

Titan T2000

Microsoft Hyper-V



Dell PowerVault ME5 Series: Microsoft Hyper-V Best Practices

March 2022

H19007

White Paper

Abstract

This document provides best practices for configuring Microsoft Windows Server Hyper-V to perform optimally with Dell PowerVault ME5 storage.

Dell Technologies Solutions

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 © 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 March 2022 H19007.

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**
- Introduction5**
- Storage and transport best practices7**
- Hyper-V best practices10**
- ME5 snapshots and storage replication with Hyper-V.....33**
- Boot from SAN39**
- Conclusion.....42**
- References.....43**

Executive summary

Overview

This document provides best-practice guidance for deploying and optimizing the Microsoft Windows Server Hyper-V hypervisor role with Dell PowerVault ME5 storage.

Hyper-V and ME5 storage are feature-rich solutions. They 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 any individuals who may evaluate, acquire, manage, operate, or design a Dell storage environment using Dell PowerVault systems. The reader should have working knowledge of Dell PowerVault ME5 storage and Microsoft Hyper-V.

Revisions

Date	Description
March 2022	Initial release

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 [PowerVault Info Hub](#).

Introduction

Dell PowerVault ME5 overview

Dell PowerVault ME5 (ME5) is a next-generation, entry-level block storage array. ME5 storage is purpose-built and optimized for SAN and DAS virtualized workloads. ME5 storage is well suited to support Microsoft workloads including the Hyper-V hypervisor role. ME5 storage arrays are available in a 2U or 5U base system with optional disk expansion by adding additional disk enclosures. ME5 simplifies the challenges of providing storage capacity, performance, expansion, and redundancy for your Microsoft Hyper-V environment.



Figure 1. Dell PowerVault ME5084 5U and ME5024 2U storage arrays

Dell PowerVault ME5 storage provides the following features:

Ease of Management: PowerVault Manager is an intuitive, all-inclusive HTML5-based management web UI that is integrated on each ME5 array. CLI and REST API are also supported management tools.

Simplicity: It is simple and quick to install and configure ME5 storage in a few minutes.

Performance: Compared to its predecessor PowerVault ME4, ME5 offers increased power and scale with updated Intel processors with double the cores. The ME5 processing power delivers significant performance gains over the ME4. ME5 also delivers increased storage capacity, higher bandwidth, and more throughput.

Connectivity: ME5 storage supports the following front-end connectivity options for presenting block storage to host servers:

- 12 Gb SAS (four ports per controller)
- 16/32 Gb FC (four ports per controller)
- Two iSCSI options:
 - 10 GbE BaseT (four ports per controller), or
 - 10/25 GbE optical (four ports per controller)

Scalability: The PowerVault ME5 base array is 2U or 5U depending on the model. The 2U models (ME5012 and ME5024) support up to 12 or 24 drives respectively in the base system. The 5U model (ME5084) supports up to 84 drives in the base system. ME5 2U and 5U base systems also support adding optional expansion enclosures of 12, 24, and 84 drives, for up to 336 drives total. ME5 supports up to six PB of storage capacity. The firmware will support up to eight PB of storage capacity once higher-capacity drives are available.

All-inclusive software: ME5 supports the following features:

- Management tools: PowerVault Manager web UI, CLI, and REST API
- Volume copy
- Snapshots (redirect on write)
- IP and FC bi-directional asynchronous replication (ME5 to ME5; ME5 to ME4; ME4 to ME5)
- VMware vCenter Server and VMware Site Recovery Manager integration
- SSD read cache
- Thin provisioning
- Three-level autotiering
- ADAPT (distributed RAID)
- Controller-based encryption (SEDs) with internal key management

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

Note: Most ME5 storage features work seamlessly in the background, regardless of the platform or workload. Usually, the default storage settings for ME5 are optimal for Hyper-V environments. This document provides configuration strategies and configuration options for ME5 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 leverages the Hyper-V role to provide virtualization technology. Hyper-V is one of many optional roles that are offered with Windows Server. ME5 supports Windows Server versions 2016, 2019, and 2022, including the Hyper-V role.

Note: Support requirements may change over time. To verify operating system compatibility with ME5 for your environment, see the latest release notes and the *Dell PowerVault ME5 Storage System Support Matrix* at [Dell Technologies Support](#).

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

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 will highlight 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.

Storage and transport best practices

Essential documentation

The following documents provide essential guidance for the planning, configuration, and deployment of PowerVault ME5. Administrators should review and follow the guidance in these documents at [Dell Technologies Support](#) to ensure a successful deployment of Windows Server and Hyper-V on ME5:

- *Dell PowerVault ME5 Owner's Manual*
- *Dell PowerVault ME5 Deployment Guide*
- *Dell PowerVault ME5 Administrator's Guide*
- *Dell PowerVault ME5 Release Notes*
- *Dell PowerVault ME5 Support Matrix*
- *Dell PowerVault Host Configuration Guide*

This white paper provides supplemental best practice guidance.

Right-size the ME5 storage array

Before deploying ME5, 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 PowerVault 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 an ME5 storage solution for your environment and 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 (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

Disk groups, pools, and RAID configuration

Choosing the type of disk, disk pool, and RAID configuration is an important part of right-sizing ME5 storage. Sizing considerations include the following:

- Storage capacity needs
- Read and write IOPS demand
- Performance and latency needs

See the *Dell PowerVault ME5 Administrator's Guide* at [Dell Technologies Support](#) for an in-depth review of the following topics.

- Linear and virtual disk groups
- Disk pools
- RAID levels
- Hot spare configuration options

From the perspective of Hyper-V, all available configuration options are supported. Choosing the best type of disk group and RAID option is a function of the workload running on Hyper-V, and the *Dell PowerVault ME5 Administrator's Guide* provides sizing guidance.

For this paper, an ME5024 array is configured with 24 spinning disks in the base enclosure, with two 12-disk pools with ADAPT. Each pool is assigned to a separate

controller to achieve balance. This configuration provides an excellent starting point for good overall performance, capacity, and expandability.

ADAPT RAID

One option discussed in the *Dell PowerVault ME5 Administrator's Guide* is the ME5 ADAPT option for RAID. ADAPT supports distributed sparing for fast rebuild times, and large-capacity disk groups. However, ADAPT requires a minimum of 12 drives to start with, and all disks must be of the same type and be in the same tier.

- If more disk performance is needed, SSDs can be used instead of spinning disks.
 - Use spinning disks for low-demand workloads, and storage capacity for archive data.
 - Use SSDs for demanding workloads that require high read and write IOPS performance and low latency.

Disk capacity and performance

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

Administrators must plan for IOPS and capacity when sizing the ME5 for Hyper-V or any other workload.

Transport and front-end connectivity

The ME5 provides block storage to host servers using direct-attached storage (DAS) or a storage area network (SAN).

The ME5 supports 12 Gb SAS, 16/32 Gb FC, and 10/25 GbE iSCSI for a DAS configuration. The ME5 supports 16/32 Gb FC, and 10/25 GbE iSCSI for a SAN configuration.

DAS may not be a practical option for Hyper-V in your environment. DAS limits the number of physical hosts to a maximum of four, assuming each host is configured to use two data paths to the ME5 for redundancy.

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

Before reading further, refer to the *Dell PowerVault ME5 Deployment Guide* at [Dell Technologies Support](#). This guide provides a thorough summary of the different DAS, SAN, host, and replication cabling options available with the ME5.

Windows Server and Hyper-V support the available front-end transport configuration options listed in the *Dell PowerVault ME5 Deployment Guide*.

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.
- If a Hyper-V environment is likely to scale beyond four physical hosts attached to the same ME5 array, start with a SAN configuration (FC or iSCSI).
 - Migrating from an FC or iSCSI DAS configuration to a SAN configuration may be disruptive.
- If the ME5 is configured to replicate to another ME5 or ME4 array, two of the four FE ports (0 and 1) on each controller are dedicated to replication traffic.
- SAS FE ports are supported in a DAS configuration only. The use of SAS FE for Hyper-V may be acceptable if the following conditions are true:
 - The Hyper-V environment will not expand beyond four hosts (assumes that two SAS paths are configured for each host).
 - PowerVault replication is not needed.

Other factors to consider include the following:

- With DAS, the hosts must be within reach of the physical cables. Place hosts in the same rack or an adjacent rack that is within convenient cabling distance.
- 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.

Hyper-V best practices

Introduction

The PowerVault ME5 is an excellent choice for external storage for stand-alone or clustered Windows Servers including servers that are configured with the Hyper-V role. Core PowerVault features such as thin provisioning, 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 on applying Hyper-V best practices.

General Hyper-V best practices

General best practices for Hyper-V (not specific to PowerVault ME5 storage) are discussed in great detail in Microsoft documentation.

For example, go to docs.microsoft.com and search on 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 additional 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 here. 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 PowerVault.
- 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.
 - PowerVault offers CLI and REST API support for additional management and scripting functionality.
- Enable monitoring and alerting features to identify and resolve issues quickly.
 - Configure ME5 email alerts.

- Enable the Dell SupportAssist feature in ME5 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 PowerVault. 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 the ME5 for each host (DAS or SAN).
3. Install and configure MPIO on each host.
4. Use PowerVault Manager to create a host group on ME5.
5. Use PowerVault 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.

Validate a Configuration Wizard

Validating

Before You Begin
Select Servers or a Cluster
Testing Options
Confirmation
Validating
Summary

The following validation tests are running. Depending on the test selection, this may take a significant amount of time.

Progress	Test	Result
100%	List Services Information	The test passed.
100%	List Software Updates	The test passed.
100%	List System Drivers	The test passed.
100%	List System Information	The test passed.
100%	List TPM Information	The test passed.
0%	List Unsigned Drivers	Test is currently running.
	List Network Metric Order	Pending...
	Validate Cluster Network Configuration	Pending...
	Validate IP Configuration	Pending...
	Validate Network Communication	Pending...
	Validate Windows Firewall Configuration	Pending...
100%	List Disks	The test passed.
100%	List Disks To Be Validated	The test passed.
100%	Validate CSV Network Bindings	The test passed.
100%	Validate CSV Settings	The test passed.
100%	Validate Disk Access Latency	The test passed.
100%	Validate Disk Arbitration	The test passed.
100%	Validate Disk Failover	The test passed.
100%	Validate File System	The test passed.
100%	Validate Microsoft MPIO-based disks	The test passed.
100%	Validate Multiple Arbitration	The test passed.
100%	Validate SCSI device Vital Product Data (VPD)	The test passed.
100%	Validate SCSI-3 Persistent Reservation	The test passed.
100%	Validate Simultaneous Failover	The test passed.
100%	Validate Storage Spaces Persistent Reservation	The test passed.
	Validate Active Directory Configuration	Pending...
	Validate All Drivers Signed	Pending...

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.

Note: 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.

100%	Validate Operating System Version	The test passed.
100%	Validate Required Services	The test passed.
100%	Validate Same Processor Architecture	The test passed.
100%	Validate Software Update Levels	The test reported some warnings.
100%	Validate System Drive Variable	The test passed.

Test is currently running.

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, integration services must be installed and updated manually on VMs. Use the **Action** menu in Hyper-V Manager 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 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.

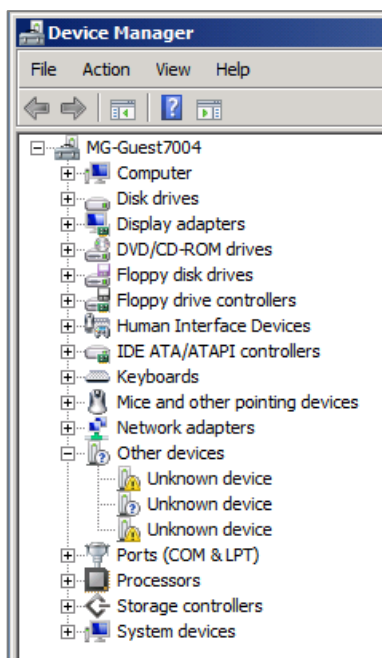


Figure 2. Unknown guest VM devices indicate missing or outdated integration services

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

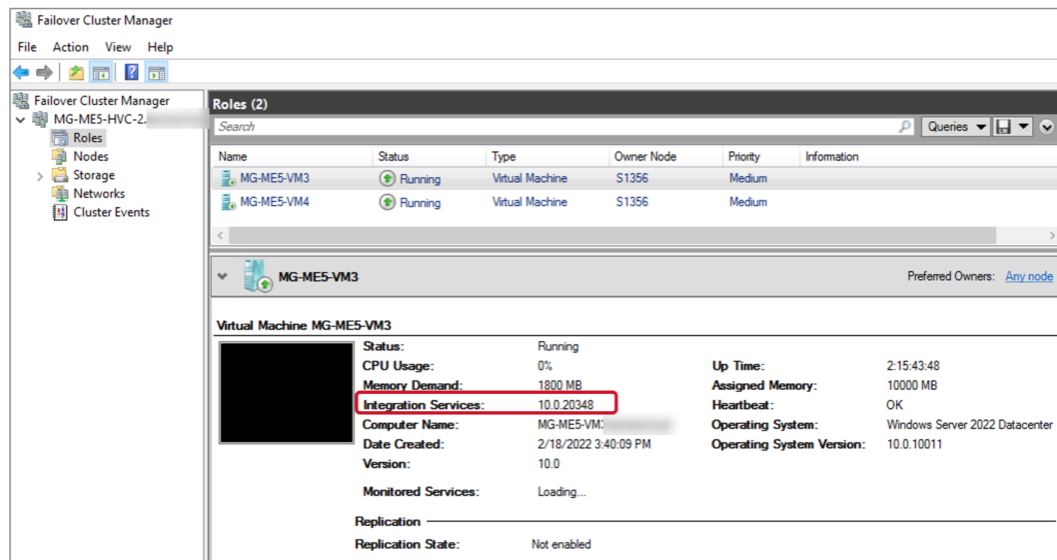


Figure 3. Verify integration services version with Failover Cluster Manager

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 eliminates 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.

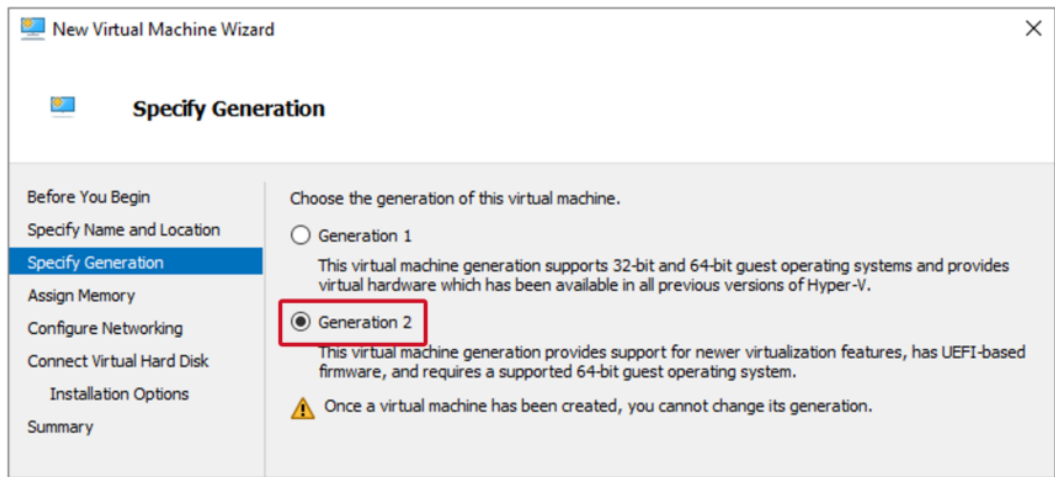


Figure 4. Guest VM generation option

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.

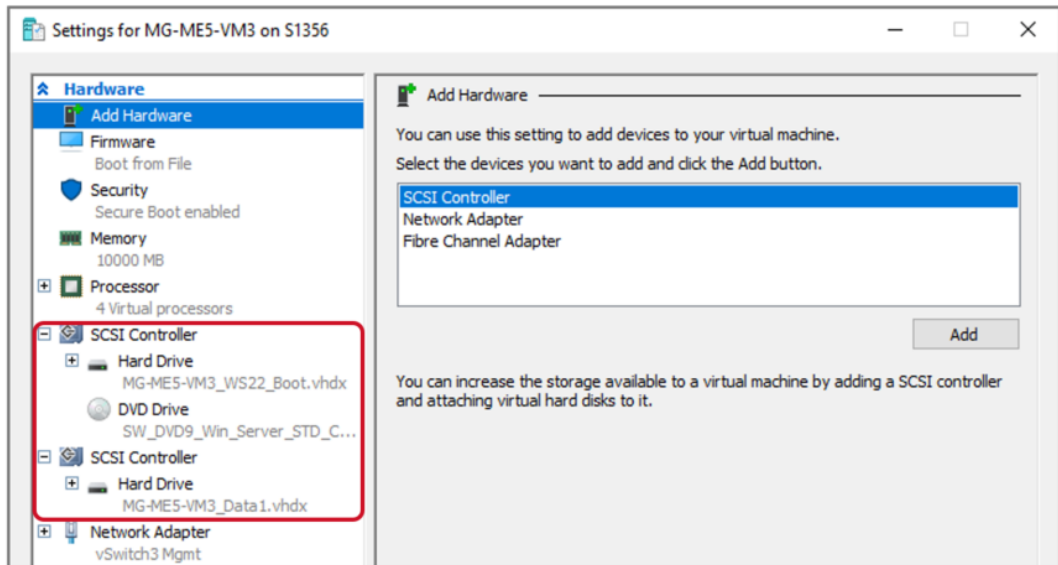


Figure 5. VM configured with two virtual SCSI controllers

Convert VMs to a newer generation

The VM generation cannot be changed once a VM has been created (see the warning message in Figure 4). 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 ME5 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, VHDX should be used 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.

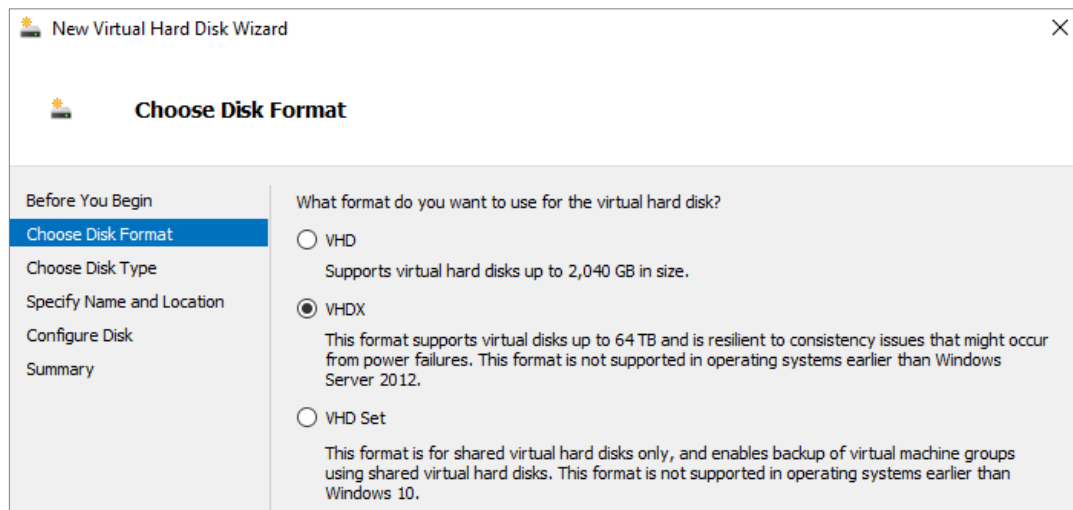


Figure 6. Virtual hard disk format options

Virtual hard disk type

In addition to the formatting options, a virtual hard disk can be designated as fixed, dynamically expanding, or differencing.

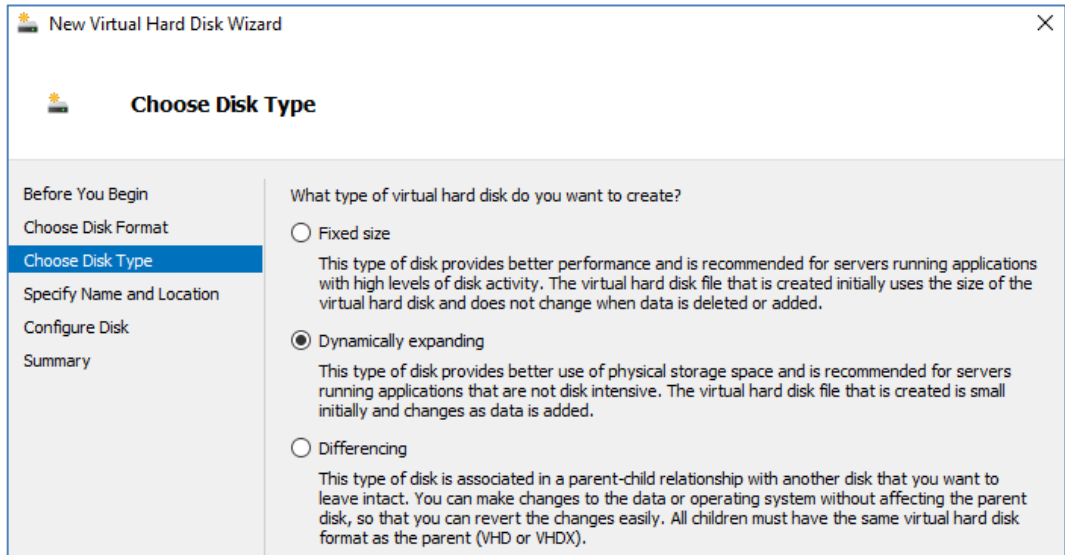


Figure 7. Options for virtual disk type

A dynamically expanding disk is the default type and will work well for most Hyper-V workloads on ME5 storage. If the ME5 is configured to use thin provisioning, 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 Figure 8 will present the full formatted size of 60 GB to the guest.




Figure 8. 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 ME5 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 leverage differencing VHDs.
 - Storage savings can be realized 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 ME5

Any virtual hard disk (fixed, dynamic, or differencing) will experience space usage efficiency on ME5 storage when the array is configured to use storage thin provisioning.

The example shown in Figure 9 illustrates 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, 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.

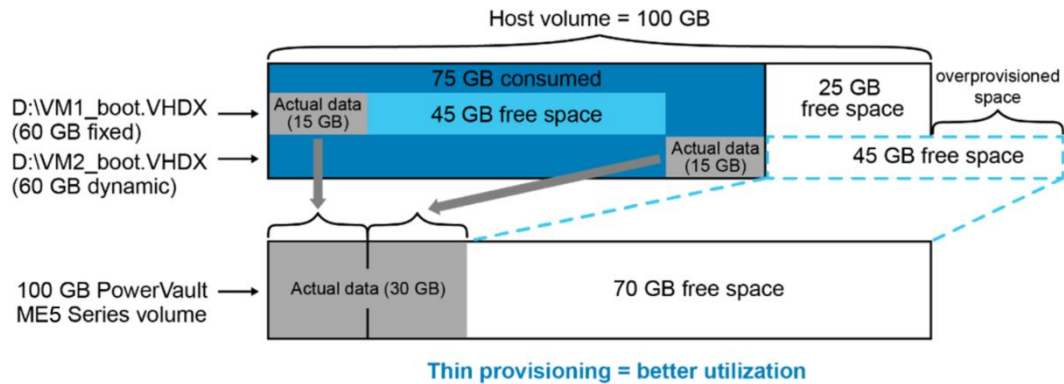


Figure 9. Thin provisioning with ME5

Comparatively, The ME5 array 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 ME5 when thin provisioning is leveraged. 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 ME5 arrays that leverage thin provisioning.
 - If 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 ME5 storage 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 ME5 storage snapshots. ME5 array-based snapshots and native Hyper-V snapshots function independently.

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




Figure 10. Location of native Hyper-V snapshots on the host or cluster volume

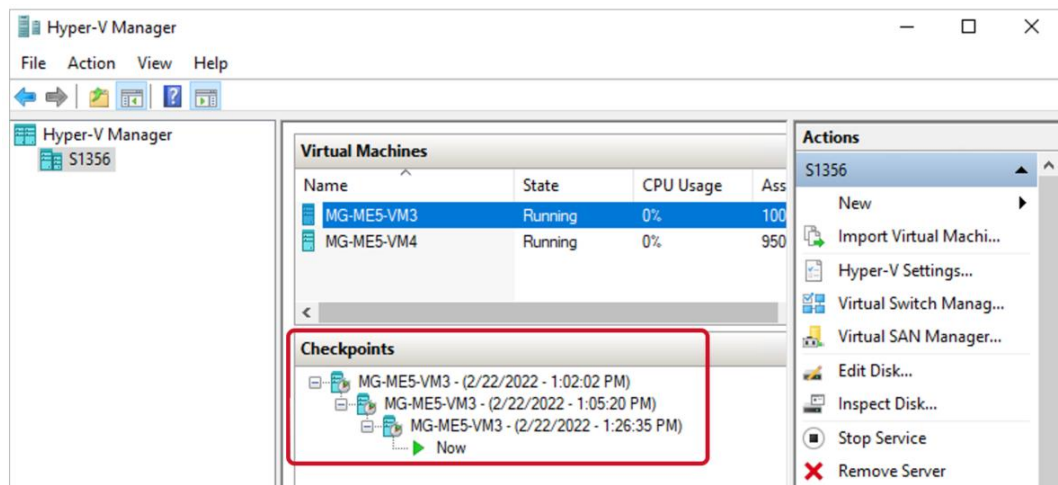


Figure 11. Hyper-V checkpoints are shown in a hierarchical tree

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.

The recommendation is to avoid using Hyper-V checkpoints, or to use them sparingly or temporarily.

Administrators can leverage ME5 array-based snapshots to protect and replicate VM data, in addition to using native Hyper-V VM replication tools.

Present ME5 storage to Hyper-V hosts and VMs

Hyper-V supports DAS (SAS, FC, iSCSI) and SAN (FC, iSCSI) configurations with ME5.

See the *Dell PowerVault ME5 Administrator's Guide* and the *Dell PowerVault ME5 Deployment Guide* at [Dell Technologies Support](#) for an in-depth review of transports and cabling options.

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.

Mixed transports

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

There is limited Microsoft support for mixing transports on the same host. Mixing transports is not recommended 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 need 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.

Volume	Layout	Type	File System	Status	Capacity	Free Spa...	% Free
(Disk 0 partition 3)	Simple	Basic		Healthy (R...	523 MB	523 MB	100 %
FC_Data1 (F:)	Simple	Basic	NTFS	Healthy (B...	1769.49 GB	361.96 GB	20 %
FC_Data2 (G:)	Simple	Basic	NTFS	Healthy (B...	1816.06 GB	475.65 GB	26 %
iSCSI_Data3 (I:)	Simple	Basic	NTFS	Healthy (B...	2142.02 GB	747.96 GB	35 %
iSCSI_Data4 (J:)	Simple	Basic	NTFS	Healthy (B...	2235.15 GB	841.08 GB	38 %
SAN_Boot_ME5 (C:)	Simple	Basic	NTFS	Healthy (B...	139.08 GB	104.67 GB	75 %
System Reserved	Simple	Basic	NTFS	Healthy (S...	100 MB	67 MB	67 %

Disk	Partitions
Disk 0 Basic 139.69 GB Online	System Reserved 100 MB NTFS Healthy (System, Active, ...) SAN_Boot_ME5 (C:) 139.08 GB NTFS Healthy (Boot, Page File, Crash Dump, Primary Partition) 523 MB Healthy (Recovery Partition)
Disk 1 Basic 1769.50 GB Online	FC_Data1 (F:) 1769.49 GB NTFS Healthy (Basic Data Partition)
Disk 2 Basic 1816.06 GB Online	FC_Data2 (G:) 1816.06 GB NTFS Healthy (Basic Data Partition)
Disk 3 Basic 2142.02 GB Online	iSCSI_Data3 (I:) 2142.02 GB NTFS Healthy (Basic Data Partition)
Disk 4 Basic 2235.16 GB Online	iSCSI_Data4 (J:) 2235.15 GB NTFS Healthy (Basic Data Partition)

Figure 12. A host with two FC volumes and two iSCSI volumes mapped concurrently

Consider the following example:

1. 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.

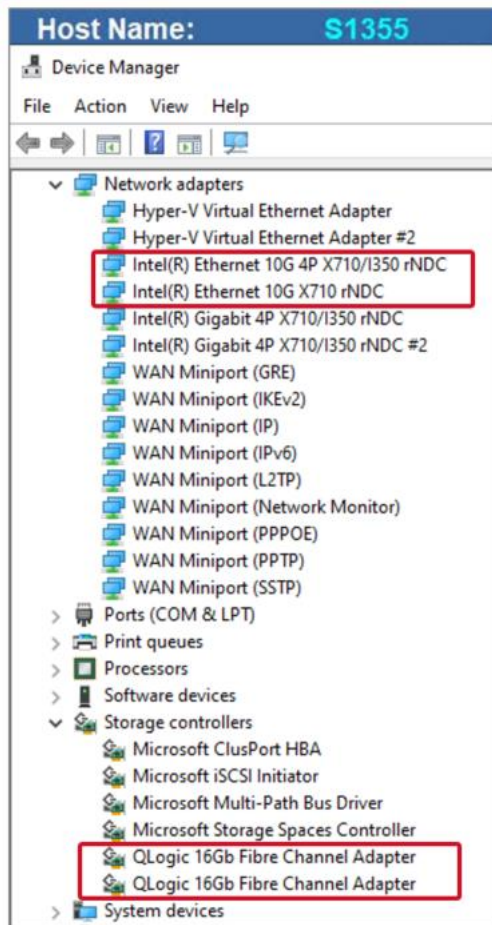


Figure 13. Host server with iSCSI NICs and FC HBAs

2. An existing FC volume is mapped to the host from a legacy storage array that is being retired.
3. Create a host object on the new ME5 array that uses iSCSI mappings.
4. Map a new volume on the ME5 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 ME5 array.
5. Migrate the workload from the existing FC volume to the new iSCSI volume on the ME5 array.
6. 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 ME5 arrays.

Windows and Hyper-V hosts default to the **Round Robin with Subset** policy with ME5 storage. Round Robin with Subset will work well for most Hyper-V environments. Specify a different supported MPIO policy if necessary.

In this example, each ME5 storage controller (Controller A and Controller B) has four FC front-end (FE) paths connected to dual fabrics, for eight paths total. Connecting fewer FE paths, such as two on each controller for four paths total, is also acceptable.

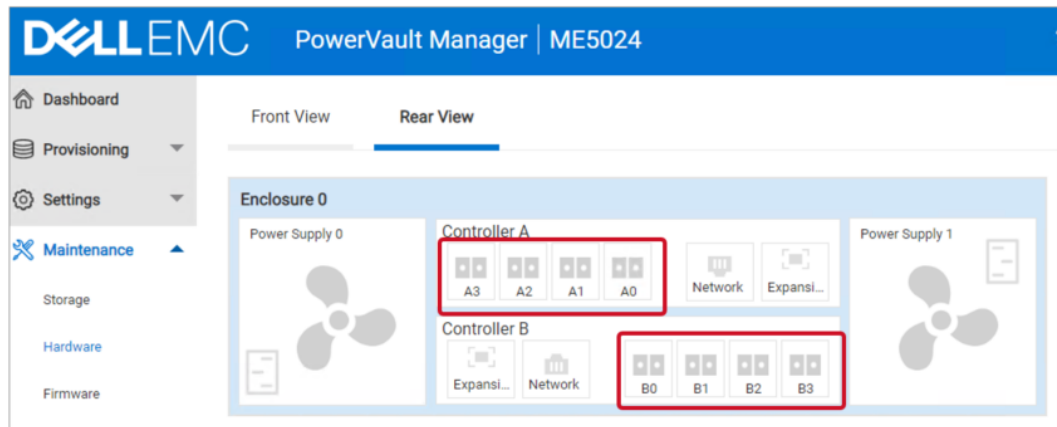


Figure 14. Controller front-end ports A0 – A3 and B0 – B3

In Figure 15, a volume mapped from ME5 to a host lists eight total paths.

- Four paths that are optimized (to the primary controller for that volume)
- Four paths that are unoptimized (to the secondary or standby controller for that volume).

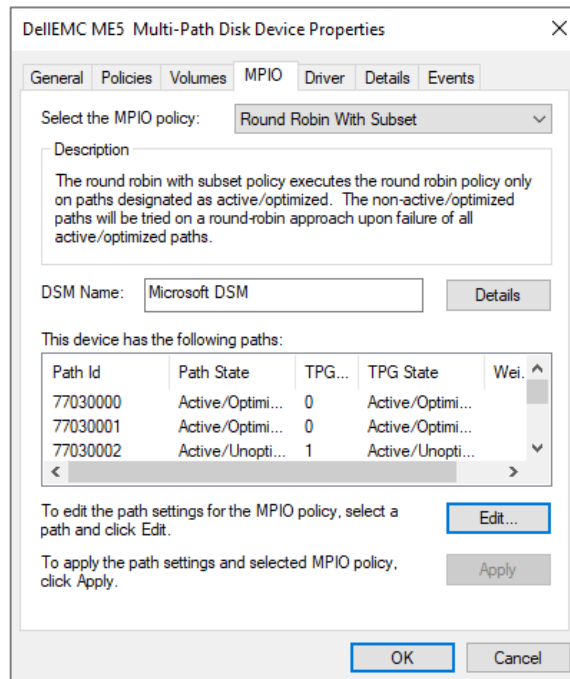


Figure 15. Verify MPIO settings (Microsoft DSM)

The **Active/Optimized** paths are associated with the ME5 storage controller that the volume is assigned to. The **Active/Unoptimized** paths are associated with the secondary or standby ME5 storage controller for that same volume.

When creating volumes on PowerVault, 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 ME5 documentation or Dell Technologies support.
- Connect all available FE ports on an ME5 array (SAN mode) to use your preferred transport to optimize throughput and maximize performance.
- Configure dual fabrics and storage networks for switch and path level redundancy.
- Configure each host to use at least two ports with a SAN or DAS configuration (iSCSI, SAS, 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.
 - ME5 controller firmware
 - 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 PowerVault ME5 Support Matrix* at [Dell Technologies Support](#).

Guest VMs and block storage options

ME5 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 ME5 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.

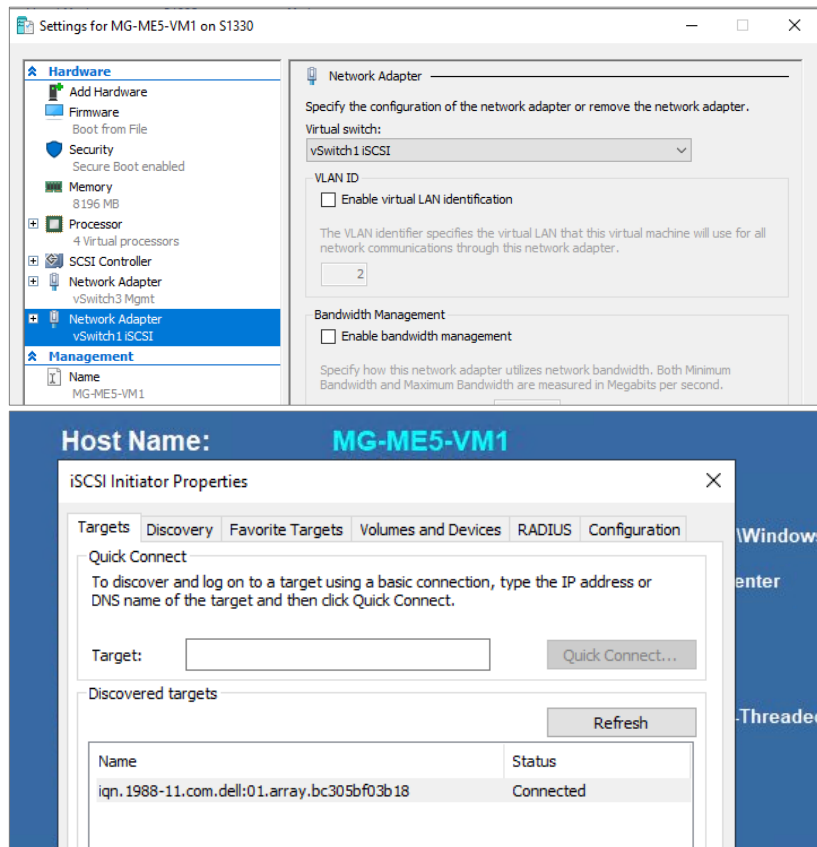


Figure 16. In-guest iSCSI

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.

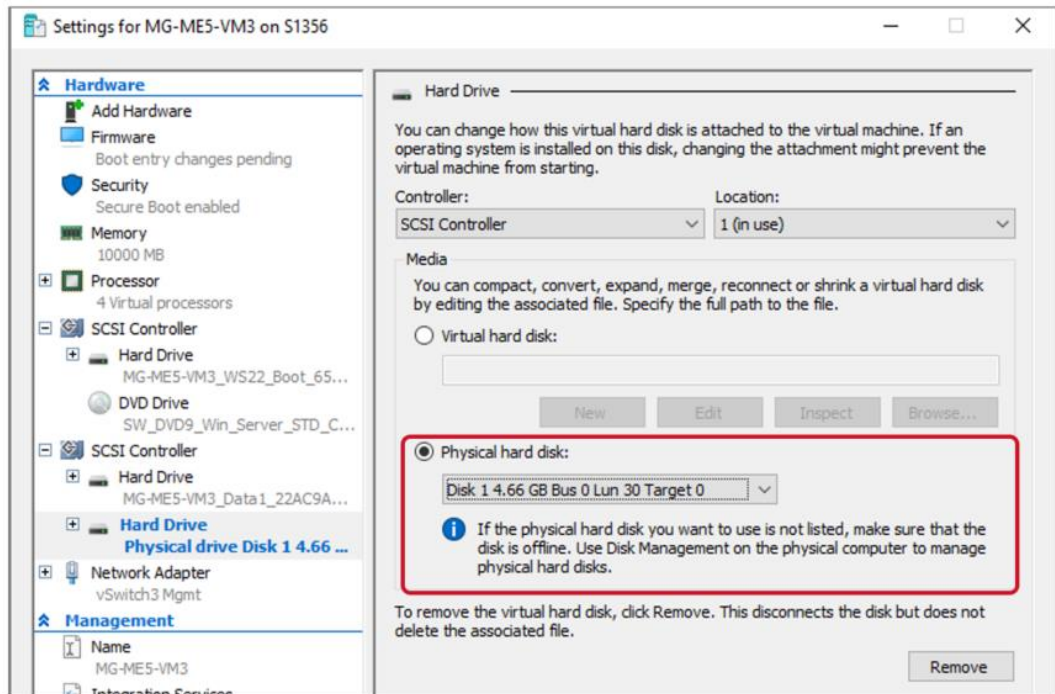


Figure 17. Hyper-V VMs support physical (pass-through) disks

In-guest iSCSI and pass-through disk use cases

ME5 arrays support in-guest iSCSI and pass-through disks (direct-attached disks) mapped to guest VMs. However, using direct-attached storage for guest VMs is not recommended as a best practice unless there is a specific use case that 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 need to troubleshoot the I/O performance of a volume and it must be isolated from all other servers and workloads.
- **Custom snapshot or replication policy:** It may be necessary in some use cases to apply a custom ME5 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 ME5 snapshots and replication to the CSV.
- **Capacity:** Legacy VHDs support a maximum size of two TB. VHDX supports a maximum size of 64 TB. If a data volume will exceed these limits, you may need 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 leverage ME5 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.

ME5 storage and Hyper-V clusters

Use a consistent LUN number when mapping shared volumes: quorum disks, cluster disks, and cluster shared volumes. Leverage host groups on the ME5 array to simplify the task of assigning consistent LUN numbers.

Note: Hyper-V hosts that use boot-from-SAN cannot be added to ME5 hosts groups. See the [Boot from SAN](#) section of this white paper for details.

Changing LUN IDs after initial assignment by ME5 may be necessary to make them consistent. By default, PowerVault 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 ME5 storage

Each cluster shared volume (CSV) will support one VM or many VMs. How many VMs to place on a CSV is a function of user preference, the workload, and how ME5 storage features such as snapshots and replication will be used. Placing multiple VMs on a CSV 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 ME5 array volumes are easier to manage.
- **Efficiency:** It is quicker and easier 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 an ME5 snapshot.
- **Replication control:** One-to-one gives administrators more granular control over what data gets replicated when ME5 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 an ME5 snapshot) avoids having to copy 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 an ME5 storage DAS or SAN 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.

For PowerVault volumes mapped to a Windows Server, the host passes a trim and unmap command to PowerVault when files are deleted. Within a few minutes, the PowerVault storage pool reflects the additional free capacity.

The ability to recover deleted disk space on PowerVault 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 PowerVault given these conditions:

- The Windows Server operating system must be version 2012 or newer (ME5 supports Server 2016 and newer).
- Volumes must be **basic** disks that are formatted as NTFS volumes. Trim and unmap is not supported with other formats such as ReFS.

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. The default settings should not be changed 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 ME5 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 ME5 snapshots of boot volumes which will conserve ME5 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, a CSV may be dedicated 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 ME5 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.
- At least one domain controller should run independent of SAN or DAS 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, SAN, or DAS resources. Place critical management tools on this resource to aid with day-to-day administration, troubleshooting, and recovery.

Domain controller placement

Avoid placing all your 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.

A target port on ME5 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 ME5 SAN configurations, configure all available front-end data (target) ports. Use of 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: Modifying queue depth settings is not advised unless there is a specific reason to do so. Queue depth changes should be tested 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 ME5 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.

ME5 snapshots and storage replication with Hyper-V

Overview

ME5 storage 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

ME5 snapshots are space-efficient as they consume no additional storage space unless they are mapped to a host or VM and new data is written.

For general use cases and best practices regarding the configuration of snapshots and replication, see the *Dell PowerVault ME5 Administrator's Guide* at Dell Technologies Support.

ME5 storage snapshots and replication allow administrators to do the following in a Hyper-V environment:

- Replicate Hyper-V volumes and snapshots to another location for DR or archive purposes.
- Perform manual recovery of hosts and VMs at a primary or alternate location.
- Provision an isolated test environment that matches a production environment.

- Provision new boot-from-SAN host servers quickly from a snapshot that contains a system-prepared (sysprep) gold image.

Crash-consistent and application-consistent snapshots

ME5 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 an ME5 snapshot.

- Leverage application-native tools to place a workload in a consistent state temporality.
- Stop application services temporarily.
- Power off the host or VM that is hosting the workload. This method is often disruptive and impractical. This method is used to create a gold image of a system-prepared (sysprep) host or VM after it is powered off.
- Leverage 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 an ME5 snapshot. Then revert the host, VM, or workload to its active state.

If possible, leverage scripting and automation tools to orchestrate a process that performs these steps automatically. Orchestration reduces administrative overhead and helps eliminate mistakes due to human error.

Guest VM recovery with ME5 snapshots

Recover Hyper-V VMs to a previous point in time by using consistent or crash-consistent ME5 snapshots. Snapshots can be used to create cloned copies of VMs in an isolated environment at the same or a different location.

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

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

Option 1: Recover the existing data volume on the host that contains the VM configuration and virtual hard disks by using an ME5 snapshot.

- If the data volume contains only one VM, recovery with a snapshot rollback may be practical. If the data volume contains multiple VMs, it will 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 one VM.
- The recovery VM can power up without any additional configuration or recovery steps required.

- It is essential to document the LUN number, disk letter, or mount-point information for the volume to be recovered, before starting the recovery.

Option 2: Map a snapshot containing the VM configuration and virtual hard disks to the host as a new volume, in a side-by-side fashion using a new drive letter or mount point. Recover the VM by manually copying the virtual hard disks from the recovery snapshot to the original location.

- Delete, move, or rename the original virtual hard disks.
- After copying the recovered virtual hard disks to their original location, they must be renamed and Hyper-V manager must be used to reassociate them with the guest VM. The guest VM can now start without permissions errors.
- If the virtual hard disks are large, copying data may not be practical. In this case, the original VM can be deleted, and the recovery VM imported or created as a new VM directly from the recovery volume. After the recovery, the original data volume can be unmapped from the host if no longer needed.
- This method also facilitates recovery of a subset of data from a VM by mounting a recovery VHD as a volume on the host server temporarily.

Option 3: Map the recovery snapshot to a different Hyper-V host and recover the VM there. Import the VM configuration or create a VM that points to the virtual hard disks on the recovery volume.

- Use this option when the original VM and the recovery VM both need to be online simultaneously. Ensure the VMs are isolated from each other to avoid name or IP conflicts, or a split-brain situation with data writes.
- Recover the VM on another host when the original host server is no longer available due to a host failure.

Before beginning a VM recovery, record VM configuration details such as the number of virtual CPUs, memory, virtual networks, and IP addresses. If importing a VM configuration fails, a new VM will need to be created using this information.

Recover guest VM on a cluster shared volume

The process of using ME5 snapshots to recover guest VMs that reside on a CSV is like the process of recovering a guest VM to a stand-alone host. However, recovering a VM from a snapshot of a CSV may require changing the disk signature first.

Windows Servers assign each volume a unique disk ID (or signature). For example, the disk ID for an MBR disk is a hexadecimal number such as 045C3E2F4. All volumes mapped to a server must use a unique disk ID.

When an ME5 snapshot is taken of a Windows or Hyper-V volume, the snapshot is an exact point-in-time copy, which includes the Windows disk ID. Recovery volumes based on snapshots will have the same disk ID as the original volume.

With stand-alone Windows or Hyper-V servers, disk ID conflicts are avoided. Stand-alone servers can automatically detect a duplicate disk ID and change it dynamically. However, host servers are not able to dynamically change conflicting disk IDs when disks are configured as CSVs due to the behavior of Windows Server clustering.

When attempting to map a copy (snapshot) of a CSV as an additional volume in that same cluster, the recovery volume will create a disk ID conflict. Disk ID conflicts can be service-affecting.

There are two methods to work around the duplicate disk ID issue:

Option 1: Map the recovery volume (snapshot) containing the CSV to another host that is outside of the cluster. Copy the guest VM files over the network to recover the guest VM.

Option 2: Map the recovery volume to another Windows host outside of the cluster and use Diskpart.exe or PowerShell to change the disk ID. Once the ID is changed, remap the recovery volume to the cluster. The following steps demonstrate how to use diskpart.exe to change a disk ID.

Change a CSV disk ID with Diskpart

Follow these steps to change a disk ID. PowerShell can also be used.

1. Access the stand-alone Windows host that the recovery volume (snapshot) containing the CSV will be mapped to.
2. Open a command window with administrator rights.
3. Run the following commands:

```
diskpart
list disk
```

4. Make note of the current list of disks (in this example, Disk 0, 1, 2, and 3).

```
Administrator: Cmd Prompt - diskpart

C:\Windows\system32>diskpart

Microsoft DiskPart version 10.0.20348.1

Copyright (C) Microsoft Corporation.
On computer: S1356


DISKPART> list disk

   Disk ###  Status              Size       Free      Dyn  Gpt
   -----  -
   Disk 0    Online              139 GB     1024 KB
   Disk 1    Reserved            4768 MB     1024 KB
   Disk 2    Reserved            512 GB     1024 KB
   Disk 3    Reserved            605 GB     1024 KB

DISKPART> _
```

5. Map the recovery volume containing the CSV to this host.
6. Run the following command:


```
rescan
```
7. Use Disk Management on the host server to bring the recovery volume online.



8. Return to Diskpart and run the following command:

```
list disk
```

9. The new disk (Disk 4 in this example) should now be listed. Usually, the bottom disk will be the new disk.

```
Administrator: Cmd Prompt - diskpart

DISKPART> list disk

Disk ###  Status         Size         Free         Dyn  Gpt
-----  -
Disk 0    Online         139 GB       1024 KB
Disk 1    Reserved      4768 MB      1024 KB
Disk 2    Reserved      512 GB       1024 KB
Disk 3    Reserved      605 GB       1024 KB
Disk 4    Online         186 GB       1024 KB

DISKPART> _
```

10. Run the following command to select Disk 4 (in this example) as the new disk:

```
select disk 4
```

11. Run the following command to view the current ID for the disk:

```
uniqueid disk
```

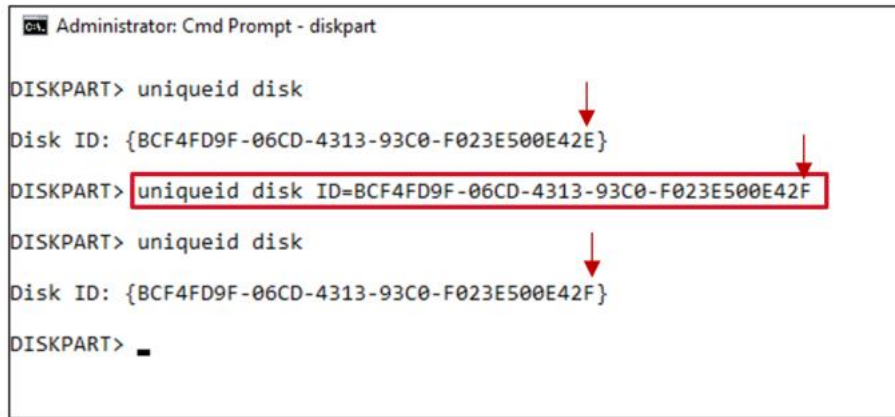
12. Change the disk ID by running this command:

```
uniqueid disk ID=<newid>
```

- a. For example, increment the last character of the string by one.
 - i For an MBR disk, the ID is an eight-character string in hexadecimal format.
 - ii For a GPT disk, the ID is a Globally Unique Identifier (GUID) in hexadecimal format.

13. To verify the new disk ID, run this command:

```
uniqueid disk
```



14. Now that the disk has a new signature, unmap it from the stand-alone host server and remap it to the cluster. The disk will no longer cause a disk ID conflict.
15. Mount the volume so it is accessible.
16. Recover the guest VM.

Create a test environment with ME5 snapshots

In addition to VM recovery, ME5 snapshots can be used to quickly create test or development environments that mirror a production environment. When volumes containing VMs are replicated to another location, it is easy to do so at a different location.

Note: To avoid IP, MAC address, or server name conflicts, copies of existing VMs that are brought online should be isolated.

The procedure to use a snapshot 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.

Migrate guest VMs with ME5 storage

Microsoft Hyper-V provides native tools to migrate VMs. Use of native Hyper-V tools is preferred. Most commonly, VMs are moved within a cluster by using Live Migration.

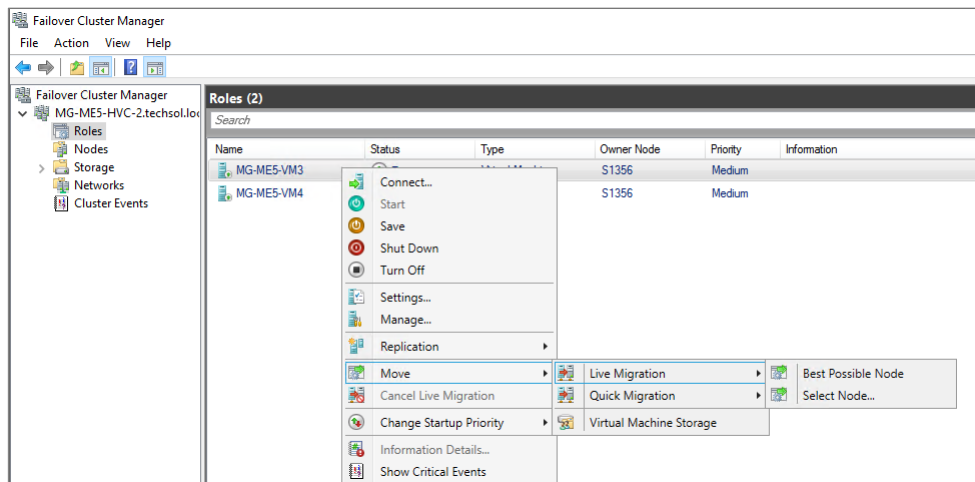


Figure 18. Hyper-V move and migration options

Moving a VM by remapping its underlying ME5 volumes to a different host or cluster may be a better choice in some situations. For example, using Hyper-V tools to move or migrate a large VM over the network may require considerable time and may not be practical. A storage-based method (remapping a volume) will involve down time, but will be quicker than waiting for a long network copy process to finish. Remapping volumes to move VMs will not consume network bandwidth.

1. Plan for a maintenance window.
2. Make a backup of the VM and its workload.
3. Take the VM and workload offline.
4. Unmap the SAN volume containing the VM configuration and virtual hard disks. A snapshot of the volume can also be used.
5. Map the volume to the new target host or cluster. Moving a VM can also be completed by using a replicated volume or snapshot at another location.
6. Mount the volume and bring the host and workload online. Verify correct operation.

Boot from SAN

Overview

Windows Server Hyper-V hosts support local boot and boot from SAN. Boot from SAN requires the use of a supported iSCSI or FC HBA that supports boot from SAN. Boot from SAN disks should be assigned LUN ID 0.

A boot from SAN disk supports MPIO. After staging a Windows Server to an MPIO-capable boot from SAN disk, install and configure the MPIO feature.

Boot from SAN allows similar hosts to be provisioned quickly by using a system-prepared (sysprep) gold image. Replicated snapshots of boot from SAN Hyper-V hosts allow for fast host recovery at an alternate location when both sites use similar host hardware.

Clustering and Boot from SAN

Boot from SAN is not preferred for large Hyper-V clusters on PowerVault.

Hosts that are configured to boot from SAN cannot be assigned to a host group in PowerVault Manager. However, cluster volumes can still be mapped to a group of clustered Hyper-V nodes in PowerVault Manager, even if they are not in a host group.

After mapping a shared volume to multiple Hyper-V nodes, verify that the LUN number is consistent on all nodes. If the LUN number is not consistent, use PowerVault Manager to change it. Click **Edit All** and specify the new LUN number. Then click **Apply**.

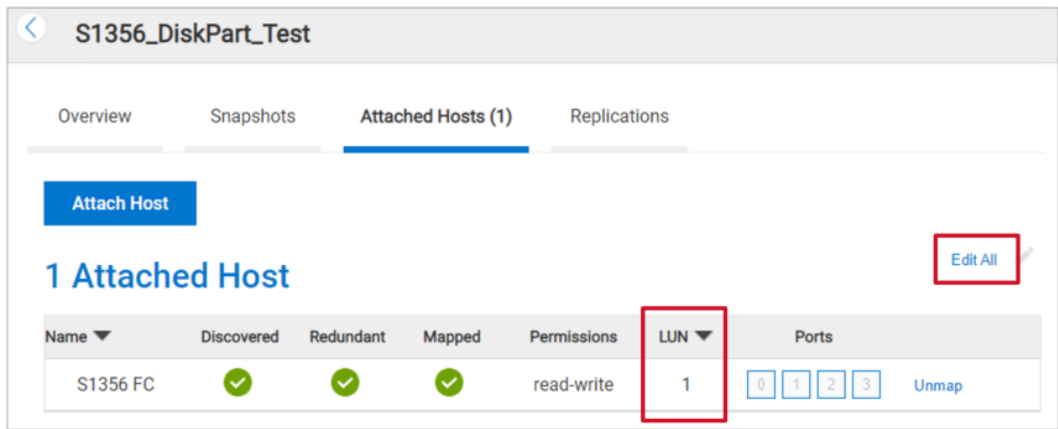


Figure 19. Change the LUN number for a volume

Perform a **Rescan Disk** on the host if the LUN number is changed in PowerVault Manager. If the LUN number does not change on the host after a rescan, reboot the host.

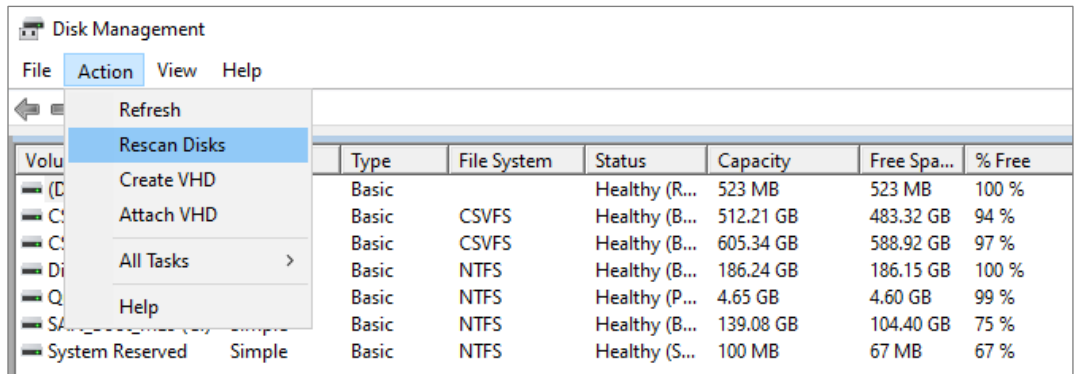


Figure 20. Rescan the disks on a host if a LUN number is changed

When a volume is mapped to a host group, or a group of hosts in PowerVault Manager, the wizard should assign a consistent LUN ID for all hosts automatically. In this example, a cluster shared volume is mapped to a host group with two member servers. The consistent LUN number is 31 in this example.

The screenshot shows a storage management interface for a host group named 'MG-ME5-HVC-2_FC_CSV1'. It features tabs for 'Overview', 'Snapshots', 'Attached Hosts (2)', and 'Replications'. An 'Attach Host' button is visible. Below the tabs, there is a section titled '2 Attached Hosts' with an 'Edit All' link. A table lists the attached hosts with columns for Name, Discovered, Redundant, Mapped, Permissions, LUN, and Ports. The LUN column for both hosts is highlighted with a red box.

Name	Discovered	Redundant	Mapped	Permissions	LUN	Ports
S1356 FC	✓	✓	✓	read-write	31	0 1 2 3 Unmap
S1357 FC	✓	✓	✓	read-write	31	0 1 2 3 Unmap

Figure 21. Cluster volume mapped to a host group consisting of two nodes

Conclusion

Careful planning, adherence to best practices, and testing are essential for a successful deployment of Microsoft Hyper-V on Dell PowerEdge ME5 storage. ME5 storage is well suited to host high-density high-demand Hyper-V virtual workloads. ME5 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