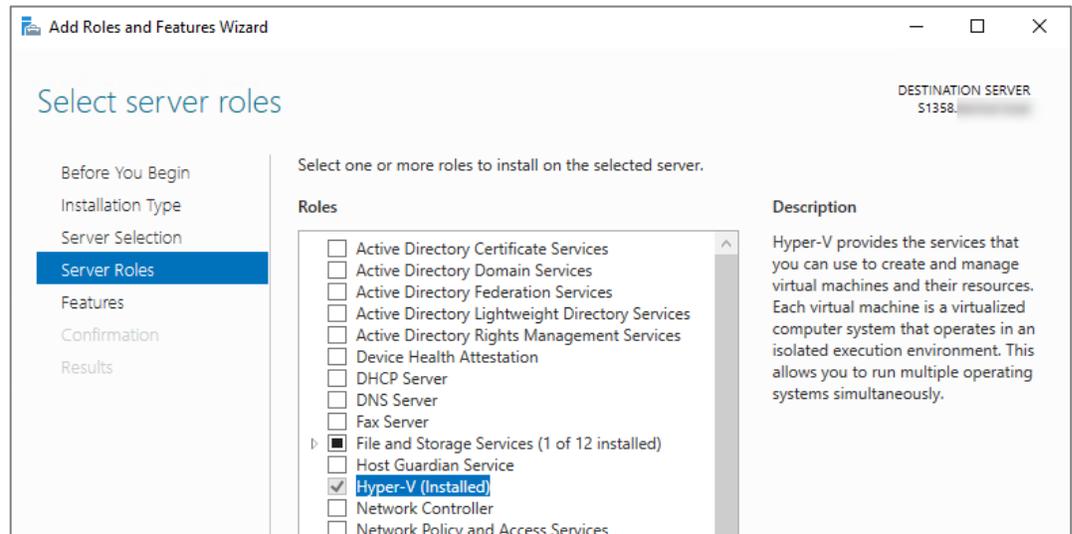


Figure 4. Server Manager > Add Roles and Features wizard

The Hyper-V role is not installed by default. Features that are typically installed to support the Hyper-V role include Multipath I/O and Failover Clustering.

To learn more about Hyper-V features, see the [Microsoft Virtualization Documentation library](#).

Hyper-V management tools

You can install and manage the Hyper-V role and other roles and features on a Windows Server using these methods:

- Windows Admin Center
- Windows Server Manager
- PowerShell

Windows Admin Center

Windows Admin Center is a free, centralized server-management tool from Microsoft. Windows Admin Center consolidates many common in-box and remote-management tools to simplify managing server environments and server clusters from one interface.

Windows Admin Center is a locally installed client that is HTML5-based and browser-accessible. Windows Admin Center is also an extensible platform allowing third parties to develop integrations for their own products or solutions.

PowerStore integration with Windows Admin Center is supported. For more information about Windows Admin Center integration with PowerStore, see the *Dell PowerStore Manager for Windows Admin Center User Guide* at [Dell Support](#).

Windows Admin Center is the recommended tool for managing Windows Server environments. However, it may not have full feature parity with the traditional management tools it replaces. Continue to use Hyper-V Manager, Failover Cluster Manager, Microsoft System Center Virtual Machine Manager (SCVMM), and PowerShell if the wanted functionality is not in Windows Admin Center.

This document includes configuration examples that use a combination of traditional tools and Windows Admin Center.

To learn more, see the [Microsoft Windows Admin Center](#) website.

For more information about Hyper-V features that are not specific to storage, see the [Microsoft Virtualization Documentation](#) library.

Supported versions

PowerStore support for different versions of Windows Server and the Hyper-V role may change over time. To verify Windows operating system and Hyper-V version compatibility with PowerStore for your environment, see the latest documentation and release notes for your version of PowerStoreOS at [Dell Support](#).

PowerStore is supported with long-term servicing channel (LTSC) releases of Windows Server. Use of semiannual channel (SAC) releases of Windows Server with PowerStore should be limited to nonproduction, test, or development use. To learn more about the differences between LTSC and SAC Windows Server versions, see this [Microsoft](#) article.

See [References](#) for a list of other resources.

Best practices overview

Best practices are derived over time from the collective experience of developers and end users. Best practices are built into the design of next-generation products. With mature technologies such as Hyper-V and Dell storage arrays, default settings and configurations typically incorporate the latest best practices.

As a result, tuning is often unnecessary and discouraged unless a specific design, situation, or workload will benefit from a different configuration. This document highlights situations where the default settings or configurations may not be optimal for Hyper-V.

Best practice design objectives commonly incorporate the following principles:

- Minimize complexity and administrative overhead
- Optimize performance
- Maximize security
- Ensure resiliency and recoverability
- Ensure a scalable design that can grow with the business
- Maximize return on investment over the life of the hardware

Best practices are baselines that may not be ideal for every environment. Some notable exceptions include the following examples:

- Legacy systems that are performing well and have not reached their life expectancy may not adhere to current best practice standards.
- A test or development environment that is not business critical may use a less-resilient design or lower-tier hardware to reduce cost and complexity.

Note: Following the best practices in this document are recommended. However, some recommendations may not apply to all environments. If questions arise, contact your Dell Technologies representative.

Terminology

The following list includes common terms that are used with PowerStore.

Table 1. Terminology

Term	Definition
Appliance	An appliance is a solution containing a PowerStore base enclosure and any attached expansion enclosures. The size of an appliance could be only the base enclosure or the base enclosure plus expansion enclosures.
Base enclosure	A 2U rackmount chassis containing the two active/active processing nodes (node A and node B) in the back, and 25 NVMe drive slots in the front.
Cluster	A PowerStore appliance is referred to as a cluster. Up to four appliances can be clustered together.
Expansion enclosure	A 2U rackmount chassis that can be cabled to a base enclosure to provide an additional 25 drive slots for storage expansion. Expansion is supported with all PowerStore models except the PowerStore 500.
Node	The component within a base enclosure that contains processors and memory. Each appliance consists of two hot-swappable nodes. Each node has an embedded module that provides management access, front-end connectivity, and ports for storage expansion.
PowerStoreOS	The PowerStore operating system.
PowerStore Manager	PowerStore Manager is an HTML5-based UI client for PowerStore configuration, management, and monitoring.
PowerStore T model	A unified (file and block) or block-optimized (block only) storage configuration option.
PowerStore X model	A VMware-based hypervisor configuration option (AppsON functionality).

Storage and transport best practices

Essential documentation

The following documents provide essential guidance for the planning, configuration, and deployment of PowerStore. Administrators should review and follow the guidance in these documents to ensure a successful deployment of Windows Server and Hyper-V on PowerStore:

- Dell PowerStore Release Notes
- Dell PowerStore Drive and OS Support Matrix
- Dell PowerStore Host Configuration Guide
- Dell PowerStore Best Practices Guide
- Dell PowerStore Installation and Service Guide
- Dell PowerStore Hardware Information Guide
- Dell PowerStore Planning Guide
- Dell PowerStore Networking Guide for PowerStore T Models

PowerStore documentation is located at [Dell Support](#) and the [Dell PowerStore Info hub](#).

After PowerStore is deployed, online help and general support documentation are available in PowerStore Manager.

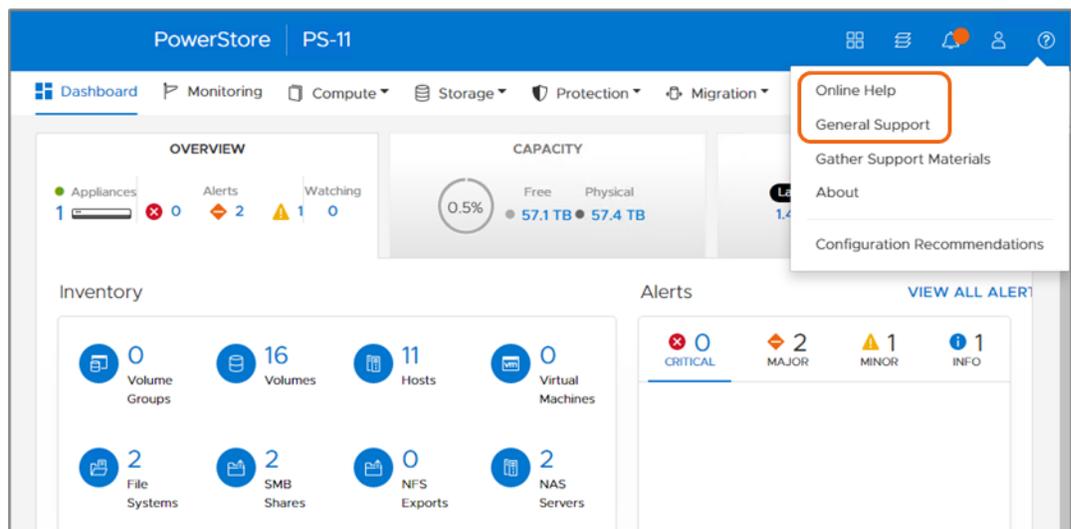


Figure 5. Help and support options in PowerStore Manager

This white paper provides supplemental guidance and best practices.

Right-size the PowerStore storage array

Before deploying PowerStore, consider the environmental design factors that impact storage capacity and performance. This planning ensures that new or expanded storage is right-sized for the Hyper-V environment. If PowerStore is deployed to support an existing Hyper-V workload, metrics such as storage capacity, bandwidth, and IOPS might already be understood. If the environment is new, these factors must be determined to correctly size the storage array, the storage fabric, and workload hosts.

Many common short- and long-term problems can be avoided by making sure the storage part of the solution will provide the right capacity and performance now and in the future. Scalability is a key design consideration.

Work with your Dell Technologies representative to complete a performance evaluation if there are questions about right-sizing a PowerStore storage solution for your environment and workload.

Disk capacity and performance

Total disk capacity does not always translate to disk performance. For example, installing a few large-capacity disks in a storage array may provide significant storage capacity, but may not support a high-IOPS workload.

Administrators must plan for IOPS and capacity when sizing PowerStore for Hyper-V or any other workload.

Avoid bottlenecks

Optimizing performance is a process of identifying and mitigating design limitations that cause bottlenecks. A bottleneck occurs when performance or functionality is negatively impacted under load because a capacity threshold is reached somewhere within the overall design. The goal is to maintain a balanced configuration end-to-end that allows the workload to operate at or near peak efficiency. The following design elements are potential bottlenecks:

- Storage performance (bandwidth; read and write I/O)
- Storage capacity
- Storage CPU and memory capacity
- Host server compute, memory, and bandwidth capacity
- Network and fabric bandwidth, throughput, and latency

Initial storage configuration

PowerStore T models support a unified configuration (file storage and block storage). If file services are not needed, choose **block-optimized** during the initial configuration of the hardware. When a PowerStore T model appliance is configured as unified, some compute and storage resources are reserved for NAS services, even if NAS is not used.

Note: If file services (NAS) may be needed later, select **Unified** during initial configuration. After an appliance is configured to support unified, the appliance must be reinitialized to change the configuration.

All PowerStore models support multipath I/O (MPIO) front-end connectivity to hosts. More information about MPIO is covered later in this document.

For more information about PowerStore host connectivity for Windows Server (including the Hyper-V role), see the *Dell PowerStore Host Configuration Guide* at [Dell Technologies Support](#).

Transport and front-end connectivity

PowerStore provides block storage to host servers using a storage area network (SAN).

PowerStore supports 16/32 Gb FC, and 10/25 GbE iSCSI. NVMe over Fibre Channel (NVMe/FC) and NVMe over TCP (NVMe/TCP) are two additional transports offered with PowerStore. NVMe/FC and NVMe/TCP are unsupported by Microsoft.

Note: A good understanding of the Hyper-V workload is essential for sizing the storage fabric correctly. PowerStore will not perform optimally if the storage fabric is inadequate for the workload.

Consider the following best practice recommendations for Hyper-V:

- Regardless of the transport used for Windows Server Hyper-V hosts, configure at least two paths to each server for redundancy in production environments.
 - Configure MPIO on each host in the environment.
 - Configuring hosts to use a single path (no MPIO) may be acceptable in test or development environments that are not business critical.

Administrators may continue using their preferred transport to maximize the return on their hardware investment, or switch to a different transport. The choice of transport is often based on personal preference or familiarity.

For more information, refer to the *Dell PowerStore Host Configuration Guide* at [Dell Technologies Support](#).

Hyper-V best practices

Introduction

PowerStore is an excellent choice as external storage for stand-alone or clustered Windows Servers including servers that are configured with the Hyper-V role. Core PowerStore features such as thin provisioning, data reduction, snapshots, and replication work seamlessly in the background regardless of the platform or operating system. Usually, the default settings for these features are optimal for Windows Server and Hyper-V. This section provides guidance for applying Hyper-V best practices for PowerStore.

General best practices for Hyper-V

General best practices for Hyper-V (not specific to PowerStore storage) are discussed in detail in Microsoft documentation.

Go to docs.microsoft.com and search for Hyper-V to view a list of technical documentation including the following:

- [Performance Tuning Hyper-V Servers](#)
- [Hyper-V Storage I/O Performance](#)
- [Hyper-V Network I/O Performance](#)
- [Detecting bottlenecks in a virtualized environment](#)

For more information about general best practices and tuning steps for Hyper-V, see the [Microsoft Windows Server Documentation](#) library.

To avoid redundancy, the general guidance in the documentation above is not duplicated in this document. This document assumes that administrators will deploy and tune Hyper-V in accordance with established Microsoft best practices.

General best practices that are common with any Hyper-V deployment include the following recommendations:

- Understand the I/O requirements of the workload before deploying it on Hyper-V.
 - Ensure the solution is adequately sized end-to-end to avoid bottlenecks.
 - Allow headroom for expansion that factors in anticipated growth.
- Keep the design simple to ease administrative overhead.
 - Adopt a standard naming convention for hosts, volumes, initiators, and so on. Consistent and intuitive naming makes administration easier.
- Configure all production hosts to use at least two data paths (MPIO) to eliminate single points of failure.
 - Use of single-path I/O may be acceptable in test or development environments that are not business critical.
- Use Windows Server Core to minimize the attack surface of a server and reduce administrative overhead.
- Use Windows Admin Center (for small deployments) or System Center Virtual Machine Manager (for large deployments) to centrally manage hosts and clusters.
- Ensure that all hosts and VMs are updated regularly.

- Provide adequate malware protection.
- Ensure that essential data is protected with backups that meet recovery time objectives (RTO) and recovery point objectives (RPO).
 - Snapshots and replication are integral to a data protection strategy with PowerStore.
- Minimize or disable unnecessary hardware devices and services to free up host resources for VMs. This action also helps to reduce power consumption.
- Schedule tasks such as periodic maintenance, backups, malware scans, and updates to run after hours. Stagger start times if maintenance operations overlap and are resource-intensive.
- Tune application workloads according to vendor recommendations to reduce or eliminate unnecessary processes or activity.
- Use PowerShell or other scripting tools to automate step-intensive, repeatable tasks to ensure consistency and avoid mistakes due to human error. This practice can also help reduce administration time.
 - PowerStore offers CLI and REST API support for additional management and scripting functionality.
- Enable monitoring and alerting features to identify and resolve issues quickly.
 - Configure email alerts.
 - Enable Dell SupportAssist in PowerStore to automatically contact support resources when events such as a drive failure occur.

Cluster validation

Run cluster validation before creating a Hyper-V cluster on PowerStore. All tests related to storage and MPIO should pass before configuring a Hyper-V cluster and deploying a workload.

1. Stage each Windows Server and configure the Hyper-V role according to Microsoft best practices.
2. Configure two or more data paths to PowerStore for each host (iSCSI or FC).
3. Install and configure MPIO on each host.
4. Use PowerStore Manager to create a host group on PowerStore.
5. Use PowerStore Manager to map at least one cluster volume to the host group using a consistent LUN ID.
6. On a host, initialize the new disk, bring it online, and format it.
7. Perform a disk rescan on each host in the host group.
8. Use Failover Cluster Manager to run cluster validation for the hosts in the host group.
9. Verify that all tests related to disk and MPIO pass.
10. If any tests fail, the configuration may not support clustering. Troubleshoot and resolve all disk or MPIO failures and run cluster validation again until they pass.

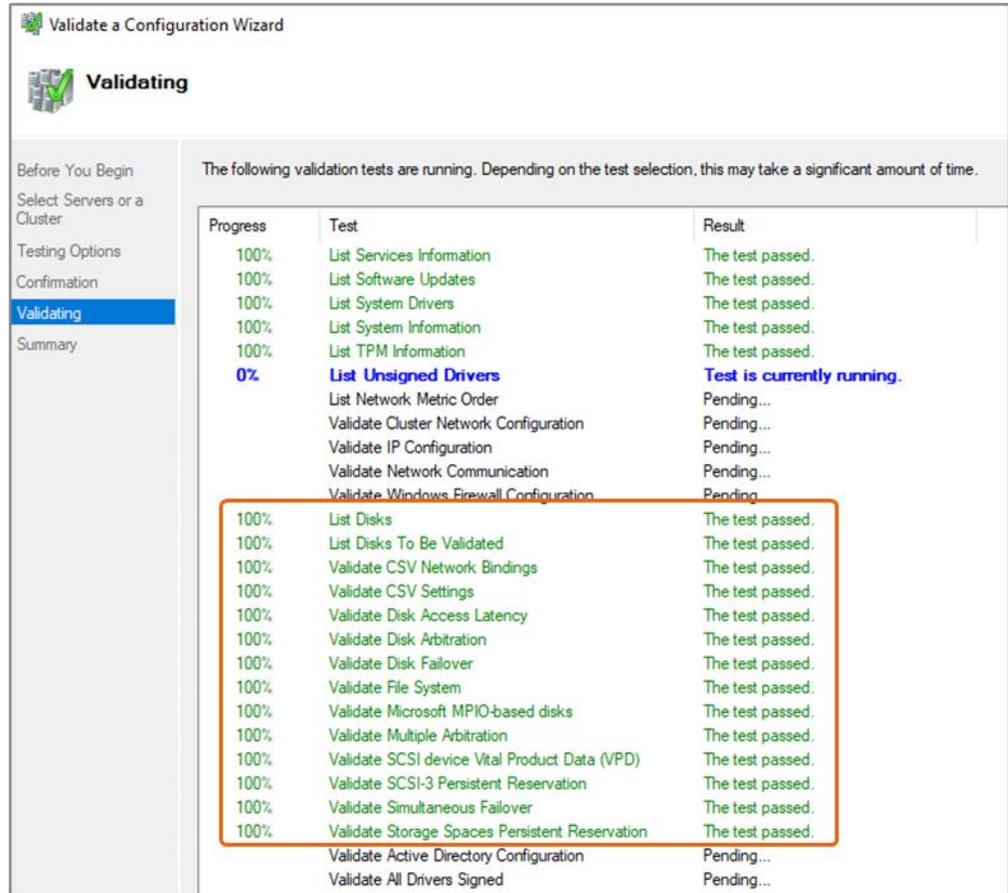


Figure 6. Failover Cluster Manager > Cluster validation wizard: Ensure that all disk and MPIO tests pass

Minor warnings will not prevent hosts from being clustered. For example, cluster validation may detect slight differences in the patch level of fully updated hosts and generate a warning. Even if updates are current on all tested hosts, this test may detect minor differences.

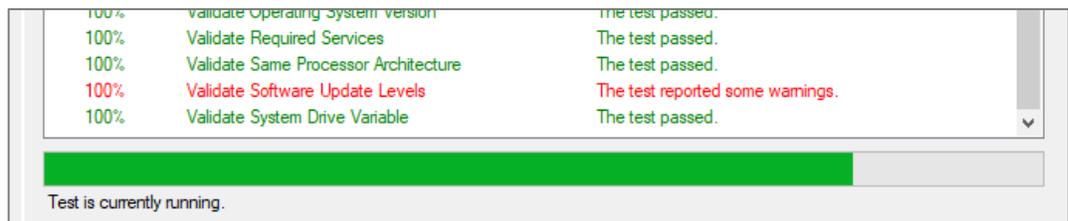


Figure 7. Cluster validation verifies software update levels

You can also use Windows Admin Center, in addition to Failover Cluster Manager, to create and manage clusters.

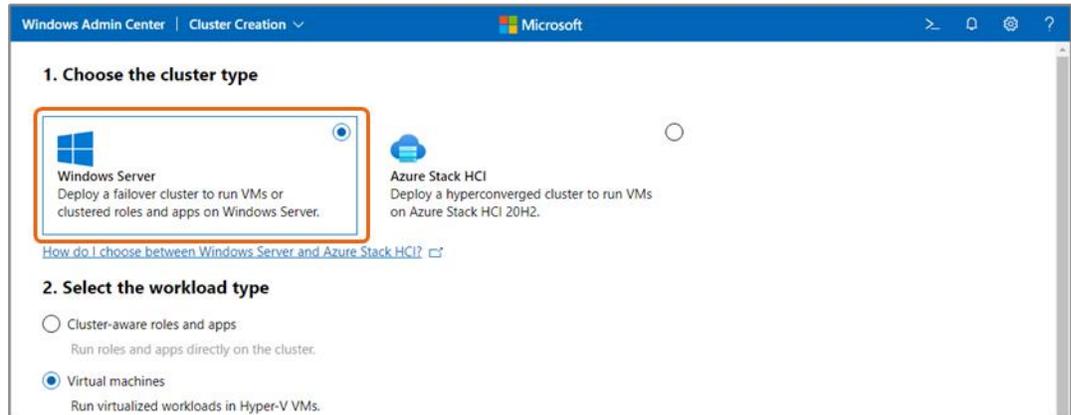


Figure 8. Windows Admin Center cluster creation tools

Guest VM integration services

Guest VM integration services are a package of virtualization-aware drivers that are installed on a guest VM. Integration services optimize the guest VM virtual hardware for interaction with the physical host hardware and with external storage.

Starting with the release of Windows Server 2016, VM integration services are installed automatically as a part of Microsoft updates.

If you have earlier versions of Hyper-V in your environment, you must install and manually update integration services on VMs. In Hyper-V Manager, use the **Action** menu to mount the Integration Services Setup Disk (an ISO file). Follow the prompts in the guest VM console to complete the installation.

Mounting the integration services ISO is not supported with Windows Server 2016 Hyper-V and newer. With newer versions of Hyper-V, integration services are provided exclusively as part of Microsoft updates.

When moving a VM from an older version of Hyper-V to a newer version, verify that the integration services get updated on the VM.

If a VM is not performing as expected (due to CPU, disk I/O, or network performance), verify that the VM integration services are current for the VM.

The presence of unknown devices on a VM may indicate that integration services are not installed or are outdated.

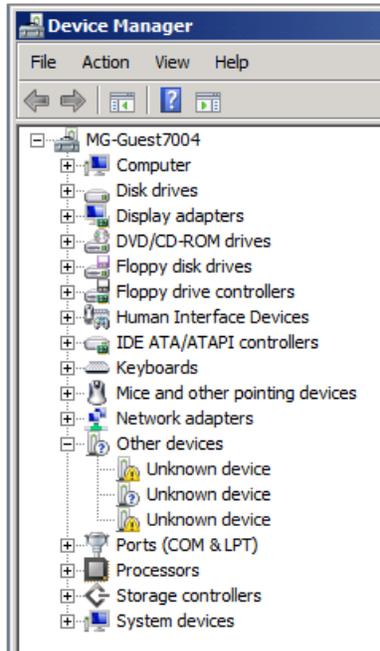


Figure 9. Unknown guest VM devices indicate missing or outdated integration services

Use tools such as Failover Cluster Manager, PowerShell, or Windows Admin Center to verify the version of integration services.

Hyper-V guest VM generations

Windows Server 2012 R2 Hyper-V introduced generation 2 VMs. When generation 2 VMs were introduced, existing VMs were designated as generation 1 VMs.

Generation 2 VMs included many new enhancements, including the following:

- Use of Unified Extensible Firmware Interface (UEFI) when booting instead of a legacy BIOS. UEFI provides better security and better interoperability between the operating system and the hardware, which offers improved virtual driver support and performance.
- Generation 2 VMs eliminate the dependency on virtual IDE for the boot disk. Generation 1 VMs require a virtual IDE disk controller for the boot disk.
 - Generation 2 guests support virtual SCSI controllers for all disks.
 - Virtual IDE is not a supported option with generation 2 VMs.

Generation 1 VMs are still supported with Hyper-V 2016 and newer. The New Virtual Machine Wizard may default to generation 1. However, all new VMs should be created as generation 2 as a best practice, if the guest operating system will support it.

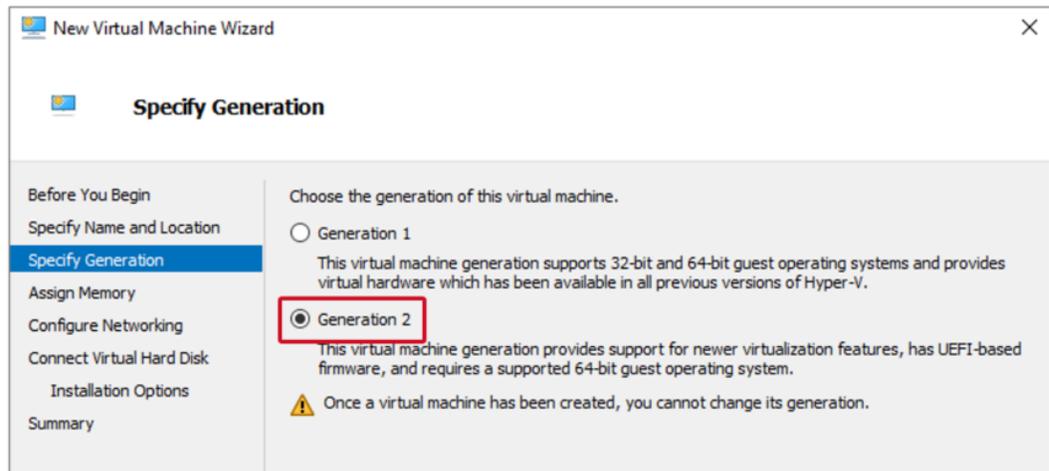


Figure 10. Guest VM generation option

For either generation of guest VM, if there are multiple disks requiring high I/O, each disk can be associated with its own virtual disk controller to maximize performance.

Convert VMs to a newer generation

You cannot change the VM generation after a VM is created. However, conversion may be possible using third-party tools (use at your own risk). The best-practice method is to migrate a workload to a generation 2 VM rather than attempting to convert a generation 1 VM to generation 2.

Virtual hard disks

A virtual hard disk is a set of data blocks that the host operating system stores as a regular Windows file with a VHD, VHDX, or VHDS extension. All virtual disk format types are supported with PowerStore storage.

Virtual hard disk format

There are three kinds of virtual hard disk formats that are supported with either VM generation:

- **VHD** is supported with all Hyper-V versions but is limited to a maximum size of two TB. VHD is a legacy format.
 - The New Virtual Hard Disk wizard may default to VHD with older versions of Hyper-V. However, use VHDX for new VM deployments when supported by the guest operating system.
- **VHDX** is supported with Windows Server 2012 Hyper-V and newer.
 - VHDX format is more resilient.
 - VHDX offers better performance and capacity up to 64 TB.
 - It is easy to convert a VHD to VHDX format using tools such as Hyper-V Manager or PowerShell.
- **VHDS** (or VHD Set) is supported on Windows Server 2016 Hyper-V and newer.
 - Two or more guest VMs can share access to a VHDS.

- Guest VMs can use VHDS disks as virtual cluster disks in highly available (HA) configurations.

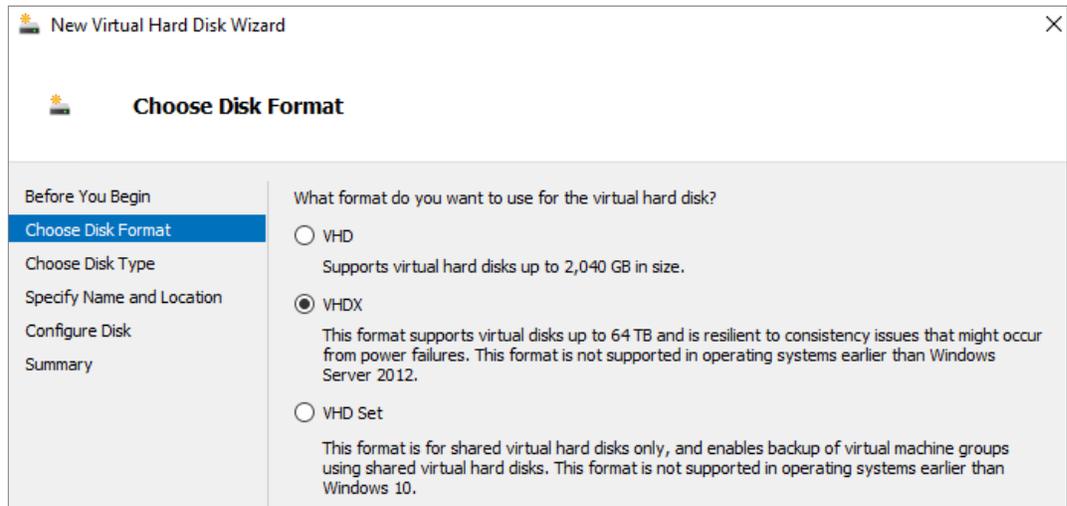


Figure 11. Virtual hard disk format options

Virtual hard disk type

Besides applying formatting options, you can designate a virtual hard disk as fixed, dynamically expanding, or differencing.

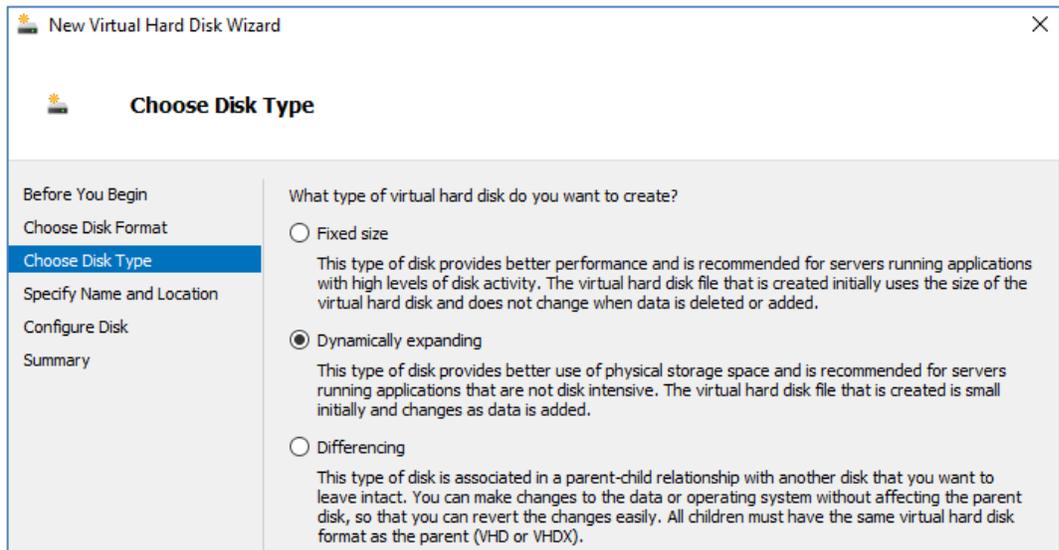


Figure 12. Options for virtual disk type

A dynamically expanding disk is the default type and will work well for most Hyper-V workloads on PowerStore storage. With PowerStore, only new data consumes storage capacity, regardless of the disk type (fixed, dynamic, or differencing). As a result, determining the best disk type is a function of the workload as opposed to how it will impact storage utilization. For general workloads, the performance difference between fixed and dynamic will usually be negligible. For workloads generating high I/O, such as Microsoft SQL Server databases, Microsoft recommends using the fixed-size virtual hard disk type for optimal performance.

A fixed virtual hard disk consumes the full amount of space from the perspective of the host server. For a dynamic virtual hard disk, the space is consumed as the VM writes new data to the disk. Dynamic virtual hard disks are more space efficient from the perspective of the host. From the perspective of the guest VM, either type of virtual hard disk shown in the following figure will present the full-formatted size of 60 GB to the guest.

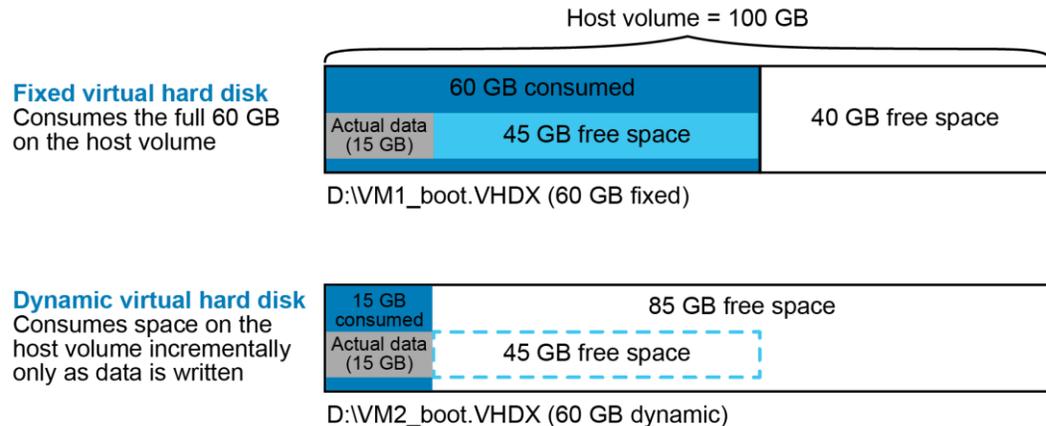


Figure 13. Fixed and dynamic virtual hard disk comparison

There are some performance and management best practices to consider when choosing a virtual hard disk type in your PowerStore storage environment.

- Fixed-size virtual hard disks:
 - Workloads or functions that generate high disk I/O experience better performance with fixed-size VHDs.
 - Fixed-size VHDs are less space efficient on the host server volume. For example, a 100 GB fixed-size VHD file consumes 100 GB on the host, even if the VHD contains no data.
 - Fixed-size VHDs are less susceptible to fragmentation.
 - Fixed-size VHDs take longer to copy to another location. The VHD file size is the same as the formatted size, even if the VHD contains no data.
- Dynamically expanding virtual hard disks:
 - Dynamic VHDs are recommended for most workloads, except for high disk I/O use cases.
 - Dynamic VHDs are space-efficient on the host, and the VHD file expands only as new data is written to it by the VM.
 - Dynamic VHDs are more susceptible to fragmentation at the host level.
 - A small amount of extra host CPU and I/O is required to expand a dynamic VHD file as it increases in size. Performance is not impacted unless the workload I/O demand is high.
 - Less time is required to copy a dynamic VHD file to another location. For example, if a 500 GB dynamically expanding VHD contains only 20 GB of data, the VHD file size when copied to another location is 20 GB.

- Dynamic VHDs allow the host disk space to be overprovisioned. Host disk space should be monitored closely. Configure alerting on the host server to avoid running volumes out of space when supporting dynamic VHDs.
- Differencing virtual hard disks:
 - Use cases are limited. For example, a virtual desktop infrastructure (VDI) deployment can use differencing VHDs.
 - You can realize storage savings with differencing VHDs by allowing multiple Hyper-V guest VMs with identical operating systems to share a common virtual boot disk.
 - All children must use the same virtual hard disk format as the parent.

Virtual hard disks and thin provisioning with PowerStore

Any virtual hard disk (fixed, dynamic, or differencing) will experience space-usage efficiency on PowerStore storage due to storage thin provisioning.

The example in the figure below shows a 100 GB volume presented to a Hyper-V host that contains two 60 GB virtual hard disks. Overprovisioning is shown in the example to demonstrate behavior, but not as a best practice. One disk is fixed, and the other is dynamic. Each virtual hard disk contains 15 GB of data. From the perspective of the host server, 75 GB of space is consumed and can be described as follows:

Example: 60 GB fixed disk + 15 GB of used space on the 60 GB dynamic disk = 75 GB total

Note: The host server will always report the formatted size as consumed for a fixed-size VHD.

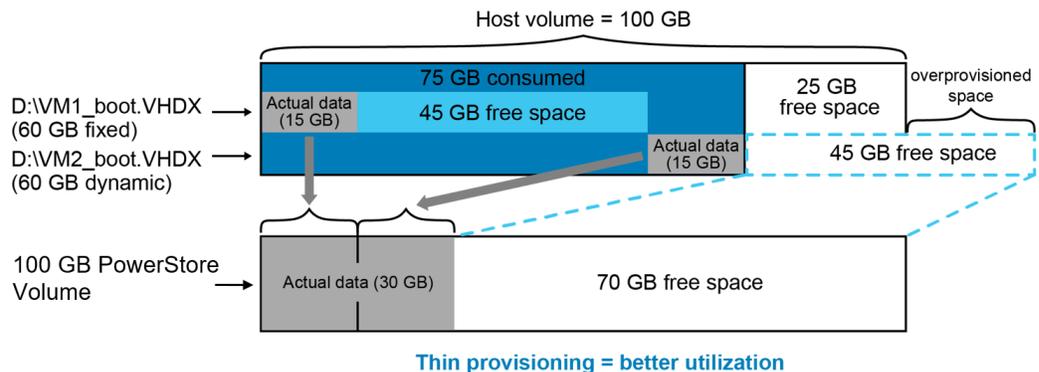


Figure 14. Thin provisioning with PowerStore

Comparatively, PowerStore reports storage utilization on this same volume as follows:

Example: 15 GB of used space on the fixed disk + 15 GB of used space on the dynamic disk = 30 GB

Note: Either type of virtual hard disk (dynamic and fixed) will consume the same space on PowerStore because of thin provisioning. Other factors such as the I/O demand of the workload would be primary considerations when determining the type of virtual hard disk in your environment.

Overprovisioning with dynamic virtual hard disks

With dynamic VHDs and thin provisioning, running the storage out of space is a concern if the storage is overprovisioned.

To mitigate risks, consider the following best practice recommendations:

- Create Hyper-V physical volumes that are large enough so that current and future expanding dynamic virtual hard disks will not fill the host volumes to capacity. Creating large Hyper-V physical volumes will not waste space on PowerStore because of thin provisioning.
 - If native Hyper-V checkpoints (snapshots) are used, allow adequate overhead on the physical volume for the extra space consumed by the snapshot data.
 - Expand existing physical volumes as needed to avoid the risks associated with overprovisioning.
 - Configure monitoring if a physical host volume with virtual hard disks is overprovisioned. For example, a percent-full threshold can generate a warning with enough lead time to allow for remediation.
- Monitor alerts on PowerStore so that warnings about disk group and pool capacity thresholds are remediated before they reach capacity.

Hyper-V Checkpoints

A native Hyper-V-based checkpoint creates a snapshot of a VM on the physical host volume or cluster volume.

Note: Native Hyper-V checkpoints (snapshots) are not the same as PowerStore storage snapshots. PowerStore array-based snapshots and native Hyper-V snapshots function independently.

Each extra Hyper-V checkpoint creates another new snapshot. They are stored in a hierarchical tree.

Hyper-V snapshots are mentioned here because of their impact on storage read performance. A long chain of Hyper-V checkpoints can degrade read performance. During a read operation, the requested blocks may reside in different checkpoints which can increase read latency enough to impact performance.

We recommend avoiding using Hyper-V checkpoints, or using them sparingly or temporarily.

Administrators can use PowerStore array-based snapshots to protect and replicate VM data, in addition to using native Hyper-V VM replication tools.

Present PowerStore storage to Hyper-V hosts and VMs

There are several ways to present PowerStore volumes as LUNs to Hyper-V hosts, nodes, and VMs.

Boot from SAN

PowerStore supports boot-from-SAN when hosts are configured with compatible adapters that also support boot-from-SAN. In Microsoft environments, boot-from-SAN is supported with stand-alone and clustered Hyper-V hosts and nodes. Guest VMs can also be

configured to boot from a pass-through disk directly from PowerStore. However, use cases for bootable pass-through disks are limited in favor of virtual hard disks (VHDs).

The following list includes some advantages of booting from a local disk or SAN disk.

Boot-from-SAN advantages:

- PowerStore snapshots of boot-from-SAN volumes provide for quick recovery for a host server.
- Boot-from-SAN volumes can be replicated to another PowerStore at a remote location for enhanced DR protection when both sites use similar hardware for server hosts.
- PowerStore supports the configuration of a base volume (as known as a gold image) that can be used as a source volume to create new servers. Thin clones of the base volume can be used to quickly provision new boot-from-SAN Hyper-V host servers.

Boot-from-local-disk advantages:

- Offline SAN or storage fabric maintenance does not affect the host boot disk.
- It allows critical roles such as AD domain controller services to stay online if the storage or storage fabric is offline.
- PowerStore host groups are supported. Hosts that are configured to boot-from-SAN are not supported as members of a PowerStore host group.

Present PowerStore boot volumes as LUN 0 to physical Hyper-V hosts or nodes that boot from SAN.

Data volumes

Present PowerStore volumes as data volumes to physical Hyper-V hosts, clusters, and VMs. This method supports the following:

- FC and iSCSI
- NTFS and ReFS disk formats
- Cluster volumes including cluster shared volumes (CSV)
- MPIO including the native Microsoft DSM
- PowerStore host groups and volume groups
- Direct-attached, in-guest iSCSI volumes for guest VMs
- Pass-through disks to VMs (for boot or data)

Note: NVMe over Fibre Channel (NVMe/FC) and TCP (NVMe/TCP) are not supported by Microsoft. Support will be extended to Windows environments as Windows drivers become available. Check the latest PowerStore documentation to verify Windows support for NVMe/FC or NVMe/TCP before choosing the NVMe initiator type.

See the *Dell PowerStore Administrator's Guide* and the *Dell PowerStore Deployment Guide* at [Dell Support](#) for an in-depth review of transports and cabling options.

Cluster shared volumes

PowerStore supports cluster shared volumes (CSV). The same principles and best practices that apply to the creation and mapping of other types of volumes apply to CSVs.

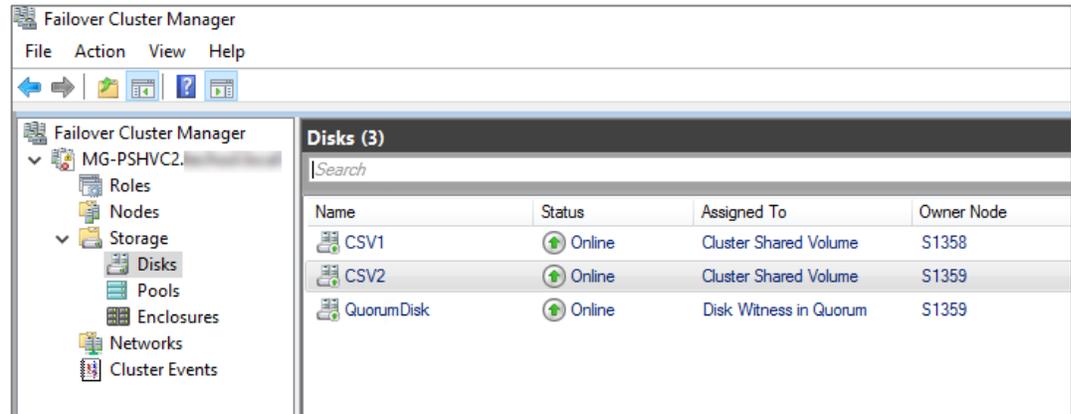


Figure 15. Cluster shared volumes

CSVs were introduced in Windows Server 2008 R2 Failover Clustering. A CSV is a volume type that allows all nodes in a Hyper-V cluster to have read/write access simultaneously.

With CSVs, clustered roles such as Hyper-V VMs can fail over quickly from one Hyper-V node to another. This action does not require a change in drive ownership or dismounting and remounting host volumes.

CSVs are most commonly used to support Hyper-V guest VMs. With each new release of Windows Server, Microsoft continues to expand the types of workloads that are supported on CSVs.

A CSV is initially formatted as an NTFS or ReFS volume.

Transport options

Deciding which transport to use is based on customer preference and factors such as the size of the environment, cost of the hardware, and the required support expertise.

iSCSI has grown in popularity for several reasons, such as improved performance with the higher bandwidth connectivity options now available. A converged Ethernet configuration also reduces complexity and cost. Small office, branch office, and edge use cases benefit when minimizing complexity and hardware footprints with converged networks.

Regardless of the transport, it is a best practice to ensure redundant paths to each host by configuring MPIO. For test or development environments that can accommodate down time without business impact, a less-costly, less-resilient design that uses single path may be acceptable.

Support for concurrent transports on the same host

In a Hyper-V environment on PowerStore, all hosts that are clustered should be configured to use a single common transport (FC or iSCSI).

There is limited Microsoft support for mixing transports on the same host. We do not recommend mixing transports as a best practice, but there are some uses cases for temporary use.

For example, when migrating from one transport type to another, both transports may be required to be available to a host during a transition period. If mixed transports must be used, use a single transport for each volume that is mapped to the host.

Consider the following example:

- A host has FC HBAs that support FC. The same host also has NICs that support iSCSI. The host is connected to both storage networks (FC and iSCSI) using MPIO.

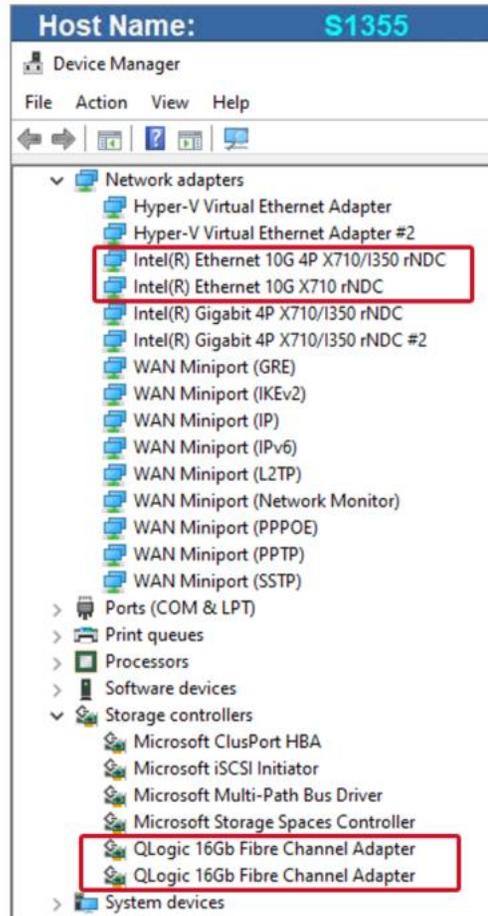


Figure 16. Host server with iSCSI NICs and FC HBAs

- An existing FC volume is mapped to the host from a legacy storage array that is being retired.

You would perform the following actions:

1. Create a host object on the new PowerStore array that uses iSCSI mappings.
2. Map a new volume on the PowerStore to the host using iSCSI. After discovery, the host will display two volumes:
 - a. The first volume is the FC volume from the legacy storage array.

- b. The new volume is the iSCSI volume from the PowerStore array.
3. Migrate the workload from the existing FC volume to the new iSCSI volume on the PowerStore array.
4. Discontinue the legacy FC volume.

Note: Do not attempt to map a volume to a Windows host using more than one transport. Mixing transports for the same volume will result in unpredictable service-affecting I/O behavior in path failure scenarios. Each volume should be mapped using a unique transport.

MPIO best practices

Windows Server and Hyper-V natively support MPIO. A Device Specific Module (DSM) provides MPIO support. The DSM that is bundled with the Windows Server operating system is fully supported with PowerStore arrays.

Windows and Hyper-V hosts default to the **Round Robin with Subset** policy with PowerStore. Round Robin with Subset will work well for most Hyper-V environments. Specify a different supported MPIO policy if wanted.

When each PowerStore storage controller (Node A and Node B) has four FC front-end (FE) paths connected to dual fabrics, eight paths are visible to the host server. Connecting fewer FE paths, such as two on each controller for four paths total, is also acceptable.

The **Active/Optimized** paths are associated with the PowerStore storage controller that the volume is assigned to. The **Active/Unoptimized** paths are associated with the secondary or standby PowerStore storage controller for that same volume.

When creating volumes on PowerStore, the wizard will alternate controller ownership in a round-robin fashion to help load balance the controllers. Administrators can override this behavior and specify a specific controller when creating a volume.

Best practices recommendations include the following:

- Do not change MPIO registry settings on the Windows or Hyper-V host (such as time-out values) unless directed by PowerStore documentation or Dell Support.
- Connect up to four FE ports on each PowerStore array node to use your preferred transport to optimize throughput and maximize performance. Connecting more than four ports per node to a fabric or storage network can result in performance degradation.
- Configure dual fabrics and storage networks for switch and path level redundancy.
- Configure each host to use at least two ports with a SAN configuration (iSCSI or FC). Configure host MPIO settings to protect against a controller or path failure.
- Verify that software versions are current for all components in the data path.
 - PowerStoreOS
 - Data and FC switch firmware
 - Boot code, firmware, and drivers for HBAs, NICs, SAS cards, and converged network adapters (CNAs)
- Verify that all hardware is supported according to the latest version of the *Dell PowerStore Support Matrix* at [Dell Support](#).

Guest VMs and block storage options

PowerStore block storage can also be presented directly to Hyper-V guest VMs using the following methods:

In-guest iSCSI: Configure the host and VM network so the VM can access PowerStore iSCSI volumes through a Hyper-V host or cluster network.

- Configure in-guest iSCSI on the VM. The setup is similar to iSCSI on a physical host.
- MPIO is supported on the VM if multiple paths are available to the VM, and the multipath I/O feature is installed and configured.

Physical disks: Physical disks presented to a Hyper-V VM are often referred to as pass-through disks. A pass-through disk is mapped to a Hyper-V host or cluster, and I/O access is passed through directly to a Hyper-V guest VM. The Hyper-V host or cluster has visibility to a pass-through disk and assigns it a LUN ID, but does not have I/O access. Hyper-V keeps the disk in a reserved state. Only the guest VM has I/O access.

- Use of pass-through disks is a legacy configuration that was introduced with Hyper-V 2008.
- Pass-through disks are no longer necessary because of the feature enhancements with newer releases of Hyper-V (generation 2 guest VMs, VHDX format, and shared VHDs).
- Use of pass-through disks is now discouraged, other than for temporary or specific use cases.

In-guest iSCSI and pass-through disk use cases

PowerStore arrays support in-guest iSCSI and pass-through disks (direct-attached disks) mapped to guest VMs. However, we do not recommend using direct-attached storage for guest VMs as a best practice unless a specific use case requires it. Typical use cases include:

- **Performance:** Direct-attached disks bypass the host server file system and so offer slightly better performance than a VHD or VHDX. There is no significant difference in performance between a direct-attached disk and a virtual hard disk for most workloads.
- **Clustering:** VM clustering on legacy Hyper-V platforms require the use of direct-attached disks. Shared VHDs are preferred for VM clustering with Server 2012 R2 and newer.
- **Troubleshooting:** Use of a direct-attached disk can be helpful if you must troubleshoot the I/O performance of a volume and isolate it from all other servers and workloads.
- **Custom snapshot or replication policy:** It may be necessary in some use cases to apply a custom PowerStore snapshot or replication policy to a specific disk (volume).
 - The preferred method is to place a virtual hard disk on a dedicated cluster shared volume (CSV) in a one-to-one configuration. Then, apply PowerStore snapshots and replication to the CSV.
- **Capacity:** Legacy VHDs support a maximum size of 2 TB. VHDX supports a maximum size of 64 TB. If a data volume will exceed these limits, you may be

required to use in-guest iSCSI or a pass-through disk. The maximum supported size of a direct-attached disk is a function of the VM operating system.

In-guest iSCSI and pass-through disk storage limitations

- **Native Hyper-V Snapshots:** The ability to perform native Hyper-V snapshots is lost. However, the ability to use PowerStore snapshots of the underlying volume is unaffected.
- **Complexity:** Use of direct-attached volumes increases complexity, requiring more management overhead.
- **Mobility:** VM mobility is reduced due to creating a physical hardware layer dependency.
- **Scale:** Each pass-through disk consumes a LUN ID on each host in a Hyper-V cluster. Extensive use of pass-through disks quickly becomes impractical and unmanageable at scale on a Hyper-V cluster. Use pass-through disks sparingly if they are required.
- **Differencing Disks:** The use of a pass-through disk as a boot volume on a guest VM prevents the use of a differencing disk.

Note: Legacy Hyper-V environments that are using direct-attached disks for guest VM clustering should consider switching to shared virtual hard disks when migrating to a newer Hyper-V version.

Virtual Fibre Channel

We do not recommend using virtual Fibre Channel (vFC) with PowerStore. If direct-attached storage is required for a Hyper-V VM, we recommend using in-guest iSCSI.

vFC requires a more complicated setup than in-guest iSCSI, and vFC can be difficult to configure and troubleshoot.

There is also limited support for a vFC configuration with PowerStore.

Note: A security parameter may be required to be disabled for a guest VM configured to use vFC adapters. See this [Microsoft KB](#) for more information.

To map WWNs to a new host (or a VM) in PowerStore Manager, the WWNs must be active. The Microsoft implementation of vFC assigns two virtual-WWN-address sets to each vFC adapter as shown in the following figure. Address set B is offline by default. The VM alternates between address set A and B when the VM fails over or live migrates. Only one set can be active at a time.

PowerStore Manager does not provide a way to map WWNs to a host or VM if the WWNs are offline. The vFC volume mapped to a VM can only be configured with half the WWNs (the active set but not the standby set). As a result, if the VM switches to the standby address set for any reason, the VM loses connectivity to the vFC volume. This occurrence causes a service interruption.

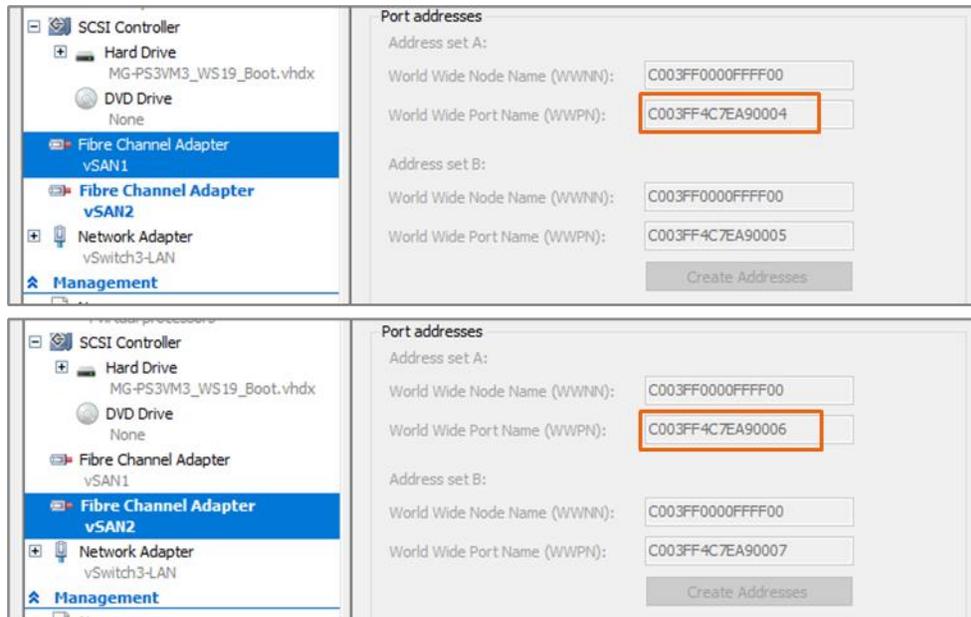


Figure 17. Hyper-V Manager vFC adapter configuration for a guest VM

In this example, the ports ending in 04 and 06 (set A) are active. The ports ending in 05 and 07 (set B) are inactive. When the VM moves to another host, the states reverse: ports 04 and 06 (set A) become inactive, and 05 and 07 (set B) become active.

There are limited use cases for vFC with PowerStore. These use cases require the administrator to acknowledge and work around the inability to assign both sets of vFC WWN addresses to a VM.

- We do not recommend this configuration for production use.
- We do not recommend this configuration for VMs on a Hyper-V cluster. Whenever a guest VM fails over or live-migrates to another physical node in a cluster, the active vFC WWN address set goes offline and the standby set becomes active.
- A VM using vFC on a stand-alone Hyper-V host (that does not have a failover or live-migrate destination) will typically always use address set A. This use applies even if the host or VM is rebooted. However, there is no guarantee that the VM will always use address set A, so the risk of a service interruption is always present.
- General use cases for vFC are similar to the use cases for in-guest iSCSI.
- MPIO is supported with vFC for the active paths that are present (address set A). Install and configure the MPIO feature on the guest VM.

PowerStore storage and Hyper-V clusters

Use a consistent LUN number when mapping shared volumes: quorum disks, cluster disks, and cluster shared volumes. Use host groups on the PowerStore array to simplify the task of assigning consistent LUN numbers.

Note: Hyper-V hosts that use boot-from-SAN cannot be added to PowerStore hosts groups.

Changing LUN IDs after initial assignment by PowerStore may be necessary to make them consistent. By default, PowerStore Manager assigns the next available LUN ID that is common when mapping a new volume to a host group or group of hosts.

Volume design considerations for PowerStore storage

Each cluster shared volume (CSV) supports one VM or many VMs. How many VMs to place on a CSV is a function of user preference, the workload, and how PowerStore storage features such as snapshots and replication will be used. Placing multiple VMs on a CSV is a good design starting point in most scenarios. Adjust this strategy for specific uses cases.

Some advantages for a many-to-one strategy include the following:

- **Avoid volume sprawl:** Fewer PowerStore array volumes are easier to manage.
- **Efficiency:** It is quick and easy to deploy a VM to an existing CSV.

Some advantages for a one-to-one strategy include the following:

- **I/O isolation:** It is easier to isolate and monitor disk I/O patterns for a specific Hyper-V guest VM or workload.
- **Ease of recovery:** It is easy to quickly restore a guest VM by recovering the underlying CSV using a PowerStore snapshot.
- **Replication control:** One-to-one gives administrators more granular control over what data gets replicated when PowerStore volumes are replicated to another location.
- **Move large VMs quickly:** Use of native Hyper-V tools to migrate VMs is preferred. However, for large VMs, it might be easier to move a guest VM from one host or cluster to another by remapping the volume. Remapping the CSV (or using a PowerStore snapshot) avoids having to copy or move a VM and its data over the network.

Other strategies include placing VHDs with a common purpose on a CSV. For example, place boot VHDs on a common CSV, and place data VHDs on other CSVs.

Optimize format disk wait time for large volumes

Formatting a PowerStore storage volume mapped to a Windows host should complete in a few seconds. If long format wait times are experienced for unusually large volumes, temporarily disable the file system **Delete Notify** attribute on the Windows host by completing the following steps:

1. Access a command prompt on the host server with elevated (administrator) rights.
2. To verify the state of the attribute, run the following command:

```
fsutil behavior query disabledeletenotify
```

3. A result of zero means the attribute is enabled. This attribute is configurable for NTFS and ReFS volumes.

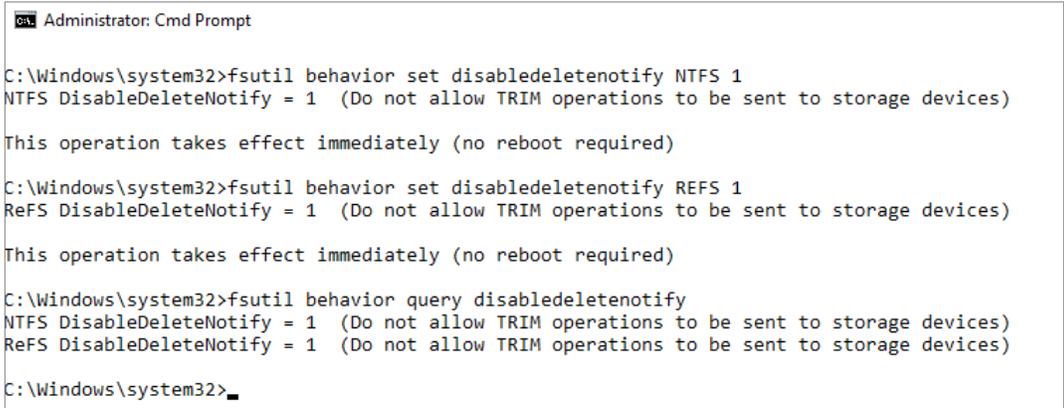
```
Administrator: Cmd Prompt
Microsoft Windows [Version 10.0.20348.558]
(c) Microsoft Corporation. All rights reserved.

C:\Windows\system32>fsutil behavior query disabledeletenotify
NTFS DisableDeleteNotify = 0 (Allows TRIM operations to be sent to the storage device)
ReFS DisableDeleteNotify = 0 (Allows TRIM operations to be sent to the storage device)

C:\Windows\system32>_
```

4. To disable the attribute, run the following commands:

```
fsutil behavior set disabledeletenotify NTFS 1
fsutil behavior set disabledeletenotify REFS 1
```



```
Administrator: Cmd Prompt

C:\Windows\system32>fsutil behavior set disabledeletenotify NTFS 1
NTFS DisableDeleteNotify = 1 (Do not allow TRIM operations to be sent to storage devices)

This operation takes effect immediately (no reboot required)

C:\Windows\system32>fsutil behavior set disabledeletenotify REFS 1
ReFS DisableDeleteNotify = 1 (Do not allow TRIM operations to be sent to storage devices)

This operation takes effect immediately (no reboot required)

C:\Windows\system32>fsutil behavior query disabledeletenotify
NTFS DisableDeleteNotify = 1 (Do not allow TRIM operations to be sent to storage devices)
ReFS DisableDeleteNotify = 1 (Do not allow TRIM operations to be sent to storage devices)

C:\Windows\system32>
```

5. When the volume is formatted, revert the setting.

Trim and unmap and for space recovery

When a file is deleted on a Windows Server, the file pointer is deleted. However, the old data remains on the disk. Over time, the operating system overwrites the old data with new data.

The host passes a trim and unmap command to external storage when files are deleted. Within a few minutes, the PowerStore storage pool reflects the additional free capacity.

The ability to recover deleted disk space on PowerStore is a key benefit of thinly provisioned volumes. In cases where trim and unmap is not supported or disabled, reclaimed space appears as free in Windows, but not on the storage.

Windows Server and Hyper-V support trim and unmap natively with PowerStore given these conditions:

- The Windows Server operating system must be version 2012 or newer.
- Volumes must be **basic** disks that are formatted as NTFS volumes. Trim and unmap is not supported with other formats such as ReFS.

ODX

Offloaded data transfer (ODX) reduces CPU and network utilization on a Microsoft Server by offloading a file-copy process from the host server to external storage. This feature is not supported with PowerStore.

For more information about ODX, see the [Microsoft Windows Server Documentation Library](#).

ReFS

The resilient file system (ReFS) was introduced with the initial release of Windows Server 2012. ReFS is a file system that is intended for managing large data volumes. ReFS uses a file-system design that autodetects data corruption and performs repairs without having to take the volume offline. ReFS eliminates the need to run time-consuming check-disk operations on large volumes to resolve disk errors or corruption. ReFS is supported with PowerStore, but trim and unmap is not supported with ReFS volumes (a Microsoft limitation).

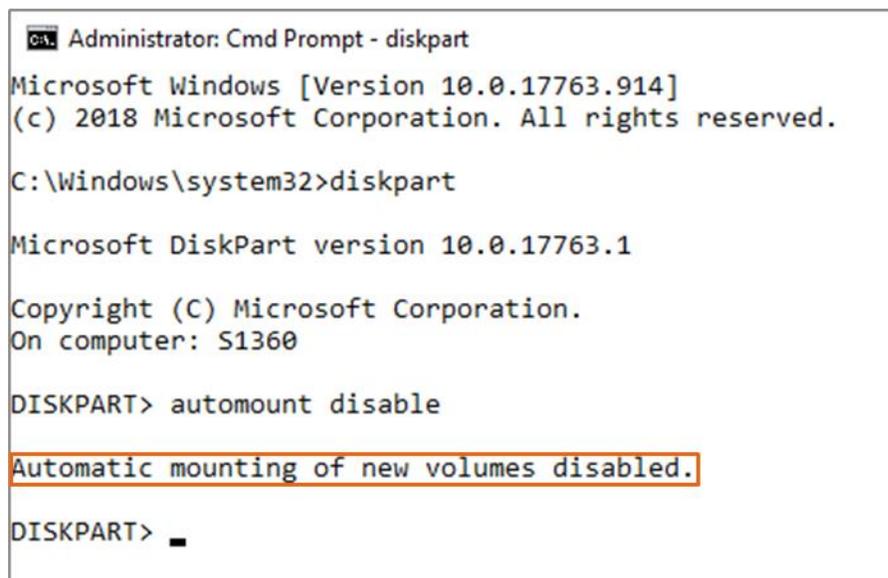
Although Microsoft recommends ReFS for large data volumes, compare feature sets for NTFS and ReFS before choosing ReFS. If trim and unmap support is needed for a volume, choose NTFS.

PowerStore features such as snapshots, thin clones, data reduction, replication, and others work equally well with NTFS or ReFS.

Disable automount

To prevent a Hyper-V host server from automatically assigning drive letters to newly mapped volumes, disable the automount feature. Automount is enabled by default. Disabling automount is beneficial in recovery scenarios when having the host automatically assign drive letters to previously formatted volumes is undesirable. These volumes could include a thin clone of an existing volume. Disabling automount can also be helpful when it interferes with a recovery process. We recommend disabling automount on Hyper-V servers in a PowerStore environment.

The following figure shows how to disable the automount feature by running **diskpart** from a command prompt with administrator privileges.



```
Administrator: Cmd Prompt - diskpart
Microsoft Windows [Version 10.0.17763.914]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Windows\system32>diskpart

Microsoft DiskPart version 10.0.17763.1

Copyright (C) Microsoft Corporation.
On computer: S1360

DISKPART> automount disable

Automatic mounting of new volumes disabled.

DISKPART> _
```

Figure 18. Disable the automount feature with diskpart

Placement of page files

Windows Servers and VMs place the page file on the boot volume by default. Windows automatically manages page file and memory settings. No user intervention is required to optimize memory management. Do not change the default settings unless required for a specific use case. For example, an application vendor may provide guidance for tuning page file and memory settings to optimize the performance of a specific workload.

With PowerStore storage, placing a page file on a separate VHD and separate CSV may provide some storage advantages. The following reasons may not be sufficiently advantageous to justify modifying the default settings. When a vendor recommends making page file changes to optimize a workload, consider the following tips as part of the overall page-file strategy.

- Move the page file to a separate dedicated volume or virtual hard disk to reduce the amount of data that is changing on the system (boot) volume. Moving the page file to a different volume reduces the size of PowerStore snapshots of boot volumes which will conserve PowerStore storage space.
- Volumes or virtual hard disks dedicated to page files usually do not require snapshot protection or replication to a remote site as part of a DR plan. Isolating page files reduces snapshot overhead and avoids replicating unnecessary data to a remote location.
- In a Hyper-V cluster environment, you may dedicate a CSV to virtual hard disks containing swap files.

Resiliency of essential services

Consider the following best practices to optimize the availability of essential services in your Hyper-V and PowerStore environment.

- Configure at least one domain controller as a physical host with local disk, or as a VM on a Hyper-V host with local disk.
- Run at least one domain controller independent of SAN storage so it will continue to provide essential services if external storage is unavailable. Essential services include AD user authentication, cluster authentication, DNS, and DHCP.
- Consider placing a management host or VM (jump box) in the environment that remains accessible regardless of the state of the storage fabric or SAN resources. Place critical management tools on this resource to aid with day-to-day administration, troubleshooting, and recovery.

Domain controller placement

Avoid placing all domain controller VMs on the same Hyper-V cluster. If the cluster service depends on AD authentication in order to start, an outage of the Hyper-V cluster will result in a recovery conundrum for the administrator. Recovery may require the following steps:

- Manually recover a domain controller VM outside of the Hyper-V cluster, and bring it online.
- With AD available, Hyper-V cluster services can now authenticate and start.
- Redesign the environment so at least one domain controller is not dependent on Hyper-V cluster services starting first.

Queue depth best practices for Hyper-V

Queue depth refers to the number of disk transactions that can be in flight from an initiator port (on a host server) to a target port (on the storage array). Host server FC and iSCSI adapters have queue depth settings that can be modified.

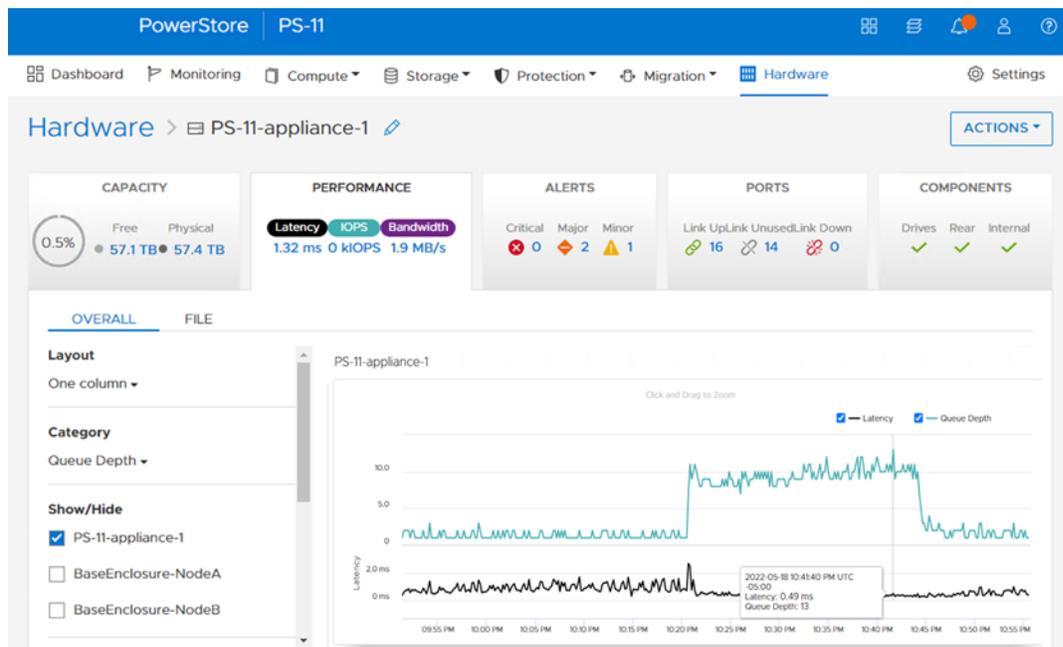


Figure 19. Monitor PowerStore queue depth with PowerStore Manager

A target port on PowerStore storage supports multiple host initiator ports sending it data concurrently. Initiator queue depth is used to limit the number of transactions an initiator can send to a target. Flooding occurs when a target port becomes saturated, and transactions are queued. Flooding causes higher latency and degraded performance for the affected workloads.

With PowerStore SAN configurations, configure up to four front-end data (target) ports per PowerStore node per fabric. Configuring more than four ports per PowerStore node per fabric may cause performance degradation. Configure two to four ports for an optimal configuration. Using multiple target ports allows I/O to be spread out, reducing the risk of port saturation.

When to change queue depth

On a Windows Server host, queue depth is a function of the Microsoft storport.sys driver and the vendor-specific miniport driver for the FC or iSCSI adapter. Default queue-depth settings provide a good starting point and are adequate for most workloads.

Note: We do not recommend modifying queue-depth settings unless there is a specific reason to do so. Test any queue-depth changes before applying them in a production environment.

Consider the following example:

- A storage array is connected to a small Hyper-V cluster consisting of a few nodes.
- The workload on this cluster is an I/O-intensive, large-block, sequential-read application.

- Increasing the queue depth settings may provide significant performance benefits for the workload on this small cluster.

However, consider the possible negative impact if many hosts are mapped to this storage array.

- Increasing host initiator queue depth may saturate the target ports on the PowerStore storage.
- All connected hosts may suffer a negative performance impact as a result.

Vendor-specific queue depth settings

See the documentation for your host adapter for information about adjusting queue depth settings.

For example, see the *Marvell QLogic Fibre Channel Adapters Users Guide* at Marvell.com.

PowerStore management and configuration best practices

PowerStore host groups

When mapping a PowerStore volume (quorum disk, cluster disk, or CSV) to multiple Hyper-V nodes, map the volume to nodes in the cluster using a consistent LUN number.

Use host groups in PowerStore to simplify the task of mapping a new volume to many nodes simultaneously. This practice can save significant time in larger environments and can help reduce the risk of user error. It also helps ensure that a consistent LUN ID is used for each Hyper-V node a volume is mapped to.

Use the **Add Host Group** wizard in PowerStore to create a host group. Only hosts that do not have any PowerStore volumes mapped can be added to a new or existing PowerStore host group.

Note: Hosts that are configured to use a PowerStore volume to boot-from-SAN are not eligible for membership in a PowerStore host group.

PowerStore volume groups

In a clustered environment, including Hyper-V, where a workload may span multiple CSVs, time consistency of snapshots and replication becomes vital in recovery scenarios. Use volume groups in PowerStore to ensure time consistency.

Run the **Create Volume Group** wizard in PowerStore to create a volume group. By default, the **Apply write-order consistency to protect all volume group members** is selected. Keep this default to ensure that PowerStore protection policies (snapshot rules and replication rules) apply to all volumes in the group simultaneously.

You can add new or existing volumes to a volume group. For a multiappliance PowerStore cluster, volumes added to a volume group must all reside on a common appliance.

If a volume is not on a common appliance, migrate the volume to the same appliance as the other volumes in the cluster.

Once the migration is complete, you can group the volume with the other volumes, and the wizard will finish without errors. Any protection policies (snapshot rules and replication

rules) assigned at the volume-group level are applied automatically to all new and existing volumes in the group. If a volume-group protection policy includes a replication rule, it may take time for initial replication to complete if volumes contain significant data.

Consistent LUN IDs

It is a best practice with cluster volumes (including CSVs) to use a consistent LUN number across all nodes in the Hyper-V cluster. In PowerStore Manager, run the **Map Hosts** wizard and select **Provide a Logical Unit Number**.

PowerStore data reduction and Hyper-V

With PowerStore, inline data reduction (compression and deduplication) is enabled by default and cannot be disabled.

Data reduction works seamlessly in the background with PowerStore. Hyper-V environments benefit from PowerStore data reduction without any extra configuration required.

Data encryption

Data at rest encryption (D@RE) is enabled by default on PowerStore. No configuration steps are necessary to protect the drives.

Intuitive naming

To ease management of a Hyper-V environment on PowerStore, use intuitive, descriptive, and consistent naming where possible for all PowerStore objects. These objects include Hyper-V server host and node names, cluster names, volume names, host groups, and volume groups.

Using intuitive naming makes it easier to find and associate related objects when applying filters. This practice also simplifies applying rules to lists of objects in PowerStore Manager, such as the volumes list. This ability becomes more important as the number of objects on a PowerStore cluster grows.

Avoid using names for objects that are too generic, such as server1 or data1.

Data protection

Introduction

PowerStore provides data protection and recovery options for Hyper-V environments with snapshots, thin clones, refreshes, restores, and replication to remote PowerStore systems.

PowerStore snapshots and storage replication support Hyper-V environments and workloads:

- Boot-from-SAN disks
- Data volumes
- Cluster volumes
- Cluster shared volumes (CSV)
- In-guest iSCSI volumes
- Physical (pass-through) disks

Snapshots and thin clones

Snapshots are space-efficient, consisting of pointers to frozen data blocks, and they consume no extra space when they are taken. When a thin clone is created from a snapshot and mapped to a host, only new data written to the thin clone will consume additional storage space.

PowerStore snapshots can be taken of volumes that are mapped as LUNs to a Hyper-V environment regardless of content. This ability applies to boot-from-SAN volumes, data volumes, CSVs, pass-through disks, in-guest iSCSI volumes, and vFC volumes. These volumes can be replicated to other PowerStore clusters for DR or archive purposes along with their snapshots and thin clones.

For more information about PowerStore snapshots, thin clones, refreshes, and restores, see the online help in PowerStore Manager.

PowerStore snapshots allow administrators to perform the following actions in Hyper-V environments:

- Recover servers to a crash-consistent state, including Hyper-V hosts that are configured to boot-from-SAN, and guest VM workloads
- Provision lab or isolated test environments that are based on a production environment by using thin clones
- Provision new servers using gold images

Application consistency with snapshots

PowerStore snapshots of Hyper-V hosts, VMs, and workloads are crash-consistent by default. Snapshots can be taken manually, or automatically as part of a recurring schedule.

When performing a recovery using a crash-consistent snapshot, it is like having the server or workload recover from a power loss.

Often, servers and nontransactional workloads can be recovered to a crash-consistent state without complications.

Transactional workloads such as Microsoft SQL Server risk data corruption and data loss if recovering to a crash-consistent state.

Consider these recommendations if application consistency is wanted before taking a PowerStore snapshot.

- Use application-native tools to place a workload in a consistent state temporality (I/O is paused, and active data is saved to disk).
- Stop application services temporarily (to pause application I/O).
- Power off the host or VM that is hosting the workload. However, this method is often disruptive and impractical in production environment. This method is used to create a gold image of a system-prepared (Sysprep) host or VM after it is powered off.
- Use a Microsoft volume shadow copy service (VSS) aware process such as backup software that can place a server or workload in a consistent state.

When the host or workload is in a consistent state, take a PowerStore snapshot. Then, revert the host, VM, or workload to its active state.

If possible, use scripting and automation tools to orchestrate a process that performs these steps automatically. Orchestration is desirable for several reasons:

- Reduce administrative overhead
- Eliminate mistakes due to human error
- Automate a process so it can run after hours without requiring administrator action

The following sections provide examples for how to use PowerStore snapshots and thin clones in a Hyper-V environment.

Use PowerStore snapshots to recover guest VMs

You can recover a Hyper-V guest VM to a previous point in time using crash-consistent PowerStore snapshots of the underlying host volume containing the VHD. You can also use snapshots to create copies of VMs in an isolated environment. You may perform this action at the same or a different location when volume replication between PowerStore appliances is used. This section provides guidance and best practices for recovery options using snapshots.

Recover a guest VM on a stand-alone Hyper-V host

In this scenario, the VHD and configuration files for a VM reside on a data volume that is mapped to a Hyper-V host or cluster.

If the VM VHD and configuration files reside on separate host data volumes, configure a PowerStore volume group for these volumes in PowerStore Manager. This action ensures that crash-consistent snapshots occur at the same exact time. For example, a boot VHD for a VM might reside on one host volume, while one or more data VHDs might reside on another host volume.

When performing a recovery of a VM with PowerStore snapshots, there are several options.

Option 1: Use PowerStore to refresh or restore the existing data volume on the host that contains the VM configuration and VHDs using a snapshot. You could also replace the volume with a thin clone.

- This option may only be practical if the data volume contains only one VM. If the data volume contains multiple VMs, the options still work if all the VMs are being recovered to the same point in time. Otherwise, option 2 or 3 would be necessary if needing to recover only one VM.
- This option allows the VM being recovered to power on without any additional configuration or recovery steps required.
- If you use a thin clone to replace the original volume, document the LUN number, drive letter, or mount point information for the volume to be recovered. Perform this action before starting the recovery.

Option 2: Map a thin clone containing the VM configuration and VHDs to the host as a new volume, side by side, using a new drive letter or mount point. Manually recover the VM by copying the VHDs from the thin clone to the original location.

- Delete, move, or rename the original VHDs.
- After copying the recovered VHDs to their original location, rename them and use Hyper-V manager to reassociate them with the guest VM. This action may be necessary to allow the guest VM to start without permissions errors.
- This option may not be practical if the VHDs are large. In this case, delete the original VM, and import the recovery VM (Hyper-V 2012 and later versions). Alternately, create the recovery VM as a new VM directly from the thin clone. After the recovery, unmap the original data volume from the host if it is no longer required.
- This method helps recover a subset of data from a VM by mounting a recovery VHD disk as a volume on the host server temporarily.

Option 3: Map the thin clone to a different Hyper-V host, and recover the VM there by importing the VM configuration. Alternately, create a VM that points to the virtual hard disks on the recovery volume.

- This method is common when the original and recovery VMs must be online simultaneously but must be isolated from each other.
- You may also use this option when the original host server is no longer available.

Before VM recovery begins, record essential details about the VM hardware configuration in case importing a VM configuration is not supported or fails. These details can include the number of virtual CPUs, memory, virtual networks, and IP addresses.

Recover a guest VM on a cluster shared volume

Previous sections describe using PowerStore snapshots to refresh or restore CSVs or using thin clones to recover guest VMs that reside on one or more CSVs. These processes are similar to recovering a guest VM to a stand-alone host. However, recovering a VM from a thin clone in this scenario may require changing the disk signature first.

Windows Server (including Hyper-V hosts and nodes) assigns each disk a unique ID (or signature). Here are two disk ID examples, each consisting of a string of characters in hexadecimal format:

- For a GPT disk: EF4F117E-6F16-4927-B696-CFA9AE988D80
- For an MBR disk: 045C3E2F4

Volumes mapped to a server must each have a unique disk ID to avoid a possible service interruption.

When a PowerStore snapshot is taken of a Windows Server Hyper-V disk, the snapshot is an exact point-in-time copy. This copy includes the disk ID that Windows Server assigns to the disk. Thin clones that are created from a PowerStore snapshot also have the same disk ID as the source volume.

Stand-alone Windows or Hyper-V servers avoid disk ID conflicts because stand-alone servers can automatically detect duplicate disk IDs and change them dynamically without user intervention. However, clustered Windows Hyper-V nodes with CSVs require administrators to manually change the disk ID of a thin clone before it is mapped to cluster nodes. This requirement is due to Hyper-V cluster nodes not being able to dynamically

change conflicting disk IDs on clustered disks. This requirement is important because a disk ID conflict results in an inability to bring the disk online and may also result in a service interruption.

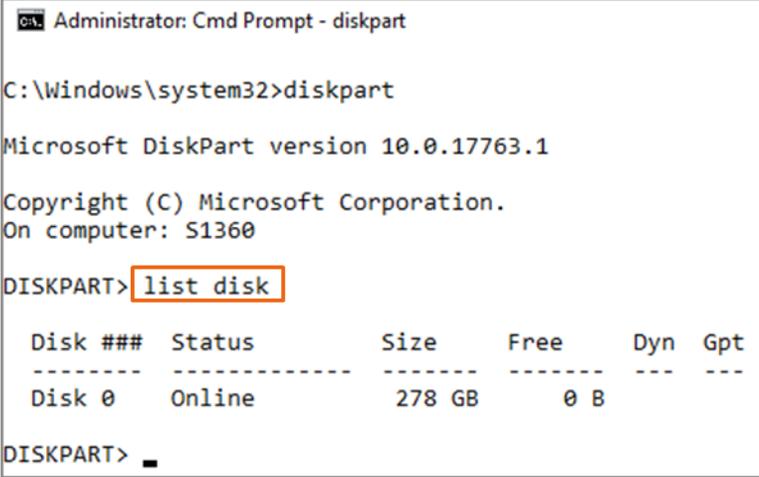
There are two options to address a duplicate disk ID with a thin clone with PowerStore:

- **Option 1:** Map the thin clone to a stand-alone Windows Server host that is not a member of the cluster. Then, copy the guest VM files over the network to recover the guest.
- **Option 2:** Map the thin clone to a stand-alone Windows Server host outside the cluster, and use diskpart to manually change the disk ID from the command line. After changing the disk ID, map the thin clone to the cluster.

Change a disk ID with diskpart

Follow these steps to change a disk ID before mapping a thin clone of a cluster volume back to that same cluster. You can also use PowerShell.

1. Log in to a stand-alone Windows Server (with or without the Hyper-V role installed) that is available in PowerStore. This server must not be a member of the Hyper-V cluster.
2. Open a command window with administrator rights.
3. Enter **diskpart**, and press **Enter**.
4. Enter **list disk**, and press **Enter**.
5. Note the current list of disks. In this example, **Disk 0** is the only disk.



```

Administrator: Cmd Prompt - diskpart

C:\Windows\system32>diskpart

Microsoft DiskPart version 10.0.17763.1

Copyright (C) Microsoft Corporation.
On computer: S1360

DISKPART> list disk

   Disk ###  Status              Size               Free              Dyn  Gpt
   -----  -
   Disk 0    Online              278 GB             0 B
DISKPART>

```

6. Use PowerStore Manager to map a thin clone of the cluster disk to this host.
7. From the diskpart command prompt, enter **rescan**, and press **Enter**.
8. Enter **list disk**, and press **Enter**.

The new disk (the thin clone) should be listed in an offline state.

```

Administrator: Cmd Prompt - diskpart

C:\Windows\system32>diskpart

Microsoft DiskPart version 10.0.17763.1

Copyright (C) Microsoft Corporation.
On computer: S1360

DISKPART> list disk

   Disk ###  Status              Size       Free       Dyn  Gpt
   -----  -
   Disk 0    Online              278 GB     0 B
   Disk 1    Offline             200 GB     0 B

```

9. To select the offline disk, enter **select disk <#>**, and press **Enter**.
10. Enter **online disk**, and press **Enter** to bring it online.
11. Type **list disk**, and press **Enter** to confirm that the disk is online.

```

DISKPART> select disk 1

Disk 1 is now the selected disk.

DISKPART> online disk

DiskPart successfully onlined the selected disk.

DISKPART> list disk

   Disk ###  Status              Size       Free       Dyn  Gpt
   -----  -
   Disk 0    Online              278 GB     0 B
   * Disk 1    Online             200 GB     0 B

```

12. Enter **uniqueid disk**, and press **Enter** to view the current ID for the disk.
13. To change the disk ID, enter **uniqueid disk ID=<newid>**, and press **Enter**.
 - In this example, only the last character of the disk ID is changed to make it unique.
 - For an MBR disk, the disk ID is an eight-character string in hexadecimal format.
 - For a GPT disk (shown in this example), the disk ID is a longer Globally Unique Identifier (GUID) that is also in hexadecimal format.

Note: If the disk is read-only, an error is returned when attempting to change the disk ID. If this error occurs, enter **attributes disk clear readonly**, press **Enter** to clear the read-only attribute, and repeat this step.

14. Enter **uniqueid disk** again, and press **Enter** to verify the new ID.

```

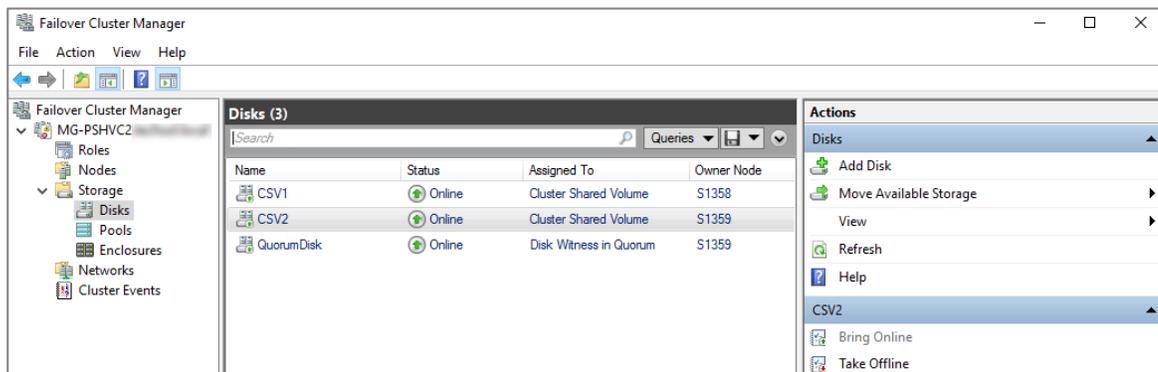
DISKPART> uniqueid disk
Disk ID: {12345678-1234-1234-1234-123456789ABC}
DISKPART> uniqueid disk id=12345678-1234-1234-1234-1234567891bd
DiskPart has encountered an error: The media is write protected.
See the System Event Log for more information.
DISKPART> attributes disk clear readonly
Disk attributes cleared successfully.
DISKPART> uniqueid disk id=12345678-1234-1234-1234-1234567891bd
DISKPART> uniqueid disk
Disk ID: {12345678-1234-1234-1234-1234567891BD}
DISKPART> _

```

15. Now that the thin clone has a new disk signature, exit from diskpart.
16. Unmap the disk from the stand-alone host server using PowerStore Manager, and map the disk to the specified Hyper-V cluster.
17. Perform a rescan disk on all nodes of the Hyper-V cluster, and bring the disk online. If Windows has automatically assigned a drive letter to any volumes on the disk, remove the drive letters, and return the disk to an offline state.

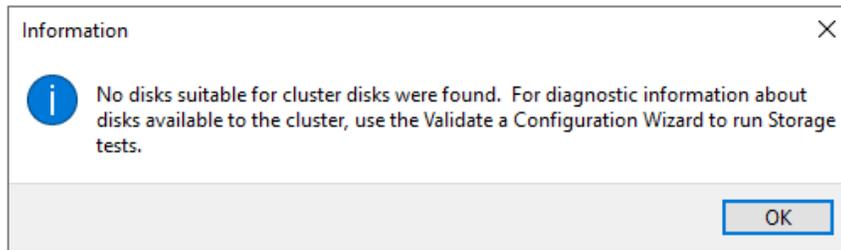
Note: We recommend disabling automount as a best practice to prevent hosts from automatically assigning drive letters to volumes in Hyper-V recovery scenarios. See [Disable automount](#) for details.

18. After making changes, put the disk into an offline state, and perform a rescan disk on each node in the Hyper-V cluster. Failure to do a rescan on all Hyper-V nodes will interfere with disk discovery in the next step.



19. Add the disk to the Hyper-V cluster. If the original disk was a CSV, convert the disk to a cluster shared volume using the **Actions** menu in Failover Cluster Manager.

Note: If the cluster is unable to discover the disk, run cluster validation and examine the report for disk errors. After resolving any errors, attempt to add the disk again.



20. It may be necessary to clear the cluster reservation attribute on the disk before the disk can be added to Hyper-V. This action can be performed with PowerShell.
 - a. Open a PowerShell window with administrator privileges. Clear the cluster reservation on the disk so that failover cluster manager can discover and import the disk.

```

Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) Microsoft Corporation. All rights reserved.

PS C:\Windows\system32> Clear-ClusterDiskReservation -Disk 1

Clear-ClusterDiskReservation
Are you sure you want to clear the disk reservation on node S1360?
[Y] Yes [N] No [S] Suspend [?] Help (default is "Y"): Y
PS C:\Windows\system32>

```

- b. Close PowerShell.
21. After the volume is online, perform the required steps, such as data or VM recovery.

Use PowerStore snapshots to create a test environment

Besides VM recovery, you can use PowerStore snapshots to quickly create a test or development environment that mirrors a production environment. To complete this step, you can map thin clones to other host servers or clusters. When PowerStore snapshots containing VMs are replicated to another location, it simplifies performing this action at a different location.

Note: To avoid IP, MAC address, or server-name conflicts, place copies of existing VMs recovered from a thin clone in an isolated environment.

The procedure to use a thin clone to create a test environment from an existing Hyper-V guest VM is similar to VM recovery. The main difference is that the original VM continues operation, and the VM copy is configured so that it is isolated from the original VM.

Use PowerStore to create gold images

An administrator can create a base operating-system image (**gold image**) using PowerStore snapshots and thin clones. This method can accelerate and simplify the process of deploying new servers. You can use gold images to quickly deploy new servers that boot-from-SAN that use the same server hardware.

Gold images provide the following benefits:

- Server provisioning is quicker with minimal reconfiguration required.

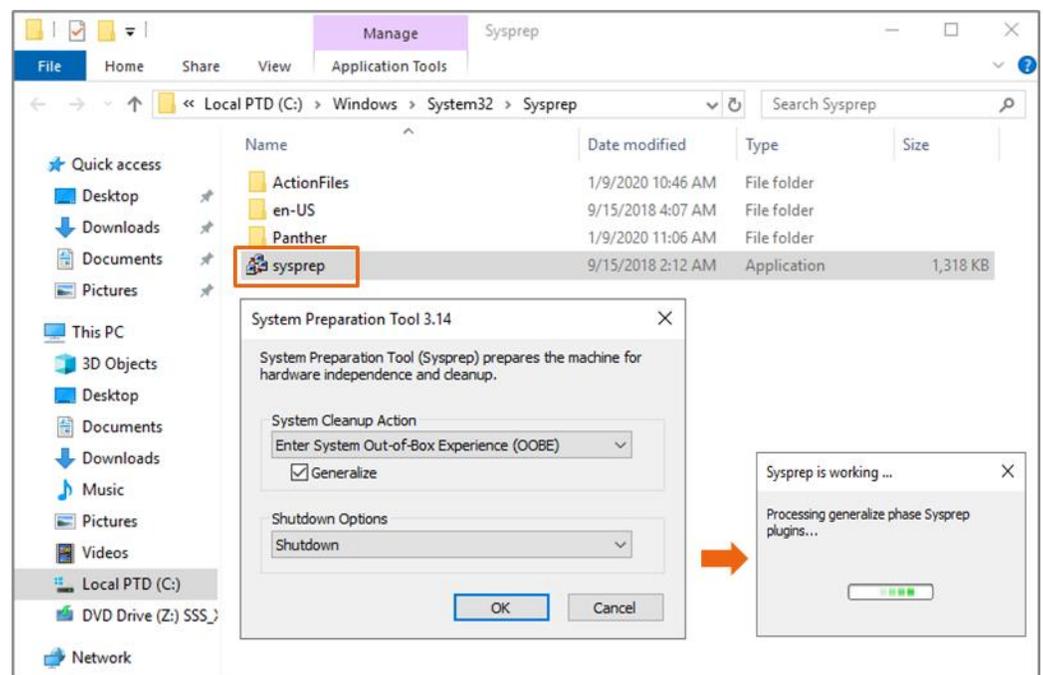
- When a host is provisioned from a thin clone of a gold image, only new data written to the thin clone consumes SAN space. Data that has not changed is read from the source volume.

The steps to configure a Windows Server boot-from-SAN gold image are as follows:

1. Create and map a PowerStore volume to a Windows host server or Hyper-V server that is configured to boot-from SAN.
2. Build the base operating-system image and install the wanted roles and features such as the Hyper-V role, failover clustering, and MPIO.
3. Configure the server, including MPIO, and patch the operating system. This action minimizes the changes that are required for each new server that is deployed from the gold image.
4. Once the operating system is fully staged, power off the operating system to put it into a consistent state. Then, take a manual PowerStore snapshot of the volume and select **No Automatic Deletion**. Use intuitive naming to simplify identifying the purpose of these snapshots later.

This snapshot represents the point in time before running the Microsoft system preparation tool **Sysprep**. The administrator can restore the parent volume using this snapshot if the system preparation process encounters an issue, or needs updating later. You do not have to create a new image from the beginning.

Power on the server and run Sysprep. Select the options **Generalize**, **Out-of-box Experience**, and **Shutdown**. Click **OK**.



5. Once Sysprep completes and the server powers off (which ensures that it is in a consistent state), manually create another PowerStore snapshot of the volume. As before, select **No Automatic Deletion**. Add a description if required to help identify the purpose of this snapshot later.

6. Select the snapshot, and click **More Actions > Create Thin Clone Using Snapshot**.
7. To view the thin clone, view the volumes list under **Storage**.

Note: If the thin clone was created from a cluster disk and is mapped back to the same cluster, change the disk signature of the thin clone. This action avoids having a disk signature conflict with the original disk. See [Change a disk ID with diskpart](#) for details about changing the disk signature of a Windows Server cluster disk.

8. Map the thin clone to the specified host or host group. Use LUN 0 if it is a boot-from-SAN volume for a physical Hyper-V host.
9. Boot the host and allow the initial boot process to complete.
10. Customize the server configuration as required.

Note: Do not use Microsoft System Center Virtual Machine Manager (SCVMM) to delete the guest VM. It also deletes the VHD. Use Hyper-V Manager instead.

11. Copy the gold source VHD file in safe location.
12. Create a VM. Make a copy of the gold VHD, and place it in the chosen location to serve as the boot disk for the VM.
13. Rename the VHD to reflect the new name of the VM or its purpose, and attach the VHD to the new VM as the boot disk.
14. Power on the VM, and customize the VM as needed.

Note: The appliance and node that owns a volume also owns all snapshots and thin clones that are created from the volume. This behavior usually does not cause significant resource-balancing issues unless many thin clones are created from the same source volume. Monitor PowerStore for imbalances, for example, if thin clones are used to configure many hosts to boot-from-SAN using a gold image.

PowerStore snapshots and Hyper-V VM migration tools

Windows Server 2012 Hyper-V and later versions offer native tools for migrating VMs. These methods may be preferred in most day-to-day situations. For example, when a guest VM is live-migrated from one node to another within the same Hyper-V cluster configuration, no data must be copied or moved. All nodes in that cluster have shared access to the underlying CSV.

However, there are many valid cases for using SAN-based snapshots to move VMs. When an administrator must migrate a guest VM from one host or cluster to another host or cluster, the data (VHDs) must be copied to the target host or cluster. This copy consumes network bandwidth and may require significant time if the VHDs are large. The copy can also consume additional SAN space unnecessarily because another copy of the data is created.

When moving VMs to another host or cluster, it may be quicker to unmap the host volumes containing the VM configuration files and VHDs. Then, map the volumes to the new target host or cluster. You can also perform this process using thin clones if preserving the state of the original environment temporarily is needed.

This procedure might involve a small amount of downtime for the VM being moved during a maintenance window. However, this approach might be more practical than waiting for a large amount of data to copy over the network. This copy would consume additional SAN space unnecessarily.

Disaster recovery with Hyper-V and PowerStore

Overview

A good business-continuity strategy always incorporates disaster recovery (DR) and disaster avoidance planning. A DR plan ensures that a company can recover as quickly as possible from data loss or from an interruption or failure preventing access to data. It is an important part of the overall IT strategy.

The DR scenarios that may be encountered are diverse and may vary by location. Disasters can be small or large. The loss of a single document that impacts one user is a disaster for that user. If not resolved quickly, a site failure might impact many users and jeopardize the future of the business.

The essential elements of DR are now commonplace, reliable, cost-effective, and easy to implement. They address or prevent most events that are most likely to occur. These protections might include moving key workloads to a cloud provider, or making tape backups with offsite storage, online backups, or disk-to-disk backups. Safeguards can also include network and physical security measures, malware protection, and redundant hardware and Internet connections. Also, DR can include SAN-based snapshots with remote replication, and battery backups or generators.

Business continuity becomes more complicated and costly with the size and number of locations. While virtualization technologies such as Microsoft Hyper-V can help ensure continuity in a disaster, they can also add complexity to the overall design.

PowerStore provides Hyper-V administrators with several options to protect, move, replicate, and recover data using snapshots, thin clones, and replications. Several practical day-to-day use cases and scenarios were described previously. This section addresses key disaster recovery and avoidance concepts that should be integral to business continuity planning with PowerStore and Hyper-V.

Cost and risk analysis

The most robust DR solutions might also be cost prohibitive. These solutions may require weighing the costs compared to the risks and determining the level of DR protection that is necessary. Factors to consider as part of a cost and risk analysis include the following:

- What are the specific regulations that apply to your service or industry?
- What are the terms of the service level agreements (SLAs) for business continuity that must be honored?
- Identify the applications and data that are the most mission-critical to the business or to customers.

Determine the recovery-time objective (RTO) for each application or service. In other words, the RTO is the time period that can something be down before the business impact becomes too great. Examples include the following:

- Practice management system: 30 minutes

- Messaging system: 4 hours
- Research and development server: 2 days

Determine the recovery point objective (RPO). In other words, the RPO is the amount of data loss that is acceptable for a subset of data. If PowerStore snapshots are used to protect a workload, and snapshots are taken every 24 hours, the RPO is up to 24 hours. An RPO of 24 hours means up to 24 hours of data loss may occur if the workload is recovered from the last snapshot. The frequency and retention of PowerStore snapshots can be increased to shorten the RPO as needed.

Determine the types of events that are most likely to occur in your area. For example, a coastal location may be prone to hurricanes or flooding.

Identify an alternate site that is distant enough so that the same event does not impact both locations.

Determine the cost (hardware, software, and staff) to design, implement, and support the necessary protections. Is the cost justified given the risk?

Disaster recovery and disaster avoidance

PowerStore provides Hyper-V administrators with many options for protecting, moving, archiving, and recovering data to protect against or recover from many DR and disaster-avoidance scenarios. In addition to the integrated redundancies with PowerStore architecture, administrators can use snapshots, thin clones, and replication in many creative ways to ensure business continuity.

Disaster recovery usually means reacting to an event that causes downtime that takes place unexpectedly with little or no warning. These events can be categorized as follows:

- Events that cause data loss such as malware infection, corruption, accidental deletion, or sabotage.
 - PowerStore snapshots and replication to a remote site are an important part of a recovery plan in these scenarios.
 - The redundancies that are built into PowerStore and Hyper-V, combined with a robust storage fabric that uses MPIO, can reduce risks associated with possible hardware failures.
- Events that interrupt the ability to access data within or between sites, such as a network or power failure (but no data is lost). Redundant Internet connections can help ensure that connectivity is maintained between sites. Battery backups can protect against a short-term power loss. A generator provides protection against longer-term power outages.
- Events that cause both loss of data and loss of access to a site. These occurrences are typically more significant and destructive such as a fire or natural disaster. If a site is at risk of these events, having another location that can carry on the business becomes critical. Using cloud-based solutions such as Microsoft Azure to protect critical VM workloads is an important strategy for businesses with a single location.

Disaster avoidance implies having enough lead time to proactively react to an impending event, such as a planned power outage, to avoid or minimize downtime. This strategy is commonly used when performing routine system maintenance with PowerStore, such as

a PowerStoreOS upgrade. With disaster avoidance, often there is time to react in a predictable and methodical way to avoid or minimize downtime. For example, an administrator may move a critical workload to another PowerStore appliance within a multiappliance cluster using integrated nondisruptive storage migration tools in PowerStore Manager. Data and workloads can be moved to an alternate location using PowerStore replication, before planned site maintenance at the main location causes an outage there.

A good business continuity plan includes both disaster recovery and disaster avoidance strategies. These strategies should use a combination of manual and automatic processes to address the widest possible range of scenarios that is feasible within the budget.

Conclusion

Summary

Careful planning, adherence to best practices, and testing are essential for a successful deployment of Microsoft Hyper-V on Dell PowerStore storage. PowerStore is well suited to host high-density high-demand Hyper-V virtual workloads. PowerStore provides Microsoft Hyper-V administrators with an all-inclusive complement of tools, options, and features. Following the guidance in this white paper will help you design and deliver a resilient, reliable, and highly performant experience for your Hyper-V users.

References

Dell Technologies documentation

The following Dell Technologies documentation provides other information related to this document. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- [Dell Technologies Storage Info Hub](#)
- [Dell Technologies Support](#)

Microsoft documentation

For Microsoft documentation, see the following resources:

- [Microsoft Windows Server Documentation Library](#)
- [Microsoft Virtualization Documentation Library](#)
- docs.microsoft.com