

# Dell EMC Unity Storage with Microsoft Hyper-V

Hybrid and All-flash arrays

## Abstract

This white paper provides best practices guidance for configuring Microsoft Hyper-V to perform optimally with Dell EMC Unity Hybrid and All-flash arrays.

June 2021

## Revisions

Date	Description
July 2017	Initial release for Dell EMC Unity OE version 4.2
October 2020	Remove reference to Dell EMC Storage Integrator (ESI) as it is end-of-life
June 2021	Update for Dell EMC Unity OE version 5.1

## Acknowledgments

Author: Marty Glaser

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.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

This document may contain language from third party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third party content is updated by the relevant third parties, this document will be revised accordingly.

Copyright © 2017–2021 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. Other trademarks may be trademarks of their respective owners. [7/7/2021] [Best Practices] [H16388.2]

# Table of contents

Revisions .....	2
Acknowledgments .....	2
Table of contents .....	3
Executive summary .....	5
Audience .....	5
1 Introduction .....	6
1.1 Essential guidelines .....	6
2 Storage features and configuration .....	7
3 Hyper-V with Dell EMC Unity overview .....	8
3.1 Hyper-V .....	8
3.2 Hyper-V with Dell EMC Unity .....	10
4 Optimize Hyper-V for Dell EMC Unity .....	11
4.1 Hyper-V general best practices .....	11
4.2 Hyper-V guest VM generations .....	13
4.3 Virtual hard disks .....	15
4.4 Present Dell EMC Unity storage to Hyper-V .....	19
4.5 Offloaded Data Transfer (ODX) .....	27
4.6 Placement of page files .....	27
4.7 Active Directory domain controller placement .....	27
4.8 Dell EMC Unity data reduction and Hyper-V .....	28
4.9 Queue depth best practices for Windows Server Hyper-V .....	29
5 Dell EMC Unity snapshots and thin clones with Hyper-V .....	30
5.1 Crash-consistent and application-consistent snapshots .....	30
5.2 Using Dell EMC Unity snapshots to recover guest VMs .....	31
5.3 Use Dell EMC Unity thin clones to create a test environment .....	36
5.4 Leverage Dell EMC Unity to create gold images .....	36
5.5 Using Dell EMC Unity for Hyper-V VM migration .....	39
6 Boot-from-SAN for Hyper-V .....	41
6.1 Boot-from-SAN advantages .....	41
6.2 When to use boot from local disk .....	41
6.3 Configure a Hyper-V host server to boot from SAN .....	41
7 Dell EMC Unity and NAS for Hyper-V .....	43
7.1 Present Dell EMC Unity SMB file shares to Hyper-V clusters .....	43
A Technical support and additional resources .....	48

A.1 Related resources ..... 48

## Executive summary

Microsoft® Hyper-V® and Dell EMC™ Unity storage are feature-rich solutions that together provide a diverse range of configuration options to solve key business objectives such as performance and resiliency.

This paper delivers straightforward guidance to customers deploying Hyper-V with Dell EMC Unity Hybrid and All-Flash storage systems, including the Unity 380/F, 480/F, 680/F, and 880/F models. It builds upon the existing [virtualization documentation provided by Microsoft](#), and the Unity documentation library at [Dell EMC Unity family technical white papers and videos](#). For more information about SMI-S integration, see the [Dell EMC Unity Family SMI-S Programmer's Guide](#).

These resources are prerequisite reading and serve as the primary source for understanding and configuring your Hyper-V and Unity environment. This document provides supplemental information.

General Hyper-V best practices that are storage agnostic are not covered in detail this guide. Knowledge of general Hyper-V best practices is assumed.

Dell EMC encourages following best practice recommendations. Some recommendations may not apply to all environments. For questions about the applicability of specific recommendations in your environment, contact your Dell EMC representative.

## Audience

This document is for Dell EMC customers, partners, and employees who want to learn more about best practices for Microsoft Hyper-V with Dell EMC Unity Hybrid and All-flash storage systems. It assumes the reader has working knowledge of Dell EMC Unity and Hyper-V.

We welcome your feedback along with any recommendations for improving this document. Send comments to [StorageSolutionsFeedback@dell.com](mailto:StorageSolutionsFeedback@dell.com).

# 1 Introduction

Best practices are commonly used when organizations face design and configuration choices affecting the cost, complexity, expandability, performance, availability, and resiliency of the environment or workload. Due to numerous considerations, a design that works well for one environment may not be ideal for another. There is no one-size-fits-all answer. However, following best practices can help you make informed choices that best suit your environment to ensure the best return on your investment.

The guidance in this document is focused on the best design when building a new solution. Legacy systems that are performing well and have not reached their life expectancy may not follow current best practices. Often, the best course of action is to run legacy configurations until they reach their life expectancy. It might be too disruptive or costly to make changes outside of a normal hardware progression or upgrade cycle.

Dell Technologies recommends upgrading to the latest technologies and adopting current best practices at key opportunities such as when upgrading or replacing infrastructure.

## 1.1 Essential guidelines

Dell EMC Unity and Hyper-V environments should follow established best practices and adhere to the following essential guidelines:

- Minimize complexity to ease administrative overhead.
- Ensure redundancies for critical workloads to protect against single points of failure.
- Ensure that disaster recovery time objective (RTO) and recovery point objective (RPO) goals can be achieved.
- Leverage the latest features in current Dell EMC Unity operating environment (OE), and Hyper-V versions.
- Evaluate a solution in a test environment before deploying in production where possible.
- Avoid or minimize the use of legacy or depreciated hardware or software especially if support or patches are no longer available.
- Avoid proprietary hardware or software that might limit future expansion, upgrades, or migration.
- Plan for scalability to accommodate anticipated future growth and other business needs over the expected life of the hardware.

## 2 Storage features and configuration

Dell EMC Unity arrays provide a unified block and file storage platform to optimize simplicity, performance, and efficiency to meet a wide range of business needs.

For information about specific Unity models, see [Dell EMC Unity XT unified storage](#).

For more information about Unity features and configurations, see [Dell EMC Unity family technical white papers and videos](#). Topics covered include the following:

- Dynamic Pools
- Data Efficiency (data reduction, compression)
- Data Protection (replication, snapshots, thin clones, MetroSync)
- NAS and File
- Migration
- Monitoring
- Performance
- Platform and Hardware
- Security
- Serviceability

## 3 Hyper-V with Dell EMC Unity overview

### 3.1 Hyper-V

The Microsoft Windows Server platform leverages Hyper-V for virtualization technology. Initially offered with Windows Server 2008, Hyper-V has matured with each release to include many new features and enhancements.

Microsoft Hyper-V is a mature, robust, proven virtualization platform. Hyper-V presents the physical host server hardware resources in an optimized and virtualized manner to one or more resources such as a guest virtual machine (VM). Hyper-V hosts are referred to as nodes when clustered.

Virtualization greatly enhances the utilization of physical host and node hardware such as processors, memory, NICs, and power. Virtualization allows many VMs to share these physical resources concurrently.

Hyper-V Manager, Failover Cluster Manager® (FCM), Microsoft System Center® (SC) Virtual Machine Manager® (VMM), and PowerShell®, are commonly used tools. They offer administrators great control and flexibility for managing host and VM resources.

#### 3.1.1 Hyper-V management tools

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 (WAC) (see Figure 1) is a free, centralized server-management tool from Microsoft. WAC consolidates many common in-box and remote-management tools to simplify managing server environments from one interface.

WAC is a locally installed stand-alone client that is HTML5-based and browser-accessible. WAC is also an extensible platform allowing third parties to develop integrations for their own products or solutions. Administration and monitoring of Unity storage from WAC is not supported.

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




Figure 1 Use Windows Admin Center to manage Hyper-V servers

WAC is now the recommended tool for managing Windows Server environments. However, it may not have full feature parity with the traditional management tools it replaced. Continue to use Hyper-V Manager, FCM, SCVMM, and PowerShell if the wanted functionality is not in WAC.

This document includes configuration examples that use a combination of traditional tools and WAC.




Figure 2 Server Manager > Add Roles and Features wizard

Many core Hyper-V features (such as dynamic memory) are storage agnostic and are not covered in detail in this guide. To learn more about Hyper-V features, there are many online resources available such as the [Microsoft Tech Community](#) and the [Microsoft technical documentation library](#).

## 3.2 Hyper-V with Dell EMC Unity

Unity is a fully unified storage solution. Unity arrays can be configured to present storage through traditional SAN-based block storage using Fibre Channel or iSCSI, and through file-based SMB file shares to Hyper-V environments. SMB is an alternative to traditional block-based SAN storage. When infrastructure to support FC or iSCSI is in place, customers can continue to realize the return on their investment by using these robust transports with Dell EMC Unity. SMB file shares can be used instead of, or concurrently with, FC and iSCSI.

Core Dell EMC Unity features such as thin provisioning, data reduction, compression, snapshots, and replication work seamlessly in the background, regardless of the platform. Often the default settings for these features work well with Hyper-V or serve as good configuration starting points. This document covers additional configuration or tuning steps to enhance performance, utilization, or resiliency.

### 3.2.1 Supported versions

See the current documentation for your Unity hardware and Unity OE to verify compatibility with Windows Server and Hyper-V versions.

Unity is supported with long-term servicing channel (LTSC) releases of Windows Server. Use of semiannual channel (SAC) releases of Windows Server with Unity 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.

**Note:** Unity support for different versions of Windows Server and the Hyper-V role may change over time. Always consult the latest documentation and release notes for your Unity OE and hardware to verify compatibility.

## 4 Optimize Hyper-V for Dell EMC Unity

The following sections include additional configuration or tuning steps to enhance the performance, utilization, and resiliency of Hyper-V with Dell EMC Unity.

### 4.1 Hyper-V general best practices

For general best practices for Hyper-V not specific to storage, see the online resources available at [Microsoft](#).

The following provides a high-level summary of common Hyper-V tuning steps:

- Minimize or disable unnecessary hardware devices and services to free up CPU cycles and reduce power consumption.
- Schedule tasks such as periodic maintenance, backups, malware scans, and updating to run after hours. Stagger start times when operations overlap and are CPU or I/O intensive.
- Tune application workloads to reduce or eliminate unnecessary processes or activity.
- Leverage Microsoft PowerShell or other scripting tools to automate step-intensive repeatable tasks, to ensure consistency, avoid human error, and to reduce administration time.

In addition to these general guidelines, the following subsections provide detailed configuration information.

#### 4.1.1 Hyper-V integration services

Guest integration services are virtualization-aware drivers that are installed on a guest VM to optimize its virtual hardware interaction with the physical host hardware and storage. This hardware includes Unity storage. Installing these drivers is typically the first step for optimizing VM performance. 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.

Installing and updating integration services are commonly overlooked steps that can ensure overall stability and optimal performance of guest VMs. Although newer Windows-based operating systems and some enterprise-class Linux-based operating systems come with integration services, updates may still be required. New versions of integration services may become available as physical Hyper-V hosts and nodes are updated.

With Hyper-V 2012 R2 and prior versions, when configuring and deploying a new VM, the configuration process does not prompt to install or update integration services. Also, the process to install integration services with Hyper-V 2012 R2 and prior versions is obscure and so it is explained further in this section. With Windows Server 2016 and newer versions, Windows Updates automatically updates integration services on a Windows guest VM, reducing administration overhead.

One common issue occurs when VMs are migrated from an older physical Hyper-V host or cluster to a newer host or cluster. For example, this issue could occur when migrating from Windows Server 2008 R2 Hyper-V to Windows Server 2012/R2 Hyper-V. The integration services are not updated automatically, which can result in degraded performance. This issue could waste administration time to troubleshoot a suspected but nonexistent performance issue with Unity or the storage fabric.

Aside from performance problems, there is another way to tell that integration services are outdated or not present on a Windows VM. One key indicator is the presence of unknown devices in Device Manager for the VM (see Figure 3).




Figure 3 Unknown devices that are listed for a guest VM indicate missing or outdated integration services

For versions of Hyper-V before 2016, use Hyper-V Manager to connect to a VM. Click **Action > Insert Integration Services Setup Disk** (an ISO file) as shown in Figure 4. Follow the prompts in the guest VM console to complete the installation.

---

**Note:** Mounting the integration services ISO is not supported with Windows Server 2016 Hyper-V and newer. With these versions, integration services are provided exclusively as part of Windows Updates.

---




Figure 4 Mount the integration services setup disk ISO

Use WAC (see Figure 5), FMC (see Figure 6), or PowerShell to verify the version of integration services for a VM.




Figure 5 Verify integration services version with Windows Admin Center




Figure 6 Verify integration services version with Failover Cluster Manager

Verification can also be performed using PowerShell, as shown in the following example:

```
PS C:\Windows\system32> get-VM | Select-Object name, integrationservicesversion
Name      IntegrationServicesVersion
-----
MG-VM12a  6.3.9600.18080
MG-VM12b  6.3.9600.18080
MG-VM12c  6.3.9600.18080
MG-VM12d  6.3.9600.18080
```

## 4.2 Hyper-V guest VM generations

When Windows Server 2012 R2 Hyper-V was released, Microsoft designated all existing VMs as **generation 1**. This designation differentiates them from a new classification of VMs that could be created as **generation 2**. Unity supports either generation of VM. If the workload supports it, the recommendation from Microsoft and Dell Technologies is to configure new guests as generation 2. See Figure 30 and Figure 31.




Figure 7 New virtual machine wizard in Windows Admin Center




Figure 8 New virtual machine wizard in Hyper-V manager

This configuration is important because generation 2 guests use 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. Also, one of the most significant changes with generation 2 guests is the elimination of the dependency on virtual IDE controllers for the boot disk. Generation 1 VMs require the boot disk to use a virtual IDE disk controller. Generation 2 guests instead use virtual SCSI controllers for all disks. Virtual IDE is not a supported option with generation 2 VMs.

### 4.2.1 Convert VMs to a newer generation

The warning message shown in the wizard (Figure 8) indicates that the VM generation cannot be changed once a VM has been created. However, third-party tools are available to convert VMs. Dell Technologies does not endorse any specific methods for converting VMs (use at your own risk).

---

**Tip:** Use Unity snapshots and thin clones to create a test environment for a production VM. In a test environment, VM conversion can be attempted without affecting the production environment.

---

Follow these steps to migrate a VM:

1. Create another instance of the generation 1 VM in an isolated test environment.

**Tip:** Use Unity snapshots, thin clones, or native Hyper-V tools to create a copy of the VM.

2. Create an instance of a generation 2 VM in the isolated test environment.
3. Migrate roles, features, workloads, and data from the generation 1 VM to the generation 2 VM.
4. Test functionality.
5. If everything works as expected, plan a maintenance window to perform the migration on the production VM.

## 4.3 Virtual hard disks

A virtual hard disk (VHD) contains data blocks and is stored as a regular Windows file on the host system. VHD files end with a **vhd**, **vhdx** or **vhds** extension depending on the type of VHD. All VHD formats are supported with Unity.




Figure 9 A boot virtual hard disk file (vhdx) for a Windows Server 2019 guest VM on a cluster shared volume

### 4.3.1 Virtual hard disk format

There are three 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 2,048 GB.
- **VHDX** is supported with Windows Server 2012 Hyper-V and newer. The VHDX format offers better resiliency if there is a power loss, better performance, and supports a maximum size of 64 TB. VHD files can be converted to the 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. The VHDS format is used when two or more guest VMs share access to the same VHD in support of VM clustering or other high-availability configurations.




Figure 10 Different formats available for virtual hard disks

### 4.3.2 Virtual hard disk type

In addition to choosing the format type, a VHD can be designated as fixed, dynamically expanding, or differencing. Since Unity uses thin provisioning, only the data that is written to a virtual hard disk, regardless of the disk type, consumes space on Unity. Determining the best disk type is more about workload requirements than how it impacts storage utilization on Unity.




Figure 11 Select a virtual hard disk type

If the environment is designed correctly, the **dynamically expanding** disk type works well for most workloads. For workloads that generate high I/O, such as Microsoft SQL Server® databases, Microsoft recommends using the **fixed size** disk type.

Figure 12 shows a fixed virtual hard disk that consumes the full amount of space from the perspective of the host server. For a dynamic virtual hard disk, the space that is consumed is slightly larger than amount of data on the virtual disk. The small amount of extra space consumed is due to metadata. This VHD type is more space efficient from the perspective of the host. From the perspective of the guest VM, either type of virtual hard disk appears as a full 60 GB of available space.




Figure 12 Fixed and dynamic virtual hard disk comparison

Unity supports any VHD type. From the perspective of the host server, there are some best practice performance and management considerations to consider when choosing the VHD type for your environment.

- Fixed-size VHDs:
  - These VHDs are recommended for workloads with a high level of disk activity, such as Microsoft SQL Server, Microsoft Exchange®, or operating-system page or swap files. For many workloads, the performance difference between fixed and dynamic is negligible.
  - When formatted, these VHDs consume the allocated space on the host server volume.
  - They are less susceptible to fragmentation at the host level.
  - They take longer to copy because the file size is the same as the formatted size.
- Dynamically expanding VHDs:
  - These VHDs are recommended for most workloads, except in cases with high disk I/O.
  - When initially formatted, the VHDs consume little space on the host, and expand only as new data is written to them by the guest VM or its workload.
  - As they expand, they require a small amount of additional CPU and I/O overhead temporarily. This overhead usually does not impact the workload except in cases where I/O demand is high.
  - They are more susceptible to fragmentation at the host level.
  - They require less time to copy than fixed VHDs.
  - They allow the physical storage on a host server or cluster to be over overprovisioned.

**Tip:** Configure alerting to avoid running physical storage out of space when using dynamically expanding VHDs.
- Differencing virtual hard disks:
  - These VHD types offer storage savings by allowing multiple Hyper-V guest VMs with identical operating systems to share a common boot virtual hard disk.
  - They are practical for limited use cases such as a virtual desktop infrastructure (VDI) deployment.
  - All children must use the same VHD format as the parent.
  - Reads of unchanged data reference the differencing VHD.
  - New data is written to a child VHD.

A native Hyper-V-based checkpoint (or snapshot; not to be confused with a Unity snapshot) of a Hyper-V guest VM creates a differencing virtual hard disk (avhdx). This differencing VHD freezes data that has changed since the last checkpoint. Each additional checkpoint of the VM creates another differencing virtual hard disk, maintained in a hierarchical chain. Note these additional points about Hyper-V checkpoints:

- Use of native Hyper-V-based checkpoints of a Hyper-V guest VM can negatively impact read I/O performance. The data is spread across the VHD and one or more differencing disks. This situation increases read latency.
- Longer chains of differencing virtual hard disks are more likely to negatively impact read performance. It is a best practice to avoid using checkpoints, and to keep native Hyper-V-based checkpoints to a minimum if a workflow or process requires them.
- Administrators can use Unity snapshots, thin clones, and replication for archive and recovery of Hyper-V guest VMs to avoid using native Hyper-V checkpoints.

### 4.3.3 Virtual hard disks and thin provisioning with Unity

Disk space utilization on Unity is optimized regardless of the type of virtual hard disk used due to the advantages of thin provisioning. For all virtual hard disk types, only new data created by a guest VM or workload consumes space on Unity.

The following example describes a 100 GB Unity volume that is presented to a Hyper-V host that contains two 60 GB virtual hard disks. The host volume is overprovisioned in this case to demonstrate behavior, but not as a general best practice. One virtual hard 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 (also see Figure 13 for a diagram):

**60 GB fixed disk + 15 GB of used space on the dynamic disk = 75 GB**

---

**Note:** The host server reports the entire capacity of a fixed virtual hard disk as consumed.

---




Figure 13 Thin provisioning with Unity

Compare this example to how Unity reports storage utilization on the same volume:

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

Also, due to Unity data reduction and compression technology, the space that is consumed on Unity is less than what is reported on the host.

---

**Note:** Dynamic and fixed virtual hard disks achieve the same space efficiency on Unity due to thin provisioning. Other factors such as the I/O performance of the workload would be primary considerations when determining the type of virtual hard disk used.

---

#### 4.3.4 Overprovisioning with dynamic virtual hard disks

With dynamic virtual hard disks and thin provisioning, there is a risk of the host volume or a storage container on Unity running out of space. See Figure 13 for an example. If the dynamic disk used by VM2 on the host volume expands far enough, it would consume the free space on the host volume and negatively impact VM1 and VM2. From the perspective of VM2, it still detects 20 GB of free space but cannot use it because the underlying host volume would be full. To resolve this issue, an administrator must move the VHD for VM1 or VM2 elsewhere to free up space on the 100 GB host volume. Otherwise, the administrator could expand the host volume. In either case, it may be difficult to identify the root cause of the problem, and resolution may require a service interruption.

To mitigate risks, consider the following best practices:

- Create a Hyper-V volume on the physical host that is large enough so that expanding dynamic VHDs do not fill the host volume to capacity. Creating large Hyper-V host volumes does not negatively impact space efficiency on Unity because of the benefits of thin provisioning. However, volumes should not be sized larger than needed as a best practice.
- Hyper-V-based VM checkpoints (snapshots) create differencing VHDs on the same physical volume as the VM parent virtual hard disk. Allow adequate overhead on the host volume for the extra space consumed by differencing VHDs if checkpoints are used.
- At the host level, set up monitoring on overprovisioned volumes. If a threshold is exceeded (such as more than 90 percent full), an alert is generated with enough lead time to allow for remediation.
- At the Unity level, configure alerting so remediation steps can be taken to address a storage-pool-capacity alert.

## 4.4 Present Dell EMC Unity storage to Hyper-V

There are several different ways to present Dell EMC Unity storage to Windows Server Hyper-V hosts, nodes, and VMs.

Typically, an environment is configured to use a preferred transport (FC or iSCSI) when it is built and will be part of the infrastructure core design. When deploying Hyper-V to existing environments, the existing transport is typically used. Deciding which transport to use is based on customer preference and factors such as size of the environment, cost of the hardware, and the required support expertise.

It is not uncommon, especially in larger environments, to have more than one transport available. Multiple transports might be required to support collocated but diverse platforms with different transport requirements.

For more information about configuring hosts to access FC or iSCSI Unity storage see the [Configuring Hosts to Access Fibre Channel \(FC\) or iSCSI Storage](#) white paper.

### 4.4.1 Supported transports

- Dell EMC Unity storage can be presented to physical Hyper-V hosts and cluster nodes using Fibre Channel, iSCSI, or SMB file shares.
- Boot-from-SAN is supported if the physical server has a compatible FC or iSCSI device installed.
  - Booting to a local onboard disk in the host and leveraging Unity for shared storage such as a cluster shared volume is also a common configuration.
- Dell EMC Unity storage can be presented directly to Hyper-V guest VMs using:
  - In-guest iSCSI (direct-attached storage)
  - Hyper-V pass-through disks (a legacy configuration that should be avoided if possible)
  - SMB file shares

### 4.4.2 Dell EMC Unity support for front-end SAS

Dell EMC Unity arrays do not support front-end SAS connectivity to host servers.

### 4.4.3 Multiprotocol support

Although Dell EMC Unity arrays support mixed transports, there is limited Microsoft support. For each LUN mapped to a Hyper-V host or node, use a single transport. Each transport by itself should present adequate bandwidth and redundant paths that leverage MPIO. Using multiple transports increases design complexity unnecessarily and can introduce unpredictable service-affecting I/O behavior in path failure scenarios.

### 4.4.4 MPIO best practices

The Windows Server operating system and Hyper-V (2008 and newer) natively support MPIO with the integrated device-specific module (DSM) that is bundled with the operating system. The basic functionality offered with the Microsoft DSM is supported with Unity. Dell Technologies recommends the use of Dell EMC PowerPath™ (licensed, at additional cost) for MPIO management on server hosts and VMs instead of the Microsoft DSM. Dell EMC PowerPath, a server-resident software solution designed to enhance performance and application availability, provides the following:

- Combines automatic load balancing, path failover, and multiple-path I/O capabilities into one integrated package
- Enhances application availability by providing load balancing, automatic path failover, and recovery functionality
- Supports servers, including cluster servers, connected to Dell EMC and third-party arrays

For complete guidance on installing and configuring PowerPath including guidance for Windows Server Hyper-V, see the *PowerPath Installation and Administration Guide for Microsoft Windows* at [Dell EMC Support](#).

For more information about MPIO setup, including changing registry values, see the [Configuring Hosts to Access Fibre Channel \(FC\) or iSCSI Storage](#) white paper.

Windows and Hyper-V Hosts will default to **Round Robin with Subset** with Dell EMC Unity storage, unless a different default MPIO policy is set on the host by the administrator. The default setting is recommended.




Figure 14 Verify MPIO settings (Microsoft DSM shown)

Note the following:

- The **active/optimized** paths are associated with the Dell EMC Unity storage processor (SP) that owns the volume. The **active/unoptimized** paths are associated with the other SP.
- If each SP has four FC paths configured (as shown in the Figure 14 example), each volume that is mapped should list eight total FC paths: 4 that are **active/optimized**, and 4 that are **active/unoptimized**.

Best practices recommendations include:

- Changes to MPIO registry settings for FC or iSCSI (such as time-out values) should not be made unless:
  - Directed by Dell EMC documentation (see [Configuring Hosts to Access Fibre Channel \(FC\) or iSCSI Storage](#)).
  - Directed by Dell EMC support to solve a specific problem.
- Configure all available data ports on a Dell EMC Unity array to use your preferred transport (FC or iSCSI) to optimize array CPU utilization and maximize performance.
- Verify that current versions of software are installed (such as boot code, firmware, and drivers) for all components in the data path:
  - Dell EMC Unity OE
  - Data switches
  - Host bus adapters (HBAs), network interface cards (NICs), and converged network adapters (CNAs)
- Verify that all hardware is supported. See the [Dell EMC Unity Simple Support Matrix](#) for software and hardware support information.

## 4.4.5 Guest VMs and in-guest iSCSI

Dell EMC Unity supports in-guest iSCSI to present block storage volumes directly to guest VMs. The setup and configuration are essentially the same as for a physical host server, except that the VM is using virtual hardware.

## 4.4.6 Direct-attached in-guest iSCSI storage use cases

Dell EMC Unity arrays support in-guest iSCSI volumes mapped to guest VMs. Direct-attached storage for guest VMs is not recommended as a best practice unless there is a specific use case that requires it.

Here are some common use cases for in-guest iSCSI:

- **Performance:** Use in-guest iSCSI if a workload has a high I/O requirement and the performance gain (even if small) compared to using a VHD may be beneficial. For most workloads, there is no notable difference in performance between a direct-attached disk or virtual hard disk if the environment is designed and sized properly. In this case, the extra complexity with in-guest iSCSI can be avoided.
- **High availability:** The use of iSCSI enables VM clustering. However, the preferred VM clustering method is to use shared virtual hard disks. Shared VHDs were introduced with Windows Server 2012 R2 Hyper-V.
- **I/O isolation:** Using in-guest iSCSI is helpful when it is necessary to troubleshoot I/O performance on a volume and it must be isolated from other servers and workloads.
- **Data isolation:** This use case is helpful when it is necessary to use Unity to snap or replicate a specific subset of data. However, this result can also be accomplished by placing a VHD on a dedicated physical host volume.
- **Large capacity volumes:** Use in-guest iSCSI when a data volume that is presented to a guest VM may exceed the maximum size for a VHD (2 TB) or VHDX (64 TB).

There are also some limitations to consider before using direct-attached storage for guest VMs:

- **Checkpoints:** The ability to perform native Hyper-V checkpoints (snapshots) is lost. However, the ability to use Unity snapshots is unaffected.
- **Complexity:** Using in-guest iSCSI adds complexity and requires more overhead to manage.
- **Mobility:** VM mobility is reduced due to creating a physical hardware layer dependency.

---

**Note:** Legacy environments that are using direct-attached disks for guest VM clustering should consider switching to shared virtual hard disks, particularly if migrating to a newer version of Windows Server Hyper-V.

---

## 4.4.7 Guest VMs and pass-through disks

A block-based pass-through disk is storage that is mapped to a Hyper-V host or cluster and passed through directly to a Hyper-V guest VM. See Figure 15. The Hyper-V host or cluster has visibility to a pass-through disk but does not have I/O access. The Hyper-V host shows the disk in a reserved state since only the guest VM has I/O access.

Pass-through disks can be presented as boot or data volumes to a VM. Pass-through disks can be mapped to stand-alone Hyper-V hosts or cluster nodes.




Figure 15 Add a Unity volume to a guest VM as a pass-through disk

Although Unity supports pass-through disks, they are discouraged unless there is a specific use case that requires them. Usually, they are no longer necessary because of the feature enhancements offered with newer releases of Hyper-V. These features include generation 2 guest VMs, the VHDX format, and shared VHDs in newer Windows Server Hyper-V versions. Use cases for pass-through disks are like those for direct-attached storage as listed in section 4.4.6).

Limitations when using pass-through disks for direct-attached storage include the following:

- **Support for differencing disks is lost:** The use of a pass-through disk as a boot volume on a guest VM prevents the use of a differencing disk.
- **Difficult to manage at scale:** The use of pass-through disks becomes unmanageable and impractical at larger scale due to the large number LUN IDs consumed on physical Hyper-V hosts and nodes. Each pass-through disk consumes a LUN ID on every Hyper-V host or node that it is mapped to. In a large clustered environment with many nodes and VMs, pass-through disks can consume dozens or even hundreds of LUN IDs on each physical Hyper-V node.

#### 4.4.8 Dell EMC Unity and Hyper-V server clusters

When mapping shared volumes to multiple hosts, ensure that the volume is mapped to all nodes in the cluster using a consistent LUN number. This guidance applies to quorum disks, cluster disks, and cluster shared volumes (CSV).

Specify a specific LUN ID when mapping a volume to a host, node, or host group with Dell EMC Unisphere™. Alternately, you can allow Unisphere to pick the next available free LUN ID. LUN IDs can be changed in Unisphere after the volume is mapped if necessary to make the LUN ID consistent across all nodes in a cluster.

As a best practice and a time-saving tip, configure host groups in Unisphere. In this way, when mapping new storage LUNs, the LUN ID will be the same on all hosts or nodes in the group.

#### 4.4.9 Volume design considerations for Dell EMC Unity

One of the design considerations for which there are often no clear answers is how many guest VMs to place on a Dell EMC Unity volume. While many-to-one and one-to-one strategies both have advantages, a many-to-one strategy presents a good design starting point in most scenarios and can be adjusted for specific uses cases.

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

- Fewer Dell EMC Unity array volumes to create and administer (avoids volume sprawl).
- Quicker VM deployment because additional guest VMs do not require creation of a new volume on the Dell EMC Unity array.

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

- Easier to isolate and monitor disk I/O patterns for a specific Hyper-V guest VM.
- Ability to quickly restore a guest VM by recovering the Dell EMC Unity volume from a snapshot.
- Administrators have more granular control over what data gets replicated when Dell EMC Unity volumes are replicated to another location.
- Makes it faster to move a guest VM by remapping the volume rather than copying large virtual hard disk files from one volume to another over the network.

Other strategies might include placing all boot virtual hard disks on a common CSV, and data volumes on other CSVs.

#### 4.4.10 Disk format time

With trim/unmap enabled (the default setting in Windows Server), significant wait time may occur when formatting a large Unity volume (several TB or larger). The larger the volume, the longer the format wait time can be. Volumes formatted as NTFS or ReFS can experience long format wait times.




Figure 16 Avoid long format wait times by temporarily disabling trim/unmap

To avoid a long format wait time when mapping and formatting a large new volume, temporarily disable trim/unmap. This setting is disabled using the `fsutil` command from a command prompt with administrator privileges.

Figure 17 shows the commands to query the state and to disable trim/unmap for NTFS and ReFS volumes on a host. A **DisableDeleteNotify** value of **1** means that trim/unmap is disabled, and long format wait times are avoided when performing a quick format.

Changing the state of `DisableDeleteNotify` does not require a host reboot to take effect.




Figure 17 Use the `fsutil` command to query or change the state of trim/unmap

Once the volume is formatted, reenable trim/unmap so the host can take advantage of deleted space reclamation for NTFS volumes.

#### 4.4.11 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 **chkdsk** (checkdisk) against large volumes. ReFS is supported with Unity, but trim/unmap is not supported with ReFS volumes.


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

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

#### 4.4.12 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. It is recommended to disable automount on Hyper-V servers in a Unity environment.

Figure 18 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

## 4.5 Offloaded Data Transfer (ODX)

Offloaded Data Transfer (ODX) reduces CPU and network utilization on a host server by offloading a file-copy process from the host server to the Dell EMC Unity array. This feature is supported in Hyper-V environments with Windows Server 2012 and newer. ODX is enabled by default and it is a best practice to leave it enabled, unless there is a need to obtain performance benchmarks or troubleshoot a problem.

SCVMM environments use ODX when it is supported. Progress bars will indicate that a copy operation is **rapid copy** when ODX is used to deploy a new VM from an SCVMM library server.

## 4.6 Placement of page files

Windows Servers and VMs typically place the page file on the boot volume by default, and automatically manage page file and memory settings without user intervention. These settings should not be changed unless an application vendor provides guidance on how to tune the page file and memory settings to optimize the performance of a workload. Ultimately, each customer will need to decide on the best strategy as they consider variables that are unique to their environment.

With Dell EMC Unity, there can be some advantages to placing a page file on a separate volume from the perspective of the storage array. The following reasons may not justify changing the defaults, but when a vendor recommends 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 to reduce the changed data on the system (boot) volume. Moving the page file can reduce the size of Dell EMC Unity snapshots of boot volumes which will conserve space in the disk pool.
- Volumes or virtual hard disks dedicated to page files typically do not require snapshot protection, and therefore do not need to be replicated to a remote site. This guidance is helpful when there is limited bandwidth for replication of volumes and snapshots to other Dell EMC Unity arrays.

## 4.7 Active Directory domain controller placement

If the Hyper-V cluster service depends on Microsoft Active Directory® (AD) to authenticate, correct placement of domain controller VMs is critical to avoid unintended service interruptions.

Consider a situation where a Hyper-V cluster goes offline that also hosts both domain controller VMs for the environment. If the cluster service depends on AD, the cluster service is unable to start if a domain controller is unavailable to provide authentication. The administrator must remediate by manually recovering a domain controller VM to a stand-alone Hyper-V host or another cluster, so AD services become available again.

To protect against this situation in Unity environments, configure at least one domain controller on a physical server with local boot (along with other critical services). Regardless of the state of external storage or the storage fabric, critical services such as AD, DNS, and DHCP remain continuously available. This availability assumes that the management network stays functional.


Other strategies for domain controllers in Hyper-V environments include the following:

- Place virtualized domain controllers on stand-alone Hyper-V hosts or on individual cluster nodes if there is an AD dependency for cluster service authentication. In this way, AD VMs start and run independent of any cluster dependencies.

- Use Hyper-V Replica (2012 and newer) to ensure that domain controller VMs can be recovered on another host.
- Configure Hyper-V so it does not have an AD dependency to authenticate cluster services (Server 2016 and newer).

## 4.8 Dell EMC Unity data reduction and Hyper-V

Data reduction can be enabled or disabled on a Unity volume on a volume-by-volume basis to increase storage efficiency.



Data reduction works seamlessly in the background, so Windows Server Hyper-V hosts and workloads require no special configuration to take advantage of Unity data reduction. The amount of data reduction achieved is a function of the type of data on the Unity volume. Datasets that exist in a compressed format natively such as video files may still achieve some additional Unity data reduction when this feature is enabled.

When an operating system or application is configured to leverage compression or deduplication on a dataset, Unity may also achieve some additional space savings when Unity data reduction is enabled.

Because performing data reduction operations can impact performance, administrators will typically enable data reduction where it provides the most benefit with the least amount of performance impact.

Option 1: enable data reduction for a Unity volume in Unisphere, but not in the host operating system or application.

Option 2: enable data reduction in the operating system or application but not on the Unity volume in Unisphere.

Option 3: enable data reduction for the Unity volume and in the operating system or application. As a best practice, this configuration is typically avoided. The performance impact may not be worth the small amount of additional space saving realized from performing data reduction on both Unity and in the operating system or application.

Testing may be necessary in order to understand the best data reduction strategy for datasets in your environment.

## 4.9 Queue depth best practices for Windows Server Hyper-V

Queue depth is defined as the total number of disk transactions that are allowed to be in flight between an initiator port (server) and a target port (storage). The initiator is typically a Windows Server HBA FC port or iSCSI initiator. The target is an FC or iSCSI port on the SAN array (in this case, Dell EMC Unity). Any given target port can have multiple initiator ports sending it data. The initiator queue depth is used to throttle the number of transactions any given initiator can send to a target from a host to keep the target from becoming flooded. When flooding happens, the transactions are queued, which can cause higher latencies and degraded performance for the affected workloads.

### 4.9.1 When to change queue depth

One question that is commonly asked is concerning best practices for queue depth settings for Windows Server hosts and nodes with Hyper-V. 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 HBA, iSCSI NIC, or CNA.

Often, there is no need to change the default queue depth, unless there is a specific use where changing the queue depth is known to improve performance. If a storage array is connected to a few Hyper-V nodes hosting a large-block sequential-read application workload, increasing the queue depth setting may be beneficial. However, if the storage array has many hosts all competing for a few target ports, increasing the queue depth on a few hosts might overdrive the target ports. The performance of all connected hosts would be negatively impacted.

Increasing the queue depth can sometimes increase performance for specific workloads. If set too high, there is an increased risk of overdriving the target ports on the storage array. Increasing the number of initiators and targets to spread out I/O can be an effective remediation if transactions are being queued and performance is impacted.

### 4.9.2 Vendor-specific HBA and CNA queue depth settings

Understanding the firmware and miniport driver registry settings for your host server FC HBA, iSCSI NIC, or CNA adapter and how these settings affect queue depth is important.

See the documentation for your particular FC HBA, iSCSI NIC, or CNA for direction on adjusting firmware or registry settings to modify queue depth.

---

**Note:** Changes to FC HBA, iSCSI NIC, or CNA firmware or registry settings that affect queue depth should be evaluated in a test environment before implementation on production workloads.

---

## 5 Dell EMC Unity snapshots and thin clones with Hyper-V

Dell EMC Unity snapshots can be used to protect any workload, including Hyper-V, on both block and file-based storage resources. Because snapshots use the same technology for block and file, managing snapshots is similar between storage resource types. Snapshots and thin clones can be created and managed with the Unisphere user interface, from the command line (CLI), or by using the RESTful API.

Dell EMC Unity snapshots leverage redirect on write (ROW) which means that new writes to snapped storage resources are written to a new location in the same disk pool. This method is space-efficient. Metadata pointers direct I/O to new data that has changed since the last storage resource snapshot, while the frozen data in snapshots is available if needed for recovery.

A thin clone is created from a snapshot. It is similar to a snapshot, except that a thin clone can be managed like a fully independent LUN. It will look and act like an ordinary LUN to the host server. A thin clone can have its own snapshot schedule and is space efficient because it uses metadata pointers to reference data from the source snapshot it was created from.

Thin clones support Dell EMC data services and features using Unisphere, the CLI, or RESTful API. In addition, thin clones:

- Can be mapped to a host as a regular LUN
- Support the configuration of host I/O limits
- Support snapshots and snapshot schedules
- Can be added to LUN consistency groups
- Support replication to other Dell EMC Unity arrays, along with their snapshot histories

Dell EMC Unity point-in-time snapshots and thin clones allow administrators to do the following in Hyper-V environments:

- Recover servers to a crash-consistent state, including Hyper-V hosts (when using boot-from-SAN) and guest VM workloads
- Provision lab, parallel, or isolated test environments (thin clones)
- Provision new servers from thin clones created from a gold image source snapshot

Point-in-time Dell EMC Unity snapshots can be taken of volumes regardless of content: boot-from-SAN disks, data disks, cluster shared volumes, pass-through disks, and direct-attached in-guest iSCSI disks. Disks, volumes, and their associated snapshots can be quickly and easily replicated to other Dell EMC Unity arrays for DR or other purposes.

### 5.1 Crash-consistent and application-consistent snapshots

A Dell EMC Unity snapshot is a space-efficient point-in-time copy of the data on a volume. A Unity snapshot is crash consistent, unless the administrator takes step to ensure in-flight data is flushed to disk (quiesced), and I/O is paused beforehand.

When recovering a server using a crash-consistent snapshot, it is similar to powering a server back on after an unexpected power outage. Servers and applications are often able to recover to a crash-consistent state without complications.

## 5.1.1 Achieving application consistency

Powering off a server will put it into an application-consistent state before taking a snapshot. Temporarily stopping application services can be used to ensure the application consistency of a workload before taking a snapshot. However, powering down a server or stopping services during the workday may be service-impacting. Tools such as PowerShell and Ansible® can help orchestrate and automate these processes so they run at a time when end users will not be impacted. Other methods for achieving application consistency include using a third-party backup application that leverages the Microsoft volume shadow copy service (VSS). VSS can briefly pause and quiesce server or application I/O to ensure consistency.

If application consistency is necessary, the administrator must ensure that the server or workload is in a consistent state before a Unity snapshot is taken. Hyper-V environments hosting transactional workloads such as Microsoft Exchange or SQL Server® have a higher risk of data loss or corruption when attempting to recover to a crash-consistent state. Ensure the consistency of a transactional workload before a snapshot is taken to avoid data loss or corruption during a recovery.

Some examples for how to configure and use Dell EMC Unity snapshots for a Hyper-V environment are provided in the following sections.

## 5.2 Using Dell EMC Unity snapshots to recover guest VMs

Hyper-V guest VMs can be recovered to a previous point in time by using Dell EMC Unity snapshots. Thin clones can be used to create copies of VMs in an isolated environment at the same or a different location when volume replication between Dell EMC Unity arrays used. This section provides guidance and best practices for several different recovery options using snapshots.

### 5.2.1 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 Unity volume that is mapped to a Hyper-V host.

A VM hard disk and configuration files may reside on separate host data volumes. It is a best practice to configure a consistency group in Unity for these volumes, so they are snapped simultaneously. For example, a boot virtual hard disk for a VM might reside on one host volume, while one or more virtual hard disks for data might reside on another volume. Database files may also span several Unity volumes.

When performing a recovery of a VM with Dell EMC Unity snapshots or thin clones, there are several options.

**Option 1:** Recover the existing data volume on the host that contains the VM configuration and virtual hard disks by using a Dell EMC Unity snapshot. A thin clone can also be used to replace the existing volume by using the same LUN number, drive letter mapping, or mount point.

- If the data volume contains only one VM, this approach 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 a single VM.
- The VM being recovered will 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 or thin clone containing the VM configuration and virtual hard disks to the host as a new volume using a new drive letter or mount point. Recover the VM by manually copying the virtual hard disks from the recovery snapshot or thin clone to the original location.

- The original virtual hard disks are deleted, moved, or renamed.
- 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. This step is necessary to allow the guest VM to start without permissions errors.
- This approach may not be practical if the virtual hard disks are so large that the time required to copy the virtual hard disks over the network is excessive. 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 virtual hard disk as a volume on the host server temporarily.

**Option 3:** Map the recovery snapshot or thin clone to a different Hyper-V host. Recover the VM there by importing the VM configuration. Optionally, create a VM that points to the virtual hard disks on the recovery volume.

- This option is common in situations where the original VM and the recovery VM both need to be online simultaneously. Isolation is necessary to avoid name or IP conflicts, or split-brain with data writes.
- This option is also a great way to recover when the original host server is no longer available due to a host failure.

Before beginning any VM recovery, record essential details about the VM hardware configuration, such as number of virtual CPUs, RAM, virtual networks, and IP addresses. In case importing a VM configuration is not supported or fails, having this information available will aid recovery.

## 5.2.2 Recover guest VM on a cluster shared volume

Use Unity snapshots or thin clones to recover guest VMs that reside on a cluster shared volume (CSV). This strategy is similar to recovering a guest VM on a stand-alone host as detailed in section 5.2.1. However, recovering a VM from a snapshot or thin clone of a CSV may require changing the disk signature to avoid a disk signature conflict.

Windows servers assign each volume a unique disk ID (or signature). For example, the disk ID for an MBR disk is an eight-character hexadecimal number such as 045C3E2F4. No two volumes mapped to a server can have the same disk ID.

When a Dell EMC Unity snapshot or thin clone is taken of a Windows or Hyper-V volume, the snapshot is an exact point-in-time copy, which includes the Windows disk ID. A recovery volume created from a snapshot or thin clone will have an identical disk ID.

With stand-alone Windows or Hyper-V servers, disk ID conflicts are avoided because stand-alone servers automatically detect duplicate disk IDs and change them dynamically with no user intervention.

However, host servers are not able to dynamically change conflicting disk IDs when disks are configured as CSVs, because the disks are mapped to multiple nodes simultaneously.

When attempting to map a copy (snapshot or thin clone) of a CSV back to any server in that same cluster, the recovery volume will cause a disk ID conflict. This situation can be service-affecting.

There are a couple of ways of working around the duplicate disk ID issue:


**Option 1:** Map the recovery volume containing the CSV to another host that is outside of the cluster and 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 has been changed, remap the recovery volume to the cluster. The steps to use Diskpart.exe to change the disk ID are detailed in section 5.2.3.

### 5.2.3 Change a cluster shared volume disk ID with Diskpart

Follow these steps to change a disk ID before mapping to the same cluster. PowerShell can also be used.

1. Log in to a stand-alone Windows Server (with or without the Hyper-V role installed) that is available in Unity. This server must not be a member of the Hyper-V cluster.
2. Open a command window with administrator rights.
3. Type **diskpart** and press **Enter**.
4. Type **list disk** and press **Enter**.
5. Make note of 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 Unisphere to map a thin clone of the cluster disk to this host.
7. From the diskpart command prompt, type **rescan**, and press **Enter**.
8. Type **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          *

DISKPART>

```

9. To select the offline disk, type **select disk <number>** and press **Enter**.
10. Type **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          *

DISKPART>

```

12. Type **uniqueid disk** and press **Enter** to view the current ID for the disk.
13. To change the disk ID, type **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, type **attributes disk clear readonly** and press **Enter** to clear the read-only attribute.

---


14. Type **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=12345678-1234-1234-1234-123456789abd
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=12345678-1234-1234-1234-123456789abd
DISKPART> uniqueid disk
Disk ID: {12345678-1234-1234-1234-123456789ABD}
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 Unisphere 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:** Disabling automount is recommended as a best practice to prevent hosts from automatically assigning drive letters to volumes in Hyper-V recovery scenarios. See section 4.4.12 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.



- b. Close PowerShell.

After the volume is online, perform the required steps, such as data or VM recovery.

### 5.3 Use Dell EMC Unity thin clones to create a test environment

In addition to VM recovery, Dell EMC Unity thin clones can be used to quickly create test or development environments that mirror a production environment. When thin clones containing VMs are replicated to another location, this strategy makes it 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 from the original VMs.

---

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.

### 5.4 Leverage Dell EMC Unity to create gold images

With Dell EMC arrays, an administrator can create gold images to accelerate and simplify the process of deploying new servers. Gold images can be used to deploy:

- New host servers when using boot-from-SAN
- Guest VMs that use pass-through disks for boot disks
- Guest VMs that boot from a virtual hard disk. In this case, using a virtual hard disk with an operating system that has been system-prepared (using Sysprep.exe) as the gold image source, would be preferred.

Using gold images provides the following benefits:

- Server provisioning is quicker, and reconfiguration steps are minimized.
- SAN utilization is more efficient when a host or VM is provisioned from a gold image. Only new data consumes space on the storage array. Data that has not changed is read from the gold image source volume.


The steps to configure a Windows Server or Hyper-V gold image are as follows:

1. Create and map a Dell EMC Unity volume to a host server as a boot volume (LUN 0).
2. Build your base operating system image, install roles and features, and fully configure and update it. This step will minimize the changes that have to be made to each new server that is deployed using the gold image.
3. Once the operating system is fully staged, run **Sysprep.exe** and choose the **Generalize, Out-of-box Experience**, and **Shutdown** options.
  - After running Sysprep.exe, if the server is a guest VM, use Hyper-V Manager to delete the guest VM. This step will delete the guest VM configuration files but will preserve the boot virtual hard disk. The virtual hard disk is the only file needed for the gold image.

---

**Note:** Do not use SCVMM to delete the guest VM because it will also delete the virtual hard disk file. Use Hyper-V Manager instead.

---




4. Allow the server to power down. This step will ensure that it is in a consistent state.
5. Manually create a Dell EMC Unity snapshot of the volume and set it to never expire. Assign it a descriptive name that clearly identifies it as a gold source.
6. From the snapshot, create one or more thin clones and map them to new host servers.
6. Complete any required configuration steps on the host such as configuring boot paths.
7. Boot the new host servers and allow the initial boot process to complete.
8. Customize the server configuration as needed. Leverage PowerShell to automate configuration steps that are repetitive.

---

**Note:** A snapshot that is designated as a gold image source supports a maximum of 16 thin clones. As a result, a maximum of 16 new hosts can be provisioned from a single gold source. To deploy more servers, create additional snapshots that can act as a gold source for more thin clones.

---



### 5.4.1 Gold images and preserving balanced Dell EMC Unity storage processors

Each Unity array contains two identical storage processors for performance and resiliency.




Figure 19 Rear view of a Unity array showing storage processor A and B

When a new volume is created on a Dell EMC Unity array, it is assigned to storage processor A (SP A) or storage processor B (SP B). All snapshots and thin clones created from a volume are also assigned to the same SP. If using gold images to deploy a large number new hosts, care should be taken to ensure that volume ownership does not become imbalanced between SP A and SP B.




Figure 20 View or change the storage processor owner for a LUN

View the properties of a volume in Unisphere to view or change the current SP owner.

## 5.5 Using Dell EMC Unity for Hyper-V VM migration

Microsoft provides native tools to move or migrate VMs with Windows Server 2012 Hyper-V and newer. As a result, so there are fewer use cases for using SAN-based snapshots to move VMs. When a guest VM is live migrated from one node to another node within the same Hyper-V cluster configuration, no data needs to be copied or moved. All nodes in that cluster have shared access to the underlying cluster shared volumes (CSV) or SMB file shares.

When a guest VM is migrated from a host or cluster, the virtual hard disks must be copied to the target host or cluster. This migration will consume network bandwidth and may require significant time if the virtual hard disks are large. Storage space is consumed unnecessarily because another copy of the data is created.

It may be quicker to leverage Dell EMC Unity when moving VMs to another host or cluster. Unmap the host volume containing the VM configuration and virtual hard disks and map the volume to the new target host or cluster. This process can also be completed using a thin clone.

This process may require down time to move the VM during a maintenance window. However, it may be a more practical approach than waiting for large virtual hard disks (when they are multiple terabytes) to copy over the network. Consuming additional SAN space unnecessarily is also avoided.

## 6 Boot-from-SAN for Hyper-V

Windows Server and Windows Server with the Hyper-V role installed support boot-from-SAN. This option can be used when a host server is not configured with a local boot disk. Usually, physical hosts will support boot-from-SAN if these requirements are met:

- The host must have an FC or iSCSI adapter with firmware that supports a boot-from-SAN configuration.
- The FC fabric or iSCSI network must be configured to allow the host adapter initiator ports to see target ports on Unity.
- The host operating system must support boot-from-SAN. Windows Server, and Windows Server with the Hyper-V role installed, support boot-from-SAN.
- If multiple boot paths are available to the host (recommended), the host operating system must support MPIO. MPIO support is enabled by installing PowerPath or a supported device-specific module (DSM) in the host operating system. Unity supports the native DSM provided with Windows Server to enable MPIO.

Deciding whether to use onboard or SAN-based LUNs for boot depends several design considerations that are unique to each environment. There are advantages and disadvantages to both options, and sometimes, booting from local disk in the host server may be preferred.

### 6.1 Boot-from-SAN advantages

Advantages of using boot-from-SAN include:

- Dell EMC Unity snapshots of boot volumes provide for quick host recovery.
- Boot volumes can be replicated to another location for enhanced disaster recovery (DR) protection when both sites use similar host server hardware.
- Dell EMC Unity thin clones can be leveraged to create boot volumes from gold-image snapshots for quick server provisioning.

### 6.2 When to use boot from local disk

It is recommended to boot from the local disk in the following scenarios:

- When critical workloads or infrastructure roles such as AD, DNS, or DHCP need to remain online during offline SAN maintenance or unplanned storage outages
- When boot-from-SAN is not an option, such as when using SMB file shares to access shared storage

### 6.3 Configure a Hyper-V host server to boot from SAN

For guidance on configuring host servers to access FC or iSCSI LUNs, see the [Configuring Hosts to Access Fibre Channel \(FC\) or iSCSI Storage](#) white paper.

In addition to the information in this guide, there are a few configuration steps and best practices to consider when configuring boot from SAN:

- If equipped, disable the onboard disk or RAID controller in the host server.
- Set the boot order on the host server to use the FC or iSCSI adapter as the first boot device.
- In Unisphere, verify that the boot LUN is mapped to the host as LUN ID = 0
- Configure a single boot path from Unity to the host when first staging the operating system. After the operating system is installed, install the MPIO feature, configure MPIO, and add additional boot paths. This step is not necessary when using a gold image that already has MPIO configured.

---

**Note:** Modify the MPIO registry settings according to the guide listed previously in this section to ensure optimal timeout settings.

---

## 7 Dell EMC Unity and NAS for Hyper-V

In addition to block storage, Dell EMC Unity arrays support file-based storage through network attached storage (NAS). Dell EMC Unity arrays support SMB 3.1.1 which was released with Windows 10 and Windows Server 2016. SMB 3.1.1 provides enhanced security and performance over previous SMB versions.

Benefits offered with NAS on Dell EMC Unity include failover and failback of NAS servers and associated file shares between arrays, including Failover with Sync.

SMB shares can be presented to stand-alone Hyper-V hosts by using Universal Naming Convention (UNC) paths or Hyper-V clusters by using Microsoft SCVMM.

For more information about NAS support with Dell EMC Unity, see the *Dell EMC Unity: NAS Capabilities* white paper at [Dell EMC Unity family technical white papers and videos](#).

### 7.1 Present Dell EMC Unity SMB file shares to Hyper-V clusters

SMB file shares can be presented as shared storage to Hyper-V cluster nodes by leveraging Dell EMC Unity SMI-S integration with Microsoft SCVMM. SCVMM does not present SMB file shares to a Hyper-V cluster as CSVs like it does with block storage. However, the functionality of an SMB file share presented to a Hyper-V cluster is similar to a CSV.


For administrators accustomed to leveraging block storage presented as CSVs, there are some differences that are important to understand. For example, SMB file shares presented to a cluster are not visible under the storage in Failover Cluster Manager. Block storage must be converted to a CSV before it can be used for VMs. SMB shares do not have to be converted to a CSV. SMB file shares presented to a Hyper-V cluster are managed transparently in the background by SCVMM.

There are a few functional differences between a CSV and an SMB file share. Using SMB file shares for shared storage on Hyper-V clusters requires an investment in SCVMM, whereas block storage does not. However, SCVMM will manage either kind of storage: block or SMB. SCVMM provides full integration with Dell EMC Unity arrays. For more information about SMI-S integration, see the [Dell EMC Unity Family SMI-S Programmer's Guide](#).


The steps in section 7.1.1 assume that an SCVMM server is already configured in the environment according to Microsoft best practices.

### 7.1.1 Configure SCVMM with SMI-S to use SMB file shares


1. Log in to the SCVMM administrator console.
2. In the **Fabric** workspace, click **Storage**. On the **Home** tab, under the ribbon bar click **Add Resources** and from the drop-down list select **Storage Devices** to launch the **Add Storage Device Wizard**.




3. Under **Select Provider Type**, select SAN and NAS devices discovered and managed by an SMI-S provider and click **Next**.




- Under **Specify Discovery Scope**, provide the management IP address of the Dell EMC Unity system with the NAS server, file system, and SMB share. Also provide an SCVMM **Run As** account that has administrator rights on the domain. Click **Next**.




- Allow the wizard to discover and import the storage device information. This step may require several minutes. Import the certificate if prompted.




- Complete the steps in the **Add Storage Device Wizard**. Verify that the SMB file shares are listed under File Servers in SCVMM. There are two NAS servers, each with one SMB share, in this example.



- In SCVMM in the **VMs and Services** workspace, expand **All Hosts**, right-click the server cluster, and select **Properties**.
- Under **File Share Storage**, click **Add**, and from the drop-down list select the file share (listed by UNC path).



- Once the SMB file share has been assigned, verify the **Access Status** column shows as green. In this example, both available SMB file shares are assigned to the Hyper-V cluster.



VMs can now be deployed to either Dell EMC Unity SMB file share in this example because they are now cluster-aware resources managed by SCVMM through SMI-S.

## A Technical support and additional resources

[Dell.com/support](https://www.dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

### A.1 Related resources

View the following resources for additional information:

- [\*Dell EMC Unity Family SMI-S Programmer's Guide\*](#)
- [\*Microsoft Tech Community\*](#)
- [\*Microsoft technical documentation library\*](#)
- [\*Microsoft Windows Admin Center\*](#)
- [\*Dell EMC Unity Simple Support Matrix\*](#)