

Dell SC Series Storage and Microsoft Hyper-V

Dell Storage Engineering
October 2016

Revisions

Date	Description
Jun 2009	Initial release
Aug 2009	Updated for Server 2008 R2
Sep 2011	Updated for Server 2008 R2 SP1 and SCOS 5.5
Oct 2012	Updated for Server 2012, Enterprise Manager 6.2 and SCOS 6.2
Dec 2012	Support for virtual Fibre Channel added with SCOS 6.3 and Enterprise Manager 6.3
Oct 2013	Updated for Server 2012 R2; convert document to new template
Oct 2016	Updated for Server 2016

Acknowledgements

This document was authored by Marty Glaser, a member of the Dell EMC Storage Applications Engineering team.

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 © 2009-2016 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [11/29/2016] [White Paper] [CML1009]

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

Table of contents

Revisions.....	2
Acknowledgements.....	2
Executive summary.....	6
1 Introduction.....	7
1.1 Audience.....	7
2 Overview.....	8
2.1 Dell SC Series storage arrays	8
2.2 Microsoft Hyper-V	8
2.3 General SC Series storage and Hyper-V best practices	9
3 Optimize SC Series storage with Hyper-V	10
3.1 Hyper-V integration services	10
3.2 Hyper-V guest VM generations	11
3.3 Virtual hard disks	12
3.3.1 Virtual hard disk format.....	12
3.3.2 Virtual hard disk type	13
3.3.3 Virtual hard disks and thin provisioning with SC Series storage	15
3.3.4 Overprovisioning with dynamic virtual hard disks.....	15
3.4 Present SC Series storage to Hyper-V.....	16
3.4.1 Transport options.....	16
3.4.2 SC Series and front-end SAS support for Hyper-V	17
3.4.3 Mixed transports	17
3.4.4 MPIO best practices	17
3.4.5 Guest VMs and in-guest iSCSI and virtual Fibre Channel disks	18
3.4.6 Direct attached storage use cases	18
3.4.7 Guest VMs and pass-through disks.....	19
3.4.8 SC Series arrays and cluster server objects	20
3.4.9 LUN limits for larger Hyper-V clusters	20
3.4.10 Volume design considerations for SC Series.....	21
3.5 Offloaded Data Transfer (ODX).....	22
3.6 Disable automount.....	22
3.7 Placement of page files	23
3.8 SC Series data reduction and Hyper-V	23

4	SC Series snapshots and Hyper-V	24
4.1	Replay Manager support for Hyper-V	24
4.2	Use SC Series snapshots to recover guest VMs	24
4.2.1	Recover a guest VM on a standalone Hyper-V host	25
4.2.2	Recover guest VM on a cluster shared volume.....	26
4.2.3	Change a cluster shared volume's disk ID with Diskpart	26
4.3	Use SC Series snapshots to create a test environment.....	28
4.4	Leverage SC Series to create gold images	29
4.4.1	Gold images and preserving balanced SC Series controllers	30
4.5	SC Series snapshots and Hyper-V VM migration	33
5	Data progression and Hyper-V.....	34
5.1	Tuning data progression settings for Hyper-V	34
5.1.1	Data Progression with archival data	35
5.1.2	Data copies and migrations	36
6	Disk space recovery with Hyper-V	37
6.1	SC Series support for Trim/Unmap with Hyper-V.....	37
6.2	Space recovery with 2008 R2 Hyper-V	37
7	Boot from SAN for Hyper-V	39
7.1	Boot from SAN advantages	39
7.2	When to use boot from local disk	39
7.2.1	Domain controllers on Hyper-V clusters	39
7.3	Configure Hyper-V hosts to boot from SAN	40
8	PowerShell integration	41
8.1	Importance of PowerShell	41
8.2	PowerShell automation with Hyper-V and SC Series.....	41
9	Business continuity	43
9.1	Cost versus risk analysis	43
9.2	Disaster recovery and disaster avoidance	44
9.3	Live Volume with Auto Failover for Microsoft	44
9.3.1	Predefined DR plans	45
9.4	Replay Manager for Hyper-V	45
9.5	Azure Site Recovery (ASR) with SC Series	45
10	Conclusion.....	46

A Technical support and additional resources.....47

A.1 Supplemental documentation.....47

Executive summary

Microsoft® virtualization technology and SAN technology continue to evolve and improve enabling robust, feature-rich environments. The two technologies provide high performance, resiliency, mobility, and reliability for a variety of workloads. However, determining best practices is not always straight forward, especially when there are many different ways to go about architecting a solution. When this is the case, determining the right configuration often requires considering factors that are unique to the environment. The intent of this document is to help readers understand the best practice configuration options available to them when configuring Microsoft Hyper-V® with Dell™ Storage SC Series SAN arrays.

1 Introduction

Dell SC Series arrays have supported Hyper-V since the release of Windows Server® 2008, and many new feature enhancements have been added to both Hyper-V and SC Series storage over time. SC Series arrays are designed from the ground up with redundancies to avoid downtime for such events as component failures, maintenance, upgrades and expansion. These redundancies benefit Hyper-V and other workloads as well. Hyper-V also includes similar features. When SC Series arrays host Hyper-V workloads, their feature sets complement each other.

There are several other means of presenting storage to Hyper-V, such as traditional on-board storage in host servers. Another is Storage Spaces, which Microsoft has promoted for Hyper-V platforms since the release of Server 2012 R2. Storage Spaces Direct and Storage Replica were introduced with Server 2016. While Storage Spaces, Storage Spaces Direct, and Storage Replica offer similar basic features as traditional SAN arrays, the SC Series SAN provides a much more powerful and complete set of integrations, management and monitoring tools that are not available with Storage Spaces or Storage Replica alone.

Determining best practices is often subject to many subtle nuances and factors that vary based on the customer environment and personal preference. These include budgetary constraints, complexity, required depth of knowledge, service level agreements, and regulations to name a few. What works well for one environment may not work well for another. On occasion, the configuration a customer chooses might not be completely harmonious with typical best practices. For example, it might be perfectly acceptable to omit some more costly design elements and redundancies in order to save money in a test, development, or DR environment that can suffer down time without significantly impacting the business. When different configuration options are possible, determining the best design is ultimately up to the customer to decide.

1.1 Audience

This document is intended for technology professionals interested in learning more about best practices for configuring Microsoft Hyper-V with SC Series arrays. It is not intended to be a how-to guide for deploying Microsoft Hyper-V or SC Series storage. It assumes the reader already has a good working knowledge of both. See [Appendix A](#) for additional documentation and how-to guidance for configuring Microsoft Hyper-V with specific SC Series configurations.

We welcome your feedback along with any recommendations for improving this document. Send comments to StorageSolutionsFeedback@dell.com.

2 Overview

This section provides a brief overview for SC Series Storage, Microsoft Hyper-V, and general best practices.

2.1 Dell SC Series storage arrays

SC Series storage solutions provide many robust features such as true flash optimization, thin provisioning, data optimization, data reduction (deduplication and compression), automated sub-LUN tiering, sub-disk RAID levels, synchronous replication with automatic failover, and intelligent read and write data placement. SC Series arrays provide a unified platform for the ultimate experience in performance, adaptability, and efficiency.

In addition to raw capacity and I/O performance, other important factors such as monitoring, reporting, trending, data protection (backups, snapshots and replication) and the ability to recover in case of a disaster are equally important. The SC Series is well suited to provide a solid, proven storage solution for Hyper-V environments to meet all of these business needs. To learn more about specific SC Series arrays, visit the [Dell SC Series Storage Solutions](#) website.

2.2 Microsoft Hyper-V

The Windows Server platform leverages Hyper-V for virtualization technology. Initially offered with Server 2008, Hyper-V has matured with each release to include many new features and enhancements. With the release of Server 2016, Hyper-V is also supported on Nano Server in addition to Server Core and Server with Desktop.

Microsoft Hyper-V has evolved to become a mature, robust, proven virtualization platform. At a basic level, it is a layer of software that presents the physical host server hardware resources in an optimized and virtualized manner to one or more guest virtual machines (VMs). Hyper-V hosts (also referred to as nodes when clustered) can greatly enhance host server hardware utilization (including processors, memory, NICs and power) by allowing many VMs to share these resources at the same time. Hyper-V Manager and related management tools (such as Failover Cluster Manager, Virtual Machine Manager, and PowerShell) offer administrators great control and flexibility for managing host and VM resources, VM mobility, SAN storage, balancing workloads, provisioning systems, and ensuring quick recovery.

Some core Hyper-V features (such as dynamic memory) are storage agnostic, and so may not be covered in any great detail this document, except as they tie to specific integrations with SC Series storage. To learn more about Hyper-V, including tools, videos, blogs, and the feature enhancements that have been made available with each new release of Hyper-V (including 2016), visit the [Microsoft TechNet Library](#) website.

In addition to this document, SC Series documentation and video library at [Dell TechCenter](#) contains deployment guides, demo videos, and reference architectures in support of single and heterogeneous application workloads running on Hyper-V and Dell SC Series arrays, including Microsoft Exchange, Microsoft SQL Server and VDI.

2.3 General SC Series storage and Hyper-V best practices

Both Hyper-V and SC Series storage are mature, feature-rich solutions that together present administrators with a diverse range of configuration options to solve key business objectives.

On occasion, legacy systems may be dependent on a strategy that would not be considered best practices today, such as using pass-through disks with generation 1 guest VMs for guest clustering. Using shared virtual hard disks would be the preferred method today. Often, the best course of action is to run legacy configurations until they reach their life expectancy, as it might not be possible or too disruptive to change or upgrade them.

Much of the best practice guidance in this document is focused on what would be the best design if building a new solution today. Administrators might be very comfortable with running, or continuing to run legacy configurations because they accomplish the work needs. However, strong consideration should be given to upgrade to the latest trends and current best practices at important crossroads such as when upgrading or replacing infrastructure.

In many cases, there are no right or wrong answers. It will often come down to choosing between two good options. When considering best practices for your SC Series and Hyper-V environment, consider a design that:

- Keeps the environment as simple as possible to minimize administrative overhead
- Ensures redundancies for critical workloads to protect against single points of failure
- Ensures DR goals can be achieved
- Takes advantage of the latest features available in the current versions of SCOS and Hyper-V
- Can be evaluated in a test environment before deploying in production where possible
- Avoids or minimizes the use of any legacy features in either SCOS or Hyper-V
- Can scale to accommodate anticipated growth and other business needs over the expected life of the hardware

There are many Hyper-V best practices recommendations not specific to storage that will be touched on only briefly in this document to avoid duplication of information. Readers are encouraged to consult [Microsoft TechNet](#) and other online resources for more information on general Hyper-V best practices that are not related to SC Series storage. From the Hyper-V perspective, common Hyper-V tuning steps include:

- Minimizing or disabling unnecessary hardware devices and services to free up CPU cycles that can be used by other VMs; this also helps to reduce power consumption
- Scheduling tasks such as periodic maintenance, backups, virus scans and patching to run after hours, and stagger them to avoid spikes in I/O
- Tuning application workloads to reduce or eliminate unnecessary processes or activity, or scheduling during off hours when possible

3 Optimize SC Series storage with Hyper-V

Any discussion about SC Series arrays and Hyper-V best practices is of course going to center on the actual storage itself – the actual LUNs that are presented to hosts, nodes, clusters and VMs. Many core SC Series features such as thin provisioning, data progression, data reduction, snapshots, and replication work seamlessly in the background regardless of the platform. In many cases, the default settings for these features will work just fine with Hyper-V, or at least serve as a good configuration starting point. This document will point out additional configuration or tuning needed to enhance performance, utilization or up-time.

3.1 Hyper-V integration services

The Hyper-V integration services installed on a guest VM include drivers for the VM synthetic I/O devices. These services decrease the CPU demand on the VM and improve performance, including disk I/O.

The latest version of integration services should be installed on every VM that has integration services available. Supported VM OSs (Windows and some versions of Linux for example) include Hyper-V integration services out of the box, but they may still need to be updated after the VM is deployed. This is normally the first step in optimizing VM performance. With Server 2016 Hyper-V, integration services are now deployed to Windows VMs as a part of Windows updates, so there is less administration required to ensure Windows VMs stay current.

With older versions of Hyper-V and guest VMs, the need to install and update the integration services may not always be obvious. This might become an issue, for example, when a VM is migrated from an older Hyper-V host or cluster to a newer one. If a VM is not performing as expected, whether it be CPU, disk I/O, or network performance, verify that the VM integration services are current.

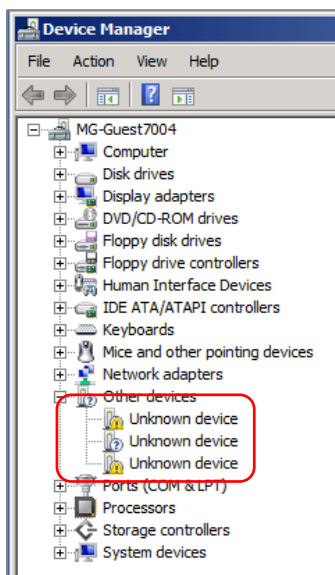


Figure 1 Unknown guest VM devices indicates missing or outdated integration services

Aside from performance problems, one of the key indications that integration services are outdated or not present on a Windows VM is the presence of unknown devices in Device Manager on the VM.

For versions of Hyper-V prior to 2016, use Hyper-V Manager to connect to a VM, mount the **Integration Services Setup Disk** as listed under the **Action** menu, and follow the prompts in the guest VM to complete the installation.

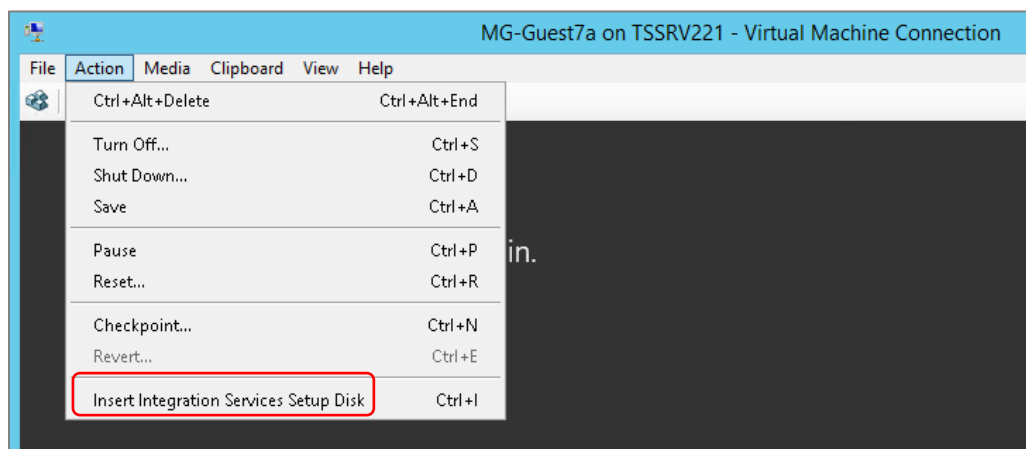


Figure 2 Mount Integration Services Setup Disk

3.2 Hyper-V guest VM generations

The type of guest VM (generation 1 or generation 2) has an impact on overall SC Series storage strategy. When Server 2012 R2 Hyper-V was released, Microsoft provided support for VMs created on previous versions of Hyper-V by designating them as generation 1 VMs. When creating new VMs on Hyper-V 2012 R2 or newer, they can be created as generation 1 or generation 2. The best practices recommendation from both Microsoft and Dell EMC is to configure new guests as generation 2 when possible.

While it is possible to convert a generation 1 VM directly to a generation 2 VM, it may require using third-party tools. Native support from Microsoft to do this is not available as of the publication date of this document. More information on this topic can be found on [Microsoft TechNet](https://technet.microsoft.com/en-us/library/8c56276e-3172-47a2-8775-2d97666d4123).

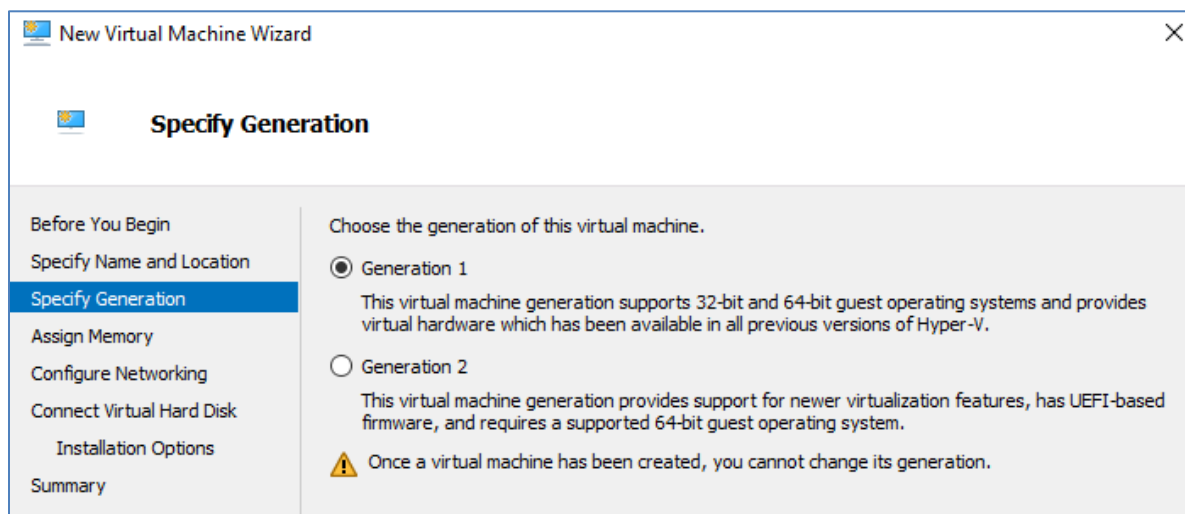


Figure 3 Specify a guest as generation 1 or generation 2

Generation 2 guests use United Extensible Firmware Interface (UEFI) when booting instead of a legacy BIOS. UEFI provides better security, and better interoperability between the OS and the hardware, which offers improved virtual driver support and performance. In addition, one of the most significant changes with generation 2 guests is the elimination of support for virtual IDE. Generation 1 VMs require the boot disk to be a virtual IDE disk. Generation 2 guests instead use virtual SCSI for all disks.

For both generations of guest VMs, if there are multiple disks requiring high I/O, each disk can be associated with its own virtual SCSI controller to further maximize performance.

3.3 Virtual hard disks

A virtual hard disk is a set of data blocks that is stored as a regular Windows file with a .vhd, .vhdx or .vhds extension, using the host operating system. It is important to understand the different format and type options for virtual hard disks and how this integrates with SC Series storage.

3.3.1 Virtual hard disk format

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

- VHD is supported with all Hyper-V versions and is limited to a maximum size of 2 TB. This is now considered a legacy format (use VHDX instead).
- VHDX is supported with Server 2012 Hyper-V and newer. The VHDX format offers better resiliency in the event of 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 Server 2016 Hyper-V and newer. VHDS is for virtual hard disks that are shared by two or more guest VMs in support of clustering (high-availability) configurations.

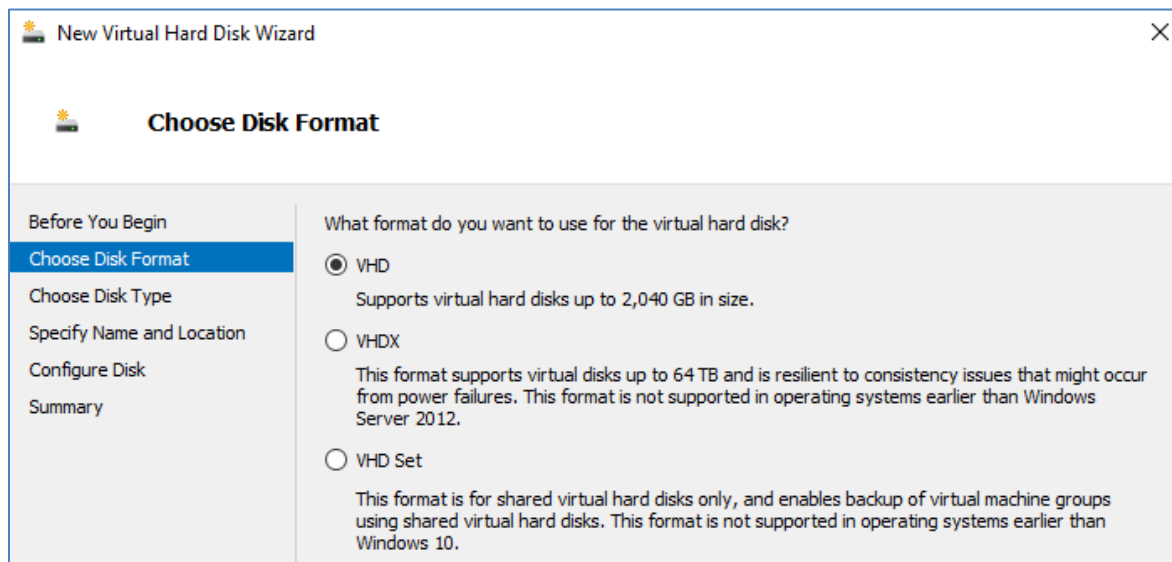


Figure 4 Different formats available for virtual hard disks

3.3.2 Virtual hard disk type

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

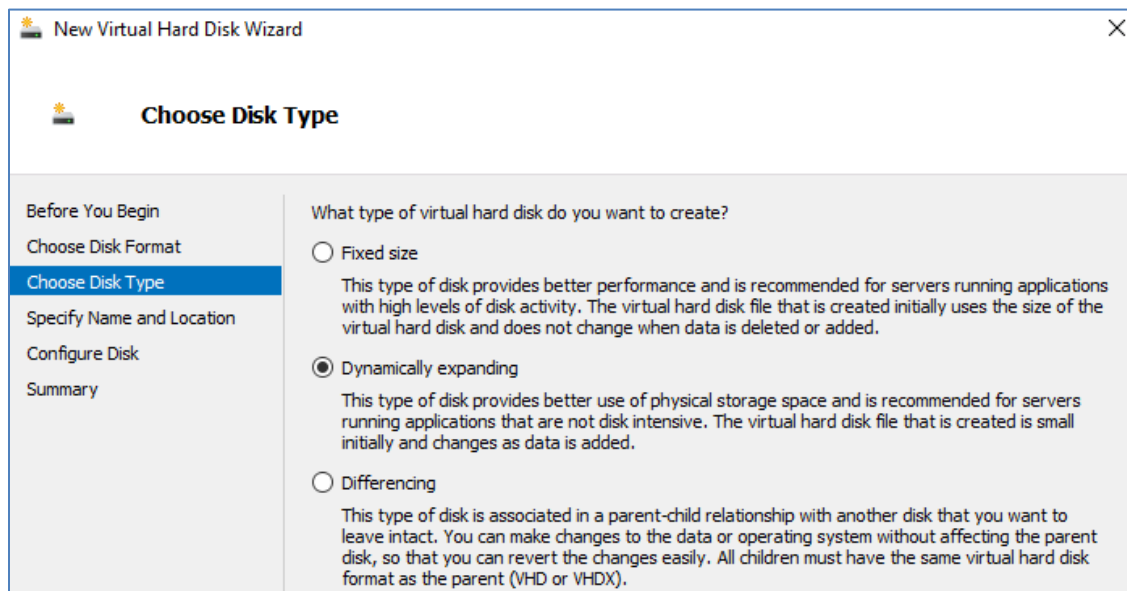


Figure 5 Select a virtual hard disk type

The dynamically expanding disk type will work well for most workloads. Since SC Series storage leverages thin provisioning, only data that is actually written to a virtual hard disk, regardless of the disk type, will consume space on the array. As a result, determining the best disk type is mostly a function of the workload as opposed to how it will impact storage capacity. For workloads generating very high I/O, such as SQL Server databases, Microsoft recommends using the fixed size disk type.

As shown below, 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 consumed is equal to amount of data on the virtual disk, and is therefore more space efficient from the perspective of the host. From the perspective of the guest VM, either type of virtual hard disk will appear as a full 60 GB of available space.

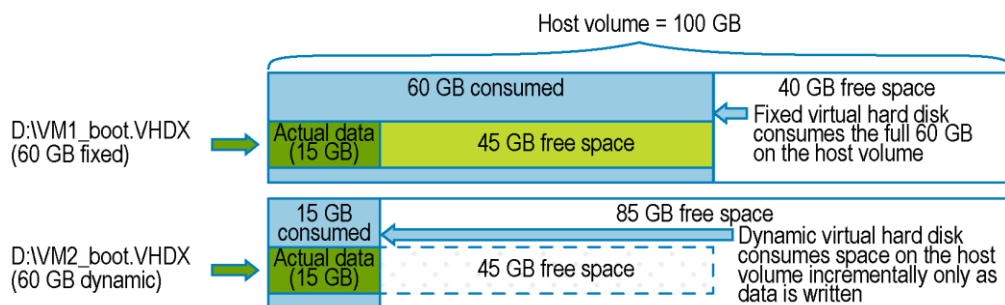


Figure 6 Fixed and dynamic virtual hard disk comparison

There are some best practice performance and management considerations to keep in mind when choosing the right kind of virtual hard disk type for your environment.

- Fixed-size virtual hard disks:
 - Recommended for virtual hard disks that experience a high level of disk activity, such as SQL Server, Microsoft Exchange or OS page or swap files. For many workloads, the performance difference between fixed and dynamic will be unnoticeable.
 - When formatted, they take up the full amount of space on the host server volume.
 - They are less susceptible to fragmentation at the host level.
 - Take longer to copy (for example, from one host server to another over the network) because the file size is the same as the formatted size.
 - With older versions of Hyper-V prior to 2012, provisioning of fixed virtual hard disks may require significant time due to lack of native Offloaded Data Transfer (ODX) support. With Server 2012 and newer, coupled with SCOS 6.3 and newer, provisioning time for fixed virtual hard disks on SC Series volumes is significantly reduced when ODX is supported and enabled.
- Dynamically expanding virtual hard disks:
 - Recommended for most virtual hard disks, except in cases of extremely high disk IO.
 - Require slightly more CPU and I/O overhead as they grow than fixed virtual hard disks. This usually does not impact the workload except in cases where I/O demand is very high.
 - They are more susceptible to fragmentation at the host level.
 - When initially formatted, they consume very little space, and expand as new data is written to them by the guest VM.
 - Take less time to copy (for example, over the network) than fixed virtual hard disks.
 - Allows the host server's volume to be over-provisioned. In this case, it is important to configure alerting on the host server's volume to avoid running it out of space unintentionally.
- Differencing virtual hard disks:
 - Offers storage savings by allowing multiple Hyper-V guest VMs with identical operating systems share a common boot virtual hard disk.
 - All children must use the same virtual hard disk format as the parent.
 - Reads of unchanged data reference the differencing virtual hard disk. Unchanged data that is read infrequently may reside in a lower tier of storage on an SC Series providing more efficient SAN utilization.
 - New data is written to the child virtual hard disk, which by default will be written to the highest performing tier on the SC Series for maximum performance.

Each native Hyper-V based snapshot of a Hyper-V guest VM creates a differencing virtual hard disk to freeze with changed data since the last snapshot, and new data is written to a new virtual hard disk file. Creating native Hyper-V based snapshots of a Hyper-V guest VM can elevate the CPU usage of storage I/O, but will probably not affect performance noticeably unless the guest VM experiences very high I/O demands.

Maintaining a long chain of native Hyper-V based snapshots of a Hyper-V guest VM can negatively affect the performance of the guest because reading from the virtual hard disk can require checking for the requested blocks in a chain of many different differencing virtual hard disks. It is a best practice to keep native Hyper-V based snapshots of Hyper-V guests to a minimum to maintain optimal disk I/O performance.

Use SC Series snapshots and Replay Manager VSS-aware backups of Hyper-V guests to avoid the negative impact on I/O that can be caused by maintaining a long chain of differencing virtual hard disks.

3.3.3 Virtual hard disks and thin provisioning with SC Series storage

It does not matter which type of virtual hard disk is used with regard to maximizing the space utilization on SC Series storage. Regardless of the virtual hard disk type, only the actual data written by a guest VM will consume space on the SC Series.

The example below illustrates a SAN-based 100 GB volume presented to a Hyper-V host that contains two 60 GB virtual hard disks. One is fixed, and the other is dynamic. Each virtual hard disk contains 15 GB of actual data. From the perspective of the host server, 75 GB of space is consumed:

$$60 \text{ GB fixed disk} + 15 \text{ GB of used space on the dynamic disk} = 75 \text{ GB}$$

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

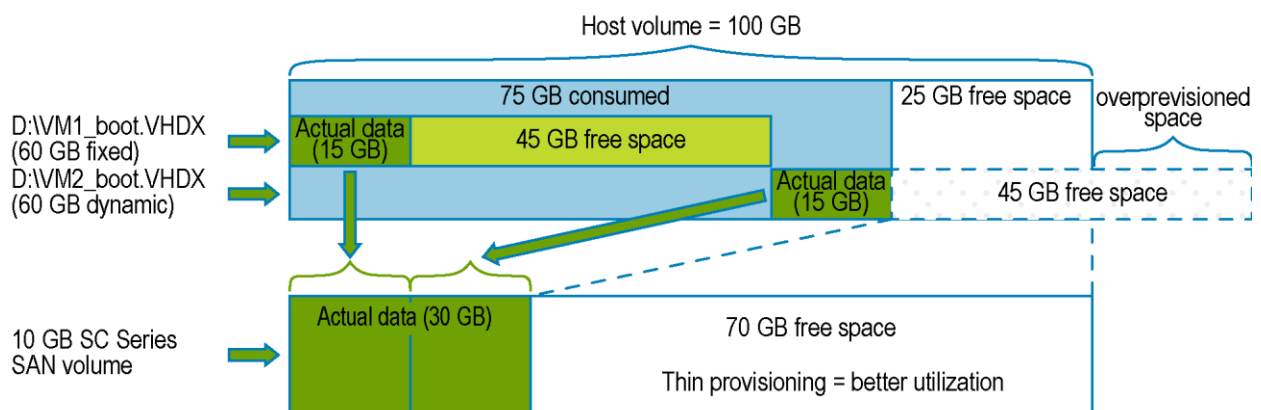


Figure 7 Thin provisioning with SC Series

Comparatively, this is how the SC Series reports storage utilization on this same volume:

$$15 \text{ GB of used space on the fixed disk} + 15 \text{ GB of used space on the dynamic disk} = 30 \text{ GB}$$

Note: Because of thin provisioning, regardless of the type of virtual hard disk, space on the SC Series is consumed only as actual data is written. Choosing dynamic over fixed therefore does not improve storage space utilization on the SC Series. Other factors such as the I/O performance of the workload would be primary considerations when determining the type of virtual hard disk in your environment.

3.3.4 Overprovisioning with dynamic virtual hard disks

With dynamic virtual hard disks and thin provisioning, there is an inherent risk of either the host volume or a storage pool on the SC Series running out of space. The figure in the preceding section shows an example of this. If the dynamic disk used by VM2 on the host volume expanded far enough, it would fill up the host volume and negatively impact one or both VM1 and VM2. From the perspective of VM2, it would still see

20 GB of free space, but would not be able to use it because the underlying host volume would be full. To resolve this, an administrator would need to expand the 100 GB host volume.

To mitigate risks, consider the following best practice recommendations:

- Create a Hyper-V host volume that is large enough so that current and future expanding dynamic virtual hard disks will not ever fill the volume completely. Creating large Hyper-V host volumes will not waste space on the SC Series because it leverages thin provisioning.
 - If Hyper-V based snapshots are used, allow adequate overhead on the host volume for the extra space consumed by the snapshots.
 - Expand existing host volumes as needed to avoid the risks associated with overprovisioning.
 - As a minimum, if a volume is overprovisioned, set up monitoring so that if a percent full threshold is exceeded, an alert is generated with enough lead time to allow for remediation.
- Configure thresholds and alerting on the SC Series so that warnings are generated before a storage tier reaches its capacity. It is important to set the alert threshold low enough to allow enough lead time for remediation so that the SC Series does not enter conservation mode.

3.4 Present SC Series storage to Hyper-V

There are a number of different ways to present SC Series storage volumes (LUNs) to Hyper-V hosts, nodes and VMs. Each will be explained in the sections that follow, along with some best practice recommendations.

- Boot from SAN is supported by all versions of physical Hyper-V hosts and nodes except for Nano server.
- Volumes can mapped directly to physical Hyper-V hosts and nodes (including to SC Series cluster server objects) using the following transports:
 - Fibre Channel
 - iSCSI
 - Front-end SAS (with select SC Series models only)
- VMs support direct-attached storage (storage presented directly to a VM):
 - In-guest iSCSI
 - Virtual Fibre Channel (FC) for Server 2012 Hyper-V and newer (exception: Replay Manager does not support virtual FC with Server 2012 Core)
- Pass-through disks

3.4.1 Transport options

SC Series storage is presented as data volumes to Hyper-V environments using iSCSI, Fibre Channel or front-end SAS, and will typically include multipath I/O (MPIO) for load balancing and failover protection.

Typically, an environment is configured to use a preferred transport 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. Since Hyper-V supports iSCSI, Fibre Channel and front-end SAS on new environments, deciding which transport to use is usually 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. This might be required to support collocated but diverse platforms with different transport requirements. When this is the case, administrators might be able to choose between several different transport options.

3.4.2 SC Series and front-end SAS support for Hyper-V

All SC Series arrays can be configured to support iSCSI or Fibre Channel for front-end connectivity. In addition, SCv2000 and SC4020 arrays can be configured to support front-end SAS connectivity. Front-end SAS may work well for small Hyper-V environments such as a branch office or other remote location where the cost and overhead of maintaining infrastructure to support Fibre Channel or iSCSI can be avoided.

A separate whitepaper is dedicated to configuring Hyper-V to work with front-end SAS and provides several different step-by-step configuration examples. See the [Configuring Microsoft Hyper-V Hosts to Access Dell SCv2000 and SC4020 Arrays with SAS Front-end HBAs](#) on Dell TechCenter for more information.

3.4.3 Mixed transports

Although SC Series storage arrays support mixed transports, there is limited Microsoft support. For example, when an SC Series volume is mapped to a Server 2012 or newer host using both iSCSI and Fibre Channel paths at the same time, the host may initially see both the iSCSI and Fibre Channel paths, but choose to use the paths for only one transport based on an enumeration algorithm. The algorithm typically results in the host choosing to use the Fibre Channel paths instead of the iSCSI paths.

While Server 2008 R2 will use all paths regardless of the transport, using mixed transports is typically unnecessary. Each transport by itself will usually present adequate bandwidth and redundant paths that leverage MPIO. Using multiple transports increases design complexity unnecessarily.

3.4.4 MPIO best practices

The Windows Server operating system and Hyper-V (2008 and newer) natively support MPIO by the built in Device Specific Module (DSM) provided by Microsoft. SC Storage supports the use of this DSM when implementing MPIO.

It is very important to adjust MPIO timeout settings for your Hyper-V environment (both hosts and VM) to avoid service outages when performing routine SAN maintenance. A separate whitepaper is dedicated to this topic and is updated frequently. See the [Dell Storage Center: Microsoft Multipath I/O Best Practices](#) guide on Dell TechCenter for more information.

Note: Adjusting MPIO settings on Windows and Hyper-V servers is sometimes overlooked as a critical part of the design. Failure to adjust MPIO timeout settings can cause unintended service outages due to loss of access to SAN storage when performing what would otherwise be non-service-affecting events such as performing an SCOS upgrade.

3.4.5 Guest VMs and in-guest iSCSI and virtual Fibre Channel disks

In most cases, storage is presented to guest VMs as a virtual hard disk. Direct-attached disks are also supported but are generally discouraged unless there is a specific use case that requires them. They are mapped directly to a guest VM and are therefore not managed by the Hyper-V host server.

Direct-attached disks are of two types:

- **In-guest iSCSI:** Simply use the iSCSI initiator software in the guest VM and configure the iSCSI target (the SC Series controllers). Native iSCSI support is provided with Server 2008 and above.
 - More information on configuring iSCSI can be found in the [Dell Storage Center: Microsoft Multipath I/O Best Practices](#) guide on Dell TechCenter.
- **Virtual Fibre Channel:** With Server 2012 Hyper-V and newer, Windows Server guest VMs (2008 R2 and newer) support virtual Fibre Channel (FC) adapters. This functionality was added by Microsoft in large part because many environments at the time used Fibre Channel exclusively and shared virtual hard disks were not yet available.
 - Virtual FC requires that all the components in the fabric (HBAs, Switches, SC Series controllers) support N_Port ID Virtualization (NPIV) and have the feature enabled.
 - There is a demo video dedicated to how virtual FC works, and provides detailed step-by-step guidance on how to configure Hyper-V nodes to support a guest VM cluster. See the [Virtual Fibre Channel for Hyper-V](#) Demo Video on Dell TechCenter for more information.

3.4.6 Direct attached storage use cases

Although SC Series arrays fully support both iSCSI and virtual FC disks mapped to guest VMs, direct-attached storage for guest VMs is generally not recommended as a best practice unless there is a specific use case that requires it. Typical use cases include:

- When the Replay Manager agent is installed directly on a guest VM in order to obtain VSS-consistent backups of SQL Server or Microsoft Exchange data, this data must reside on a direct-attached volume that uses in-guest iSCSI or virtual Fibre Channel (with the exception that Replay Manager does not support virtual FC with Server 2012 Core).
- Any recovery volumes presented to a guest VM by Replay Manager must be a direct-attached volume that uses in-guest iSCSI or virtual FC.
- Situations where a workload has very high I/O requirements and the little bit of extra performance gained over using a virtual hard disk matters. Direct-attached disks bypass the host server file system. This reduces host CPU overhead for managing guest VM IO. For many workloads, there will be no notable difference in performance between direct-attached and virtual hard disks.
- VM clustering on legacy platforms prior to support for shared virtual hard disks, which became available with the 2012 R2 release of Hyper-V, and enhanced with Hyper-V 2016.
- When needing to troubleshoot I/O performance on a particular volume and it must be isolated from all other servers and workloads. This can be done by copying the data to a dedicated direct-attached disk temporarily.
- When there is a need to create a custom SC Series storage profile, snapshot profile, or replication profile for just that subset of data. This can also be accomplished by placing a virtual hard disk on a host volume that is assigned the desired SC Series profiles.

- When a data volume exceeds the maximum size for a VHD (2 TB) or VHDX virtual hard disk (64 TB).

There are also some disadvantages to using direct-attached storage for guest VMs:

- The ability to perform native Hyper-V snapshots is lost. However, the ability to leverage SC Series snapshots is unaffected.
- More complicated and therefore requires more management overhead. This is especially true for virtual FC.
- VM Mobility is reduced due to creating a physical hardware layer dependency.
- Replay Manager does not support virtual Fibre Channel with Server 2012 Core.
- Virtual FC requires that the host servers and SC Series use soft zoning (by WWN) instead of hard zoning (by switch port). Attempting to boot a guest VM with a virtual FC adapter may crash the host server if the fabric is hard-zoned instead of soft-zoned. Verify the configuration of your fabric before creating virtual FC adapters.
- SC Series does not support using a cluster server object for guest VM clusters using Virtual FC adapters because of the unique way that Microsoft leverages two sets of WWNs for each virtual FC adapter.

Note: Legacy environments that are using direct-attached disks for guest VM clustering should consider switching to shared virtual hard disks.

3.4.7 Guest VMs and pass-through disks

A pass-through disk is a SAN volume mapped to a Hyper-V host, node or cluster that is “passed through” directly to a Hyper-V guest VM. The Hyper-V host, node or cluster has visibility to a pass-through disk but does not have I/O access. Hyper-V keeps it in a reserved state as only the guest VM has I/O access.

Using pass-through disks is discouraged unless there is a specific use case that requires it. Valid use cases are similar to the list provided for direct attached storage in the above section.

A common use case for pass-through disks is when Replay Manager is used to protect SQL Server or Microsoft Exchange workloads running on a guest VM, and the Replay Manager agent is installed directly on the VM. Although pass-through is supported, using in-guest iSCSI is recommended, since in-guest iSCSI is required regardless in order to perform a recovery. For more information, see the [Replay Manager 7 and Hyper-V Best Practices Guide](#) on Dell TechCenter.

Some reasons to avoid pass-through disks include:

- Similar to direct-attached storage, the ability to perform native Hyper-V snapshots is lost.
- The use of a pass-through disk as a boot volume on a guest VM prevents the use of a differencing disk. However, SC Series View Volumes (created from a gold image volume) can still be used to maximizing SAN space utilization.
- In large environments with a high number of cluster nodes and guests, it is possible to exhaust the pool of available LUN numbers on your hosts when using pass-through disks. Cluster server objects on a Dell SC Series support a maximum of 254 LUNs. Avoid this limitation by using direct-attached disks.
- Reduces VM mobility by creating a dependency on the physical layer.

3.4.8 SC Series arrays and cluster server objects

With SC Series storage, create and use cluster server objects on the array for your Hyper-V and guest VM clusters (both physical and virtual) whenever possible, except in the case of VM clusters that use virtual FC adapters.

In a clustered Hyper-V environment, each cluster volume needs to be mapped to all nodes in the cluster using the same LUN number. When a cluster object represents two or more nodes, mapping the storage to the cluster object provides the following benefits:

- The volume is mapped to all nodes in the cluster automatically so none get missed
- The volume is mapped to each node using a consistent LUN number
- Reduces the chance for errors or inconsistencies with volume mappings
- Saves time, particularly as the number of nodes increases
- Cluster volumes are automatically mapped to new cluster nodes that are added to the SC Series cluster object

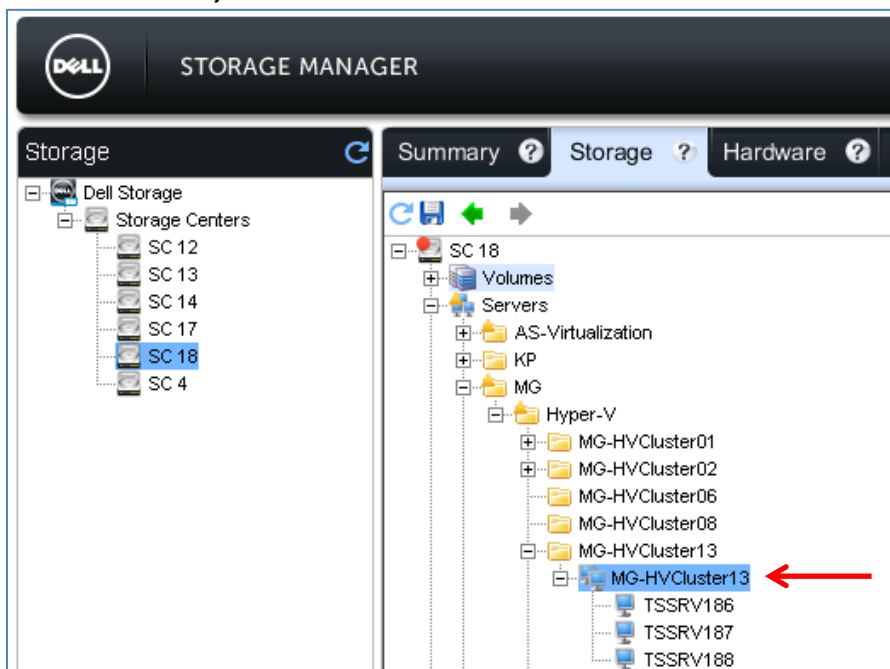


Figure 8 Map storage to a cluster object to ensure consistent LUN numbers on all nodes

3.4.9 LUN limits for larger Hyper-V clusters

For large Hyper-V clusters with many nodes, there are two LUN limits to be aware of:

- The functional LUN limit: Although the SC Series supports up to 254 LUNs per cluster server object, resources on Hyper-V server nodes might be consumed before the physical limit of 254 LUNs reached. This will vary depending on the capacity of the hardware and the workload.
- The physical LUN Limit: Depending on the Hyper-V cluster design (for example, if using pass-through disks), it is possible to consume many LUN numbers quickly, exhausting the pool of free LUN numbers.

It is also important to note that a small number of available free LUN numbers must be kept in reserve for an administrator to use for scratch, temporary volumes, and SAN maintenance. SAN maintenance might include operations such as expiring snapshots or Replay Manager restore operations that require LUN numbers on a temporary basis when presenting View Volumes to a host using in-guest iSCSI.

To avoid reaching the LUN number functional or physical limit for a Hyper-V cluster, consider some of the following strategies:

- Use a many-to-1 guests-to-CSV strategy using virtual hard disks
- Use direct-attached storage (iSCSI or virtual Fibre Channel) to present SAN volumes directly to guest VMs instead of pass-through disks
- Add physical FC or iSCSI cards to the cluster nodes to expand the functional limit. This will not increase the 254 physical LUN limit per SC Series cluster
- If using pass-through disks, create smaller Hyper-V clusters with fewer nodes and fewer guest VMs instead of larger clusters with more guests

3.4.10 Volume design considerations for SC Series

One of the design considerations for which there are often no clear answers is how many guest VMs to place on an SC Series volume, including cluster shared volumes (CSV). In most cases, a many-to-one strategy is a good starting point, and adjust for specific uses cases.

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

- Fewer SC Series volumes to create and administer (avoid volume sprawl).
- Easier to provision additional guest VMs due to not having to create new SC Series volumes each time.
- Avoid the physical and functional LUN number limits for Hyper-V clusters.

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

- Easier to track disk I/O patterns per Hyper-V guest VM.
- Ability to quickly restore a guest VM by simply replacing the original host volume or CSV with a View Volume from an SC Series snapshot.
- If SC Series volumes are replicated to a 2nd location, an administrator has more granular control over what data is getting replicated.
- It is often quicker to move a guest VM from one host or cluster to another by remapping the volume rather than copying 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.

3.5 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 SC Series controllers. This feature is fully supported in Hyper-V environments. Dell SC Series controllers running SCOS 6.3.1 and newer support ODX with 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.

ODX is also leveraged by Microsoft System Center Virtual Machine Manager (SCVMM) environments. Progress bars will indicate that a copy operation is rapid copy when ODX is being leveraged to deploy a new VM over the network from a template on the library server.

For more information, such as how to enable or disable ODX and how to establish performance benchmarks, view the [Windows Server Best Practices Guide for Dell SC Storage](#) on Dell TechCenter.

3.6 Disable automount

To prevent a Hyper-V host server from automatically assigning drive letters to newly mapped volumes, disable the automount feature, which is enabled by default. Having automount disabled is also beneficial when recovering a volume using an SC Series snapshot.

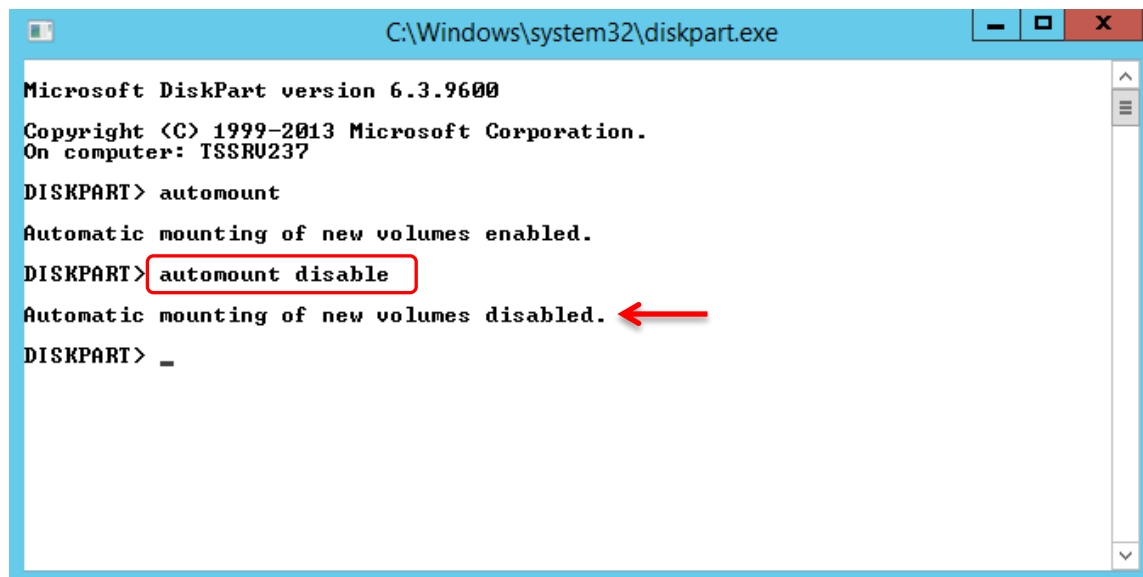


Figure 9 Disable the automount feature

To verify or change a server automount configuration, open a command prompt window and run Diskpart. Use the **automount** command to verify the current state of automount, and the **automount disable** command to disable it.

3.7 Placement of page files

Windows Servers and VMs place the page file on the boot volume by default, and automatically manage page file and memory settings without user intervention. In most cases, these settings should not be changed unless, for example, an application vendor provides specific guidance on how to tune the page file and memory settings to optimize the performance of a particular workload.

Ultimately each customer will need to decide on the best strategy as they consider variables that are unique to their environment.

With SC Series, there can be some advantages to placing a page file on a separate volume from the perspective of the storage array. The following reasons are probably not sufficiently advantageous by themselves to justify varying from the defaults, but in cases where a vendor recommends making changes to optimize a workload, consider the following tips as part of the overall page file strategy.

- Moving the page file to a separate dedicated volume reduces the amount of data that is changing on the system (boot) volume. This can help reduce the size of SC Storage snapshots of boot volumes which will conserve SAN space.
- 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 SC Series. This is especially beneficial in cases where there is limited bandwidth for replication of volumes and snapshots to other SC Series.
- Since page file data is constantly changing, it benefits from staying in Tier 1 for maximum performance. To prevent data progression of page file data to lower tiers, volumes containing only page files can be assigned a storage profile that keeps the data in Tier 1.

3.8 SC Series data reduction and Hyper-V

Data compression was introduced with SCOS version 6.5.1, and enhancements were added with the SCOS 6.5.10 and 6.7 releases. Data deduplication was introduced with SCOS 7.0, and works in tandem with compression to further reduce the amount of data stored on an SC Series.

Data reduction works seamlessly in the background as part of the SC Series daily data progression cycle each evening. Hyper-V environments benefit from data reduction without any additional configuration required.

Data reduction can be enabled, paused or discontinued on a volume by volume basis. It can also be paused system wide. For more information on how data reduction with deduplication and compression work with SC Series storage, see the [Dell Storage Center OS 7.0 Data Reduction with Deduplication and Compression Guide](#) on Dell TechCenter.

4 SC Series snapshots and Hyper-V

SC Series point-in-time snapshots allow administrators to do the following in Hyper-V environments:

- Recover servers to a crash-consistent state, including Hyper-V hosts and guest VM workloads
- Provision lab or isolated test environments using View Volumes
- Provision new servers using gold images

Point-in-time SC Series snapshots can be taken of SC Series volumes regardless of content. This applies to boot-from-SAN disks, data disks, cluster shared volumes containing virtual hard disks, pass-through disks, direct-attached in-guest iSCSI disks and virtual FC disks. Not only can snapshots be taken, but these disks and volumes and their associated snapshots can be quickly and easily replicated to other SC Series arrays for DR or other purposes.

Unless the server is powered off at the time the snapshot is taken, or put into a consistent state by Replay Manager or some other volume shadow copy (VSS) aware application, SC Series snapshots are considered to be crash-consistent. When recovering a server using a crash consistent snapshot, it is similar to having the server recover from a power outage at that point in time. In most cases, servers and applications are resilient enough to recover to a crash consistent state without any issues, whether the cause is an unexpected power outage, or the server is being recovered to a previous point in time. An exception to this is when the Hyper-V environment hosts a transactional workload such as Microsoft Exchange or SQL Server. With transactional workloads, the risk of data corruption is higher when attempting to recover to a crash consistent state.

Some examples for how to configure and use SC Series snapshots for Hyper-V environment are provided below.

4.1 Replay Manager support for Hyper-V

Replay Manager can be used to obtain application-consistent snapshots (Replays) of Hyper-V guests. The Hyper-V extension for Replay Manager leverages VSS to ensure that a guest VM is in a consistent state before a snapshot is taken. Replay Manager also includes extensions to protect Microsoft Exchange, SQL Server and VMware VMs.

There are many documents and videos on Dell TechCenter that showcase the features of Replay Manager. Two that are specific to Hyper-V include the [Replay Manager 7 and Hyper-V Demo Video](#) and the [Replay Manager 7 and Hyper-V Best Practices Guide](#).

4.2 Use SC Series snapshots to recover guest VMs

Hyper-V guest VMs can be recovered to a previous point in time by using crash-consistent SC Series snapshots. Snapshots can also be used to create copies of VMs in an isolated environment at the same or a different location when volume replication between SC Series is used. This section provides guidance and best practices for several different recovery options using snapshots.

4.2.1 Recover a guest VM on a standalone 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.

If a VM hard disk and configuration files reside on separate host data volumes, then it is a best practice to configure a consistency group for these volumes on the SC Series so that crash-consistent snapshots occur at exactly the same time. 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.

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

- Option 1: Replace the existing data volume on the host that contains the VM's configuration and virtual hard disks with a View Volume created from the desired SC Series snapshot. In this scenario, the VM in question is powered down, the original volume is unmapped from the server, and the View Volume is mapped to the host using the same LUN number, drive letter mapping, or mount point.
 - This may only be practical if the data volume contains only one VM. 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 just one VM.
 - This will allow the VM being recovered to 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 the View Volume containing the VM's 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. The VM can be recovered by manually copying the virtual hard disks from the View Volume to the original location.
 - Delete, move, or rename the original virtual hard disks.
 - After copying the recovered virtual hard disks to their original location, rename them and use Hyper-V manager to re-associate them with the guest VM. This is necessary to allow the guest VM to start without permissions errors.
 - This may not be practical if the virtual hard disks are extremely large. In this case, the original VM can be deleted, and the recovery VM imported (Hyper-V 2012 and newer) or created as a new VM (Hyper-V 2008 R2) directly from the View 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 View Volume to a different Hyper-V host and recover the VM there by importing the VM configuration or creating a new VM that points to the virtual hard disks on the recovery volume.
 - This is common in situations where the original VM and the recovery VM both need to be online at the same time, but be isolated from each other, or when the original host server is no longer available.

If possible, before beginning any VM recovery, record essential details about the VM's hardware configuration (such as CPUs, RAM, Virtual networks, IP addresses) in case importing a VM's configuration is not supported (Hyper-V 2008 R2) or fails.

4.2.2 Recover guest VM on a cluster shared volume

The process of using SC Series snapshots to recover guest VMs that reside on a cluster shared volume (CSV) is similar to the process of recovering a guest VM to standalone host, as detailed in the preceding section. However, recovering a VM from 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 an 8-character hexadecimal number such as 045C3E2F4. No two volumes mapped to a server can have the same disk ID.

When an SC Series 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. Therefore, View Volumes created from an SC Series snapshot will also have the same disk ID.

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

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

When attempting to map a View Volume of a CSV back to any server in that same cluster, the View Volume will cause a disk ID conflict, which can be potentially service-affecting.

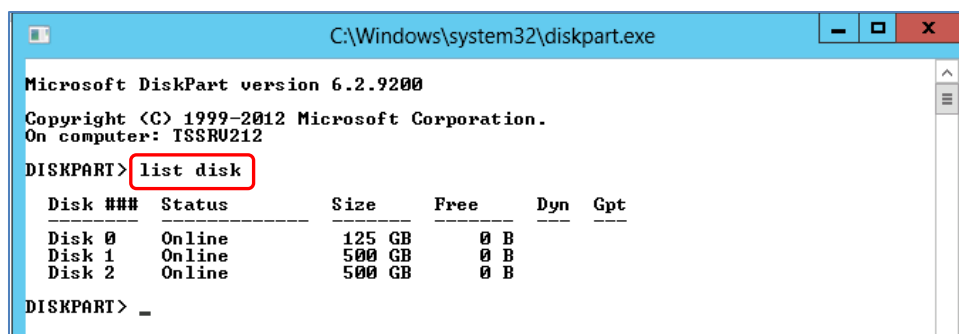
There are a couple of ways of working around the duplicate disk ID issue as detailed below.

- Option 1: Map the View Volume of the CSV to another host that is outside of the cluster and copy the guest VM files over the network to recover the guest.
- Option 2: Map the View Volume to another Windows host outside of the cluster and use Diskpart.exe to change the disk ID. Once the ID has been changed, re-map the View Volume to the cluster. The steps to use Diskpart.exe to change the disk ID are detailed below.

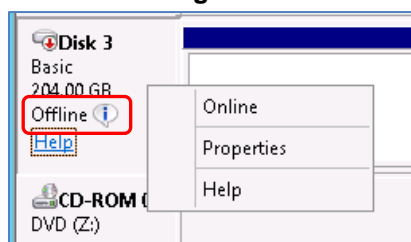
4.2.3 Change a cluster shared volume's disk ID with Diskpart

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

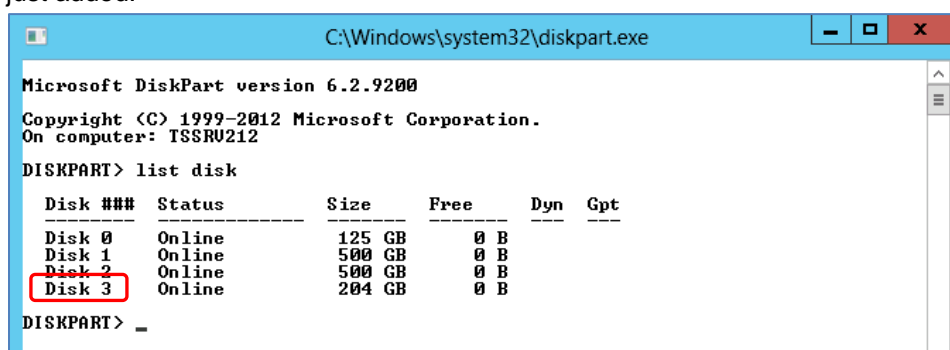
1. Access the standalone Windows host that the View Volume of the CSV will be mapped to.
2. Open a command window with administrator rights.
3. Type **diskpart.exe** and press **[Enter]**.
4. Type **list disk** and press **Enter**.
5. Make note of the current list of disks (in this example, Disk 0, Disk 1, Disk 2).



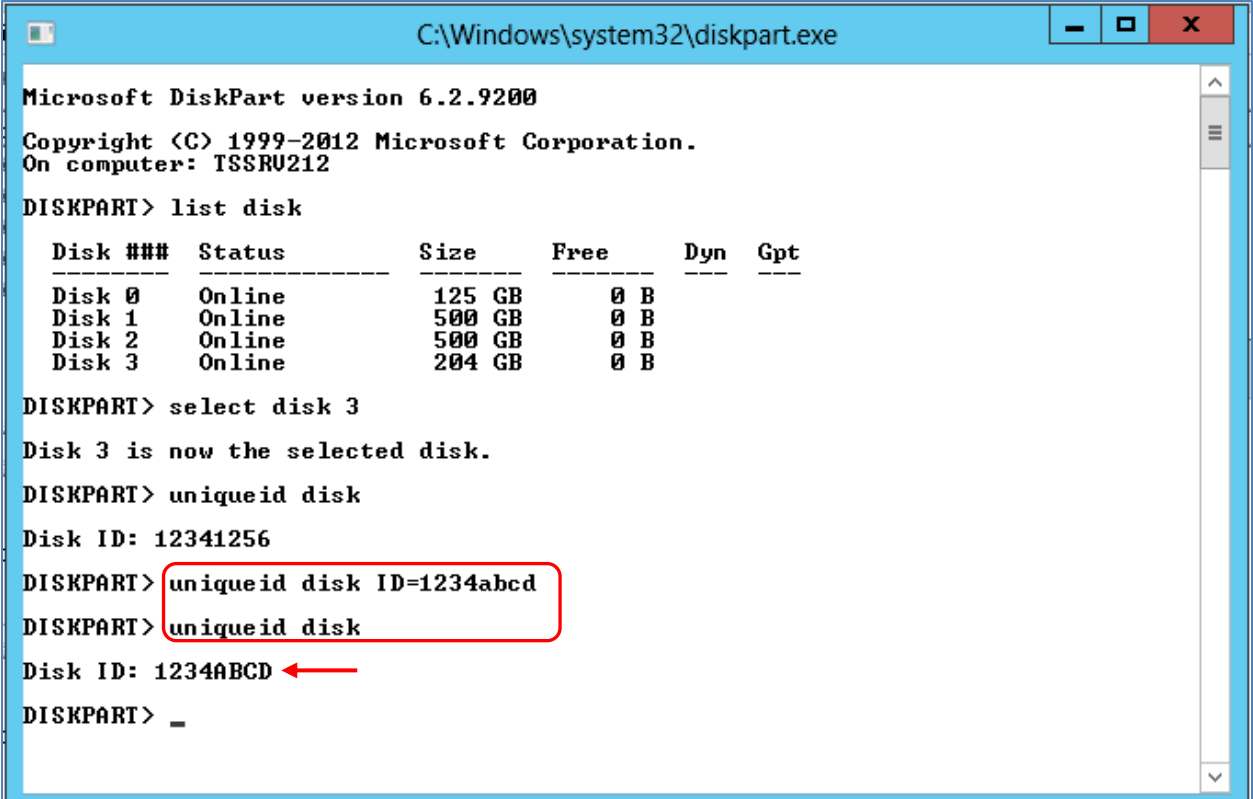
6. Using the Dell Storage Manager Client, map the View Volume of the CSV to this host.
7. From the **Diskpart** command prompt, type **rescan** and press **Enter**.
8. Use **Disk Management** on the host server to bring the View Volume online.



9. Return to the **diskpart** command prompt window and type **list disk** and press **[Enter]**.
10. The new disk (Disk 3 in this example) should now be listed. Usually, the bottom disk will be the one just added.



11. Type **select disk #** (where # represents the number of the new disk, in this example, disk 3) and then press **Enter**.
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**.
 - a. For **<newid>** provide a random ID of your choice. For an MBR disk, the new ID must be an 8-character string in hexadecimal format using a mix of the numbers 0 – 9 and the letters A – F.
 - b. For a GPT disk, the new ID must be a Globally Unique Identifier (GUID).
14. Type **uniqueid disk** again and press **Enter** to verify the ID is now changed.
15. Now that the View Volume has a new signature, it can be unmapped from the standalone host server and re-mapped to the cluster without causing a disk ID conflict.
16. Recover the guest VM.



```
C:\Windows\system32\diskpart.exe

Microsoft DiskPart version 6.2.9200
Copyright (C) 1999-2012 Microsoft Corporation.
On computer: TSSRU212

DISKPART> list disk

   Disk ###  Status       Size      Free      Dyn  Gpt
   -----  -
   Disk 0      Online        125 GB     0 B
   Disk 1      Online        500 GB     0 B
   Disk 2      Online        500 GB     0 B
   Disk 3      Online        204 GB     0 B

DISKPART> select disk 3
Disk 3 is now the selected disk.

DISKPART> uniqueid disk
Disk ID: 12341256
DISKPART> uniqueid disk ID=1234abcd
DISKPART> uniqueid disk
Disk ID: 1234ABCD
DISKPART> _
```

Figure 10 Change a disk ID with Diskpart

4.3 Use SC Series snapshots to create a test environment

In addition to VM recovery, SC Series snapshots can be used to quickly create test or development environments that mirror a production environment by mapping View Volumes to other server hosts or clusters. When SC Series snapshots containing VMs are replicated to another location, this makes it very easy to do this at a different location.

Note: To avoid IP, MAC address or server name conflicts, copies of existing VMs recovered from a View Volume should be placed in an isolated environment.

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

4.4 Leverage SC Series to create gold images

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

- Host servers (when using Boot from SAN)

Note: Nano Server does not support Boot-from-SAN. For more information on Nano Server and SC Series storage, see the [Windows Server Best Practices Guide](#) on Dell TechCenter.

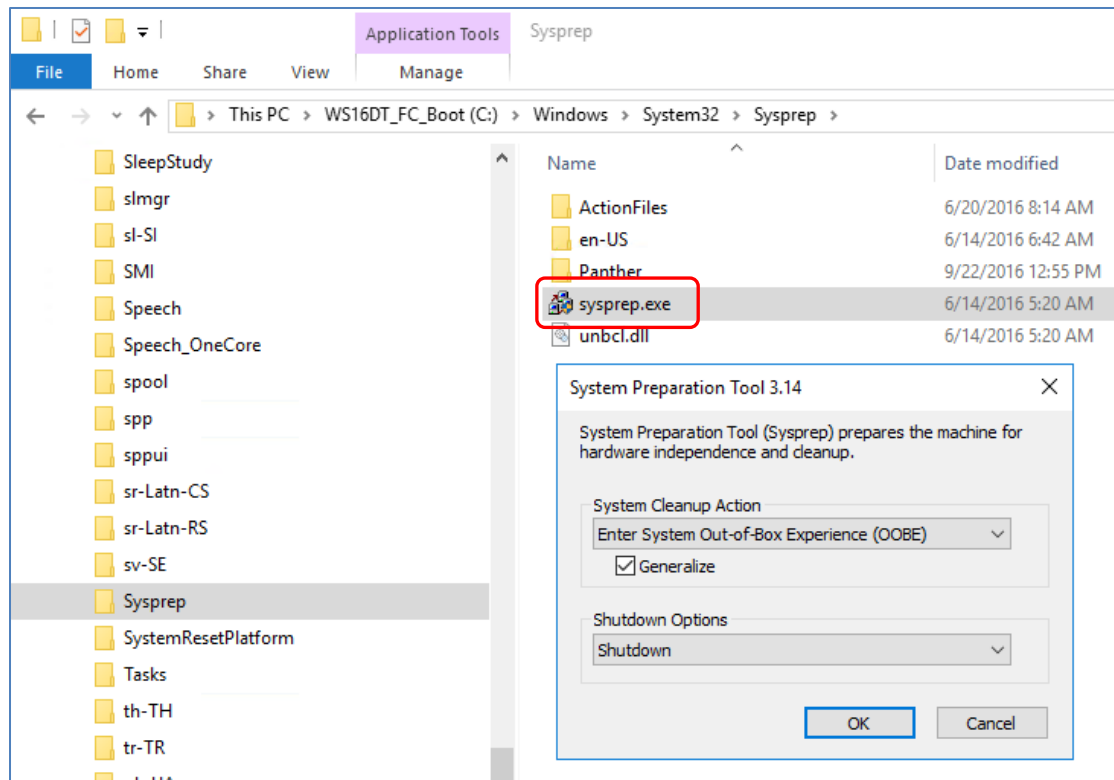
- Guest VMs that use pass-through disks for boot disks
- Guest VMs that boot from a virtual hard disk (although, in this case, using a sysprepped virtual hard disk as the gold image source would be preferred).

Using gold images provides the following benefits:

- Faster server provisioning with minimal reconfiguration.
- Better SAN utilization. When a host or VM is provisioned from a gold image, only new data consumes SAN space. Data that has not changed is read from the gold image source volume.
- Disk tiering and data progression optimize data placement for best performance and utilization.

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

1. Create and map an SC Series volume to a host server.
2. Build your base OS image, install roles and features, and fully configure and patch it. This will minimize the changes that have to be made to each new server that is deployed using the gold image.
3. Once the OS is fully staged, power down the OS to put it into a consistent state and then take a manual SC Series snapshot of the volume and set this snapshot to never expire. This represents the point in time prior to running Sysprep. If the OS image needs to be updated in the future (for example, to apply patches), it can then be used to create a new gold image source without having to stage the OS from scratch.
4. Power on the server and run Sysprep, choosing the Generalize, Out-of-box Experience, and Shutdown options.
 - a. After running Sysprep, if the server is a guest VM, Use Hyper-V Manager to delete the guest VM. This will delete the guest VM's configuration files but will preserve the boot virtual hard disk. The virtual hard disk is the only file needed for the gold image.
 - b. Do not use Microsoft System Center Virtual Machine Manager (SCVMM) to delete the guest VM as it will also delete the virtual hard disk file. Use Hyper-V Manager instead.



5. Once the server is powered down (which will ensure that it is in a consistent state), manually create another SC Series 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 this 2nd snapshot as the gold source, create a View Volume and map it to the desired host server.
If mapping View Volumes from a gold source to a Hyper-V cluster, duplicate disk IDs will cause a conflict. Change the disk ID of the volume prior to mapping it to the cluster. See Section [put section number here once heading number issues are fixed in this doc] for more information.
7. Boot the host server (or create a new VM pointed to the Sysprepped pass-through or virtual hard disk as appropriate) and allow the initial boot process to complete.
8. Customize the server configuration as needed.
9. Leverage PowerShell to automate the workflow if desired.

4.4.1 Gold images and preserving balanced SC Series controllers

When a new volume on an SC Series array consisting of two controllers is created and mapped to a host server, cluster or VM, volume ownership is assigned to one controller or the other - not both. As additional volumes are created and mapped to servers from the SC Series, controller ownership is alternated automatically in a round robin fashion so that both controllers stay evenly balanced with each controller owning roughly the same number of volumes. Administrators can override this default behavior and assign volumes to the controller of their choice when mapping them to servers, but this is usually not required.

During SC Series maintenance that requires staggered controller reboots, ownership of all volumes is temporarily moved to one controller while the other controller is rebooted. Once both controllers have been rebooted, they are rebalanced using the Dell Storage Manager client.

When using gold images to deploy new host servers, there is a limitation to be aware of. Since a View Volume is created from a snapshot of the source volume, the SC Series controller that owns the source volume will also own all View Volumes created from it. If many new hosts are deployed from the same gold image, this can result in one controller owning many more volumes than the other controller head, creating an imbalance.

Use the Dell Storage Manager client to view the summary information for a volume that is the gold source to see which controller owns it. Change to Tree View under the Snapshots tab in the Dell Storage Manager client to see a graphical representation of the relationship between the gold image source volume and any View Volumes created from it.

In the example below, Storage Center **SC18** is comprised of two controllers: SN 716 and SN 717. Two new server hosts named TSSRV200 and TSSRV201 have been provisioned from a boot-from-SAN gold image source volume that is owned by controller 717. As a result, the View Volumes created for hosts TSSRV200 and TSSRV201 are also owned by controller 717.

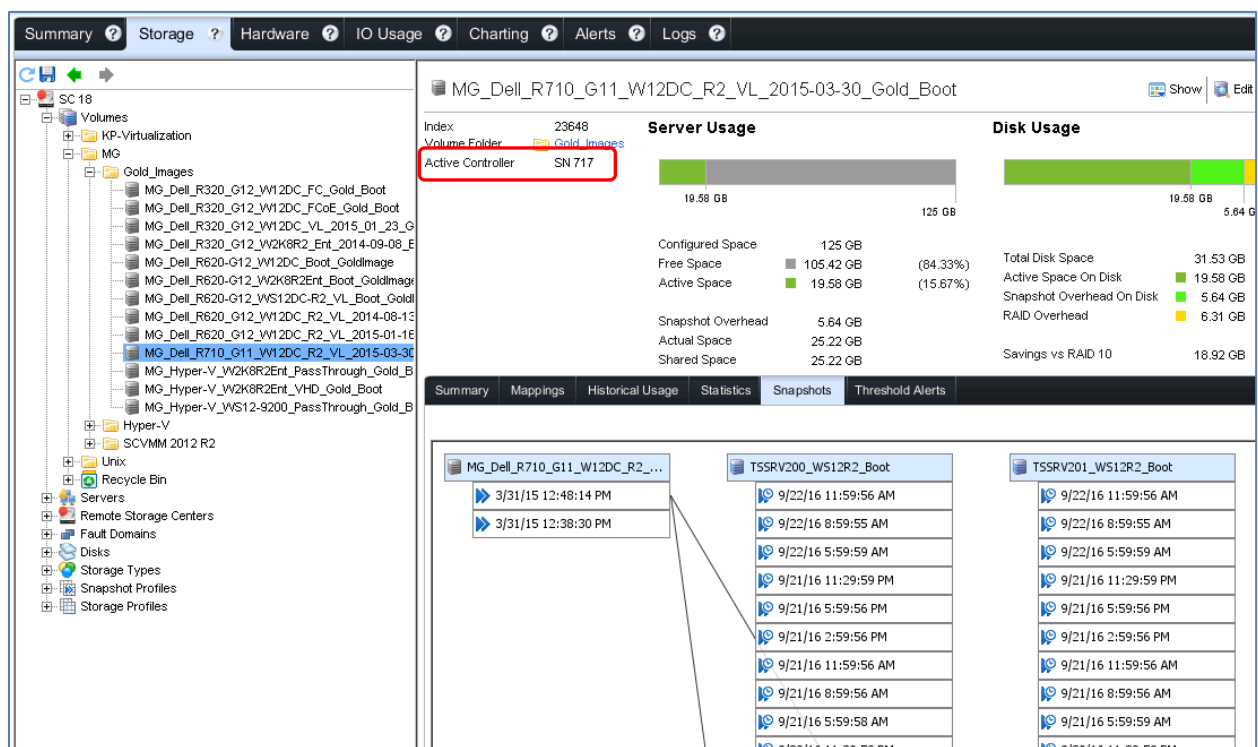


Figure 11 Verify controller ownership for a gold image and associated View Volumes

If a gold image will be used to deploy a large number of new host servers, avoid causing a controller imbalance by creating an additional gold image source volume and manually assigning it to the other controller head in the pair, in this example, SN 716. All new hosts provisioned from it will then be owned by SN 716.

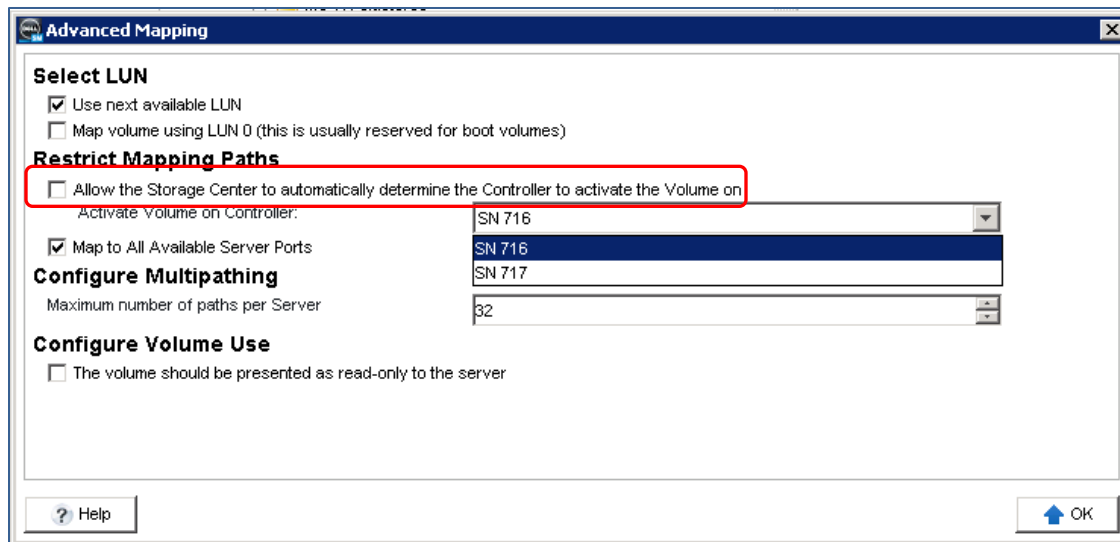


Figure 12 Specify controller ownership for a new gold image volume

As a best practice, add the controller SN to the name of the gold image volume.

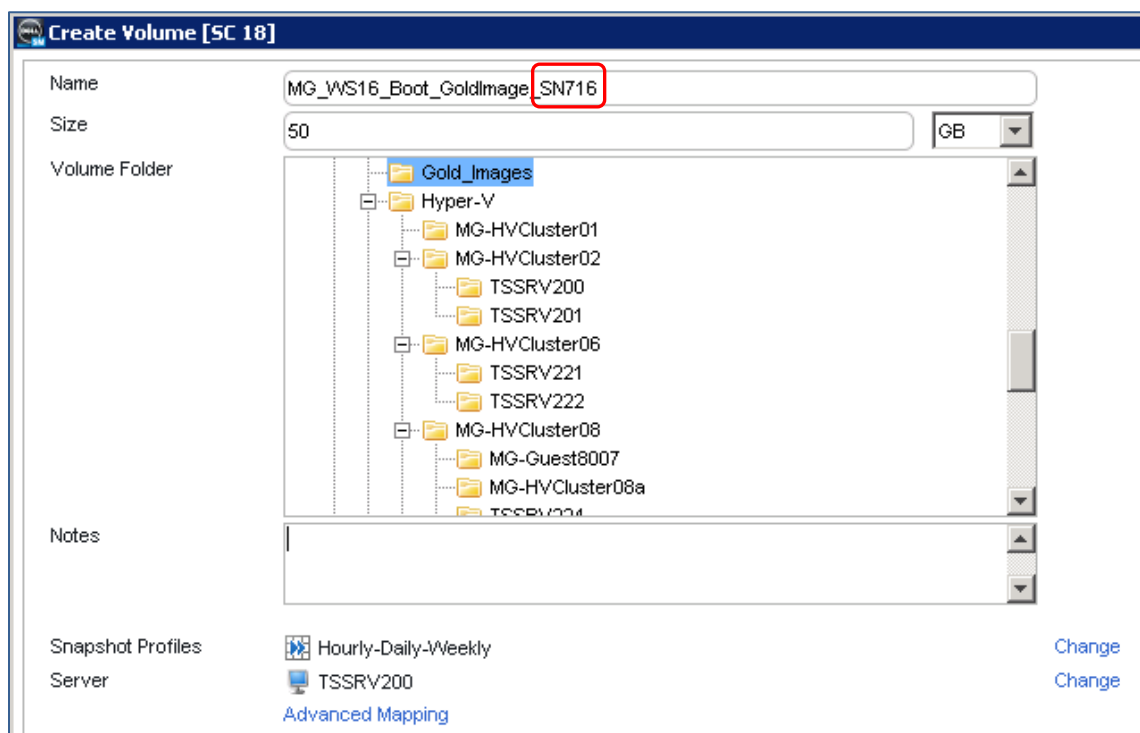


Figure 13 Add controller SN to the volume name

As new servers are deployed from View Volumes, alternate between gold image source volumes so that each controller head stays balanced by owning roughly the same number of volumes.

4.5 SC Series snapshots and Hyper-V VM migration

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 because all nodes in that cluster have shared access to the underlying cluster shared volumes (CSV).

However, when an administrator needs to migrate a guest VM from one host or cluster to another host or cluster, the data (the virtual hard disks) must be copied to the target host or cluster, and this will consume network bandwidth and may require significant time if the virtual hard disks are extremely large. This can also consume additional SAN space unnecessarily because another copy of the data is created.

When moving VMs to another host or cluster, it may be much quicker to leverage the SC Series to simply unmap the host volumes containing the VM's configuration and virtual hard disks and map them to the new target host or cluster. This can also be done using a View Volume from a point-in-time snapshot of the volume.

While this might involve a small amount of down time for the VM being moved during a maintenance window, it might be a much more practical approach than waiting for terabytes of data to copy over the network, consuming additional SAN space unnecessarily.

To avoid or minimize down time when multiple SC Series are involved, consider leveraging SC Series replication and Live Volume features. For more Information on replication between SC Series including Live Volume with automatic failover for Microsoft, see the [Synchronous Replication and Live Volume Solutions Guide](#) on Dell TechCenter.

5 Data progression and Hyper-V

Data progression is a core SC Series feature. While this feature is most commonly described in terms of how it can optimize SC Series arrays that are comprised of multiple disk tiers, it will also optimize data placement on SC Series arrays comprised of a single disk tier. With data progression, data is automatically and intelligently placed in the optimal storage tier and RAID level based on usage and performance metrics.

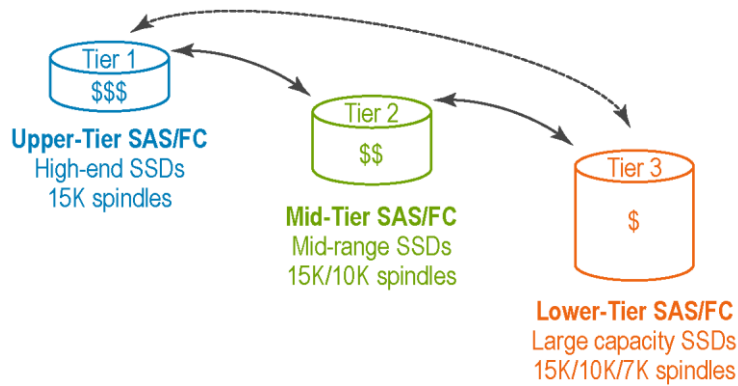


Figure 14 SC Series with data progression

The highest tier in multiple-tier SC Series arrays is typically comprised of high performance, smaller capacity, and more expensive disks. The lower tiers are typically comprised of slower, larger capacity, less expensive disks. SC Series arrays now support all-flash, a mix of spinning and flash, and all spinning configurations, depending on the performance and capacity needs of the workload.

In most environments, about 80% of all data is less active or archival in nature and therefore, providing a lower tier comprised of large capacity media (spinning or SSD) for data progression to use is an important part of the array design.

New data is automatically written to the highest storage tier for maximum performance (Tier 1-RAID 10). Data progression will gradually move inactive data to the lowest tier. Conversely, if data in Tier 3 begins to experience frequent activity, data progression will automatically move it to a higher performing tier.

5.1 Tuning data progression settings for Hyper-V

Since data progression is platform agnostic, there are no extra steps required in order for Hyper-V to take full advantage of data progression and data tiering. Choosing the **Recommended (All Tiers)** storage profile when creating new volumes (including cluster shared volumes) generally works well and is recommended for Hyper-V environments.

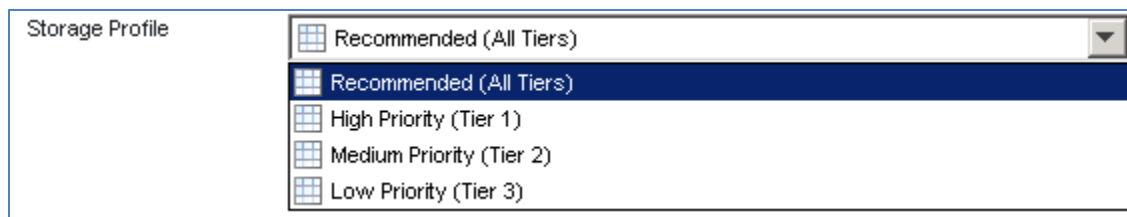


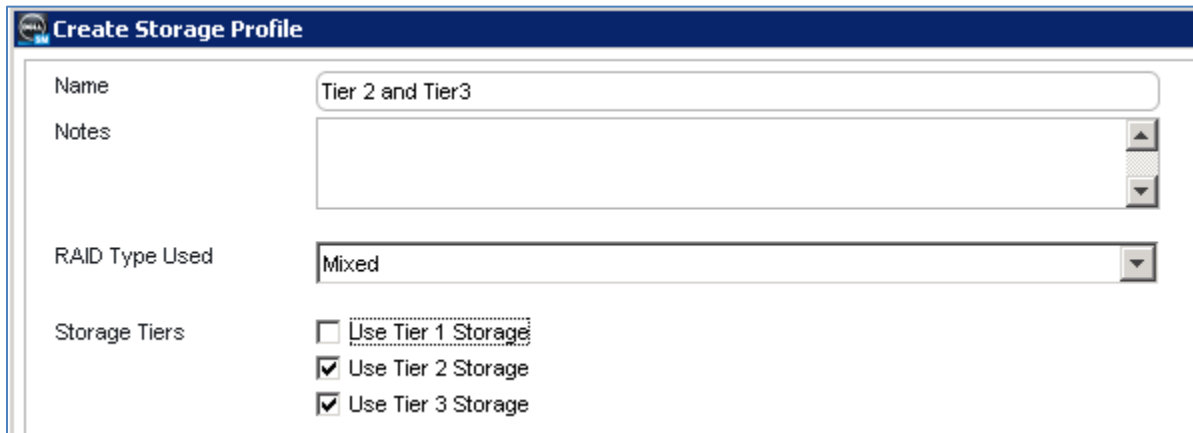
Figure 15 Default storage profiles

Choosing a different storage profile for a volume might be advisable for some Hyper-V configurations. Consider the examples in sections 5.1.1 and 5.1.2.

5.1.1 Data Progression with archival data

In this example, a Hyper-V VM workload creates a large amount of archival data, such as image files or video. This data is stored on separate virtual hard disks on one or more dedicated CSVs, and once this data has been written to disk, it is infrequently accessed.

- Option 1: Leverage the **Recommended (All Tiers)** storage profile. New writes will go to Tier 1, and over time (about 12 days between each tier) data progression will move the data to Tier 2 (if a Tier 2 exists) and then to Tier 3. If writing a large amount of archival data to Tier 1 does not negatively impact the performance or capacity of Tier 1 that might be needed for other workloads, then the recommended storage profile would work well.
- Option 2: Configure the CSV to use only the **Low Priority (Tier 3)** storage profile. This will ensure that all new data to the CSV is written to Tier 3 from the start. This is helpful when the performance of Tier 1 is needed for other workloads, or has a limited capacity that would be negatively impacted by ingesting a lot of new data that is essentially archival. With this design, the performance of Tier 3 would need to be sufficient to ingest the data without negatively impacting the performance of the application that is generating the data.
- Option 3: Create a custom storage profile that includes Tier 2 and Tier 3 only. This assumes the array has three tiers of storage. If the performance of Tier 3 is inadequate for ingesting the new data, this ensures that Tier 2 receives the data. The data is still kept out of Tier 1, and data progression will move the data to from Tier 2 to Tier 3 in about 12 days.



The screenshot shows a 'Create Storage Profile' window. It has a title bar with a Dell logo and the text 'Create Storage Profile'. Inside, there are four main sections: 'Name' with a text box containing 'Tier 2 and Tier3'; 'Notes' with a large empty text area and scrollbars; 'RAID Type Used' with a dropdown menu showing 'Mixed'; and 'Storage Tiers' with three checkboxes. The 'Use Tier 1 Storage' checkbox is unchecked, while 'Use Tier 2 Storage' and 'Use Tier 3 Storage' are both checked.

Figure 16 Create a custom storage profile

This strategy can also be applied to workload elements that require the maximum performance of a higher tier. Dedicate a CSV to the virtual hard disks that host these elements, and select the **High Priority (Tier 1)** storage profile to keep the data in Tier 1, or create a custom profile that allows Tier 1 or Tier 2, but not Tier 3. This might also apply to volumes containing gold images from which many VMs or hosts will be provisioned.

5.1.2 Data copies and migrations

When copying or moving data from one location to another, it is possible to inadvertently consume all of the available capacity in Tier 1. This is because by default, new data is written using RAID 10 in the highest Tier, and Tier 1 is often comprised of a smaller number of lower capacity, high performance disks.

Filling up Tier 1 is undesirable because new writes are then forced to occur in a lower tier which can result in significantly degraded performance for the SC Series storage and any hosted workloads. Normally, alert thresholds would notify administrators with enough lead time to remedy a Tier 1 capacity issue (by adding disks for example), but large data copy or migration operation might consume Tier 1 capacity before an administrator has time to respond.

SC Series arrays do provide some protections against this scenario. For example, when replicating a volume from one SC Series array to another, by default the replicated data is placed into the lowest tier. However, if performing a file level copy in Hyper-V at the host or VM level, new writes will occur in the highest tier allowed by the storage profile that is applied to the target volume, which will usually be the **Recommended (All Tiers)** profile.

To avoid exhausting Tier 1 capacity when copying or migrating a large amount of data, consider the following options:

- Option 1: Modify the storage profile for the target CSV or other volume to temporarily exclude Tier 1. Then copy the data. Then change the storage profile back again. This assumes that any workloads that use the target volume during the copy operation are not negatively impacted by lack of Tier 1 performance.
- Option 2: Copy data in smaller batches. Monitor the capacity status of Tier 1 and allow data progression to digest and move the data to a lower tier to free up space in Tier 1 before copying more data.
- Option 3: Create a new CSV or other volume as a target for the copied or migrated data and assign a storage profile that excludes Tier 1. Once the data is copied, then adjust the storage profile to include Tier 1 if necessary.

6 Disk space recovery with Hyper-V

In Microsoft Windows, deleting a file just removes the file pointer, not the actual data itself. The operating system will report this deleted space as free to be overwritten by new data. However, there are situations where the SC Series will not return the deleted space to the SC Series array page pool to be used elsewhere. This results in inefficient use of the SAN storage capacity.

There are two ways to automatically recover deleted space on the array:

- Native Trim/Unmap support
- Install the SC Series server agent on the host server (works in concert with the daily data progression cycle to recover space)

For more information on Trim/Unmap and the SC Series server agent, see the [Windows Server Best Practices Guide](#) on Dell TechCenter.

6.1 SC Series support for Trim/Unmap with Hyper-V

In order for Hyper-V environments to support TRIM/Unmap natively, the following conditions must be met:

- The Server OS (for hosts or clustered nodes) must be 2012 or newer.
- Physical volumes (including boot-from-SAN disks, cluster shared volumes, direct-attached and pass-through disks) must be basic disks formatted as NTFS volumes. TRIM/Unmap is not supported with other formats such as ReFS.
- SC Series SCOS must be version 6.3.1 or higher.
- Virtual hard disks support TRIM/Unmap so long as:
 - The cluster shared volume (or other data volume) that hosts the virtual hard disk is a basic disk formatted as an NTFS volume.
 - The guest OS is Server 2012 or newer.
 - The guest is a generation 2 VM.
 - The guest's virtual hard disks are basic disks formatted as NTFS volumes (fixed or dynamic).

Note: The Dell Storage Manager (DSM) server agent can be installed on a Server 2012 or newer host, but the disk space recovery feature is disabled by default as native support for Trim/Unmap is assumed.

6.2 Space recovery with 2008 R2 Hyper-V

The DSM server agent can be installed on Server 2008 physical hosts or guest VMs. It enables disk recovery to a limited extent if the following conditions are met:

- Recovery is possible from physical volumes (including boot-from-SAN disks, direct-attached and pass-through disks)
- Disks must be basic disks formatted as NTFS volumes. Other formats are not supported.

Disk space recovery will not work on the following types of drives/volumes with Server 2008 R2 Hyper-V:

- Virtual hard disks (dynamic or fixed). If free space recovery is highly desirable for a particular guest VM scenario, then present storage (basic disk formatted as NTFS) to guest VMs as pass-through or direct-attached disks.
- Space recovery from a pass-through disk is supported only if the following conditions are met:
 - Server Agent is installed on the guest VM
 - The pass-through disk is presented to the guest as a virtual SCSI device (IDE not supported)
 - The pass-through disk is a basic disk formatted as an NTFS volume.
- Even if a cluster shared volume is a basic disk formatted as NTFS, disk space recovery is not supported on cluster shared volumes with Server 2008 R2 Hyper-V hosts.

7 Boot from SAN for Hyper-V

In some cases, such as blade servers that do not have onboard disks, booting from SAN may be the only option available. Many physical Hyper-V hosts have the ability to boot locally or from an SC Series array. Deciding which option to use is dependent on a number of design considerations that are unique to each environment. There are advantages and disadvantages to both options, and in some cases, using local disk is preferred.

Note: Boot from SAN is not supported with Nano Server.

7.1 Boot from SAN advantages

- SC Series snapshots of boot volumes provide a quick recovery.
- Replicate boot volumes to another SC Series array at a remote location for enhanced disaster recovery (DR) protection when both sites use similar hardware for server hosts.
- SC Series gold image boot volumes can be leveraged to quickly provision new Hyper-V host servers.

7.2 When to use boot from local disk

- When critical roles such as a primary or secondary backup domain controller, DNS and DHCP need to remain online during offline SAN maintenance or unplanned outages. However, Live Volume with automatic failover can help ensure that critical roles stay online when there is more than one SC Series array available.
- The Dell Storage Manager Data Collector should ideally be on a server with local boot and data disks so it remains available regardless of the state of the SC Series controllers.

7.2.1 Domain controllers on Hyper-V clusters

It is a best practice to avoid configuring an Active Directory (AD) domain controller as Hyper-V guest VM on a Hyper-V cluster when the cluster service requires AD authentication in order to start. Consider this scenario: a service outage takes the cluster offline (including the domain controller VM). When attempting to recover, unless there is another domain controller available outside of the affected cluster, the cluster service will not start because the domain controller VM it needs to authenticate against is offline.

Note: This order dependency is eliminated with server 2016 Hyper-V because the cluster service uses certificates to authenticate instead of AD. With Server 2016, Hyper-V clusters can be comprised of nodes that are in workgroups or domains.

Encountering this scenario may be service affecting depending on how long it takes to recover. It may be necessary to manually recover the domain controller VM to a standalone Hyper-V host outside of the cluster, or to another cluster. If the domain controller VM exists on an SC Series volume, create a View Volume and map it to another host outside of the affected cluster, and manually recover the domain controller VM there.

This situation can be avoided by doing the following:

- Configure at least one domain controller as a physical server booting from local disk.

- Place virtualized domain controllers on standalone Hyper-V hosts or on individual cluster nodes (which avoids the cluster service dependency on AD), and use Hyper-V Replica (2012 and newer) to ensure that the guest VM can be recovered on another host.
- Place virtualized backup domain controllers on separate clusters, so that a service-affecting event with any one cluster does not result in all domain controllers becoming unavailable. This does not protect against cases where there is a site outage that takes all the clusters (and therefore all the virtualized AD servers) offline.
- Leverage Server 2016 Hyper-V, which does not have an AD dependency to authenticate cluster services.

7.3 Configure Hyper-V hosts to boot from SAN

With the exception of Nano Server, Windows Servers including Hyper-V hosts and nodes can be configured to boot from SAN using physical Fibre Channel or iSCSI host bus adapters (HBAs). To learn more, see the [*Microsoft Windows Server Best Practices Guide for Dell SC Series*](#) on Dell TechCenter.

8 PowerShell integration

Dell SC Series storage has incorporated Windows PowerShell for many years and supports a growing library of cmdlets. The SC Series PowerShell SDK command set provides administrators with the ability to perform many SC Series tasks from the command line and create automated scripts.

To learn more about PowerShell integration with SC Series stories including many examples, see the [Dell Storage PowerShell SDK Cookbook](#) on Dell TechCenter.

8.1 Importance of PowerShell

With newer versions of Hyper-V, particularly with 2016, the use of PowerShell is required for more and more day-to-day tasks for which there may not be a GUI equivalent. This is particularly true with Server Core and Nano Server. While the real power of PowerShell lies in being able script and automate complex or repeatable processes, it is often more convenient and efficient to complete simple tasks using PowerShell instead of a GUI.

Along with increased focus on Server Core and Nano Server, PowerShell reflects a trend by Microsoft to move away from servers with desktops, minimizing the malware attack surface, saving disk space, and reducing the number of disruptive host reboots to apply critical patches. It also reflects a trend by Microsoft to adopt a strategy of providing “just-enough” server and “just-enough” administrative rights to solve business needs and perform tasks. While it presents a learning curve for administrators who are not familiar with or comfortable with using the command line, PowerShell is the way of the future for server administration and so learning PowerShell is becoming increasingly important.

While initial creation and testing of PowerShell scripts may require significant up-front time in some cases, when the time savings is significant, it can be very much worth the effort. As scripts are built, they can be saved for future use, and used as building blocks to create additional scripts. Many online resources exist to aid administrators with learning to use PowerShell and develop their own scripts. See for example the [PowerShell Developer Network](#) on the Microsoft MSDN website.

8.2 PowerShell automation with Hyper-V and SC Series

PowerShell for SC Series and PowerShell for Hyper-V can be used together to script processes that involve the host servers, the Hyper-V layer and the SC Series layer. This provides administrators great control and flexibility to automate repeatable tasks and solve complex problems. Scripting also reduces the risk of user error when having to complete repetitive or tedious tasks, ensuring that steps are not missed and completed in the right order, and that naming is consistent.

While many simple and routine administrative tasks can still be accomplished adequately and efficiently through a GUI in most cases, in some situations using GUI-based tools may be inefficient or inadequate. Some examples include:

Example 1: An administrator needs to quickly modify the same setting on 500 guest VMs. Using a GUI to modify each guest VM one at a time would prove to be very impractical and onerous. By leveraging PowerShell, an administrator could script an automated process that would accomplish the same task in a fraction of the time.

Example 2: An administrator needs to provision 100 new guest VMs from an SC Series volume snapshot that contains a gold image. Native SC Series, Hyper-V and SCVMM GUI tools make it difficult to create more than one VM at a time. With PowerShell, the process can be scripted and automated.

Example 3: A backup process that involves SC Series storage and Hyper-V guest VMs runs overnight that requires some manual sequenced steps in order to complete properly. With PowerShell, the process can be automated so that an administrator does not have to be available to manage these steps after hours.

Example 4: As part of a DR plan, a number of manual order-dependent steps that involve the SC Series and recovery hosts at a DR site must be completed to bring a critical workload online when invoking a DR plan. Following the steps manually increases the time required, and increases the risk of user error, particularly if the situation is stressful. By leveraging PowerShell, many of these manual steps could be automated making it easier to meet recovery time objectives (RTO).

Because PowerShell is an incredibly powerful tool, it can be incredibly dangerous to use if administrators do not know exactly what they are doing. GUI interfaces for the most part keep administrators safe by guiding their steps with wizards that provide warning messages to keep them out of trouble. However, the trade-off is that GUI functionality is often limited. PowerShell on the other hand allows administrators great control of their environment – often beyond what can be done in a GUI. But with that power comes the inherent risk of undesired and unintended consequences if mistakes are made. PowerShell will not always provide warning messages or prevent destructive commands or scripts from running.

The customer is strongly advised to test PowerShell scripts in a non-production environment, and to use extreme caution when configuring scripts that involve SC Series cmdlets to avoid unintentional data loss.

A number of basic script examples specific to SC Series storage are provided in the [Dell Storage PowerShell SDK Cookbook](#) on Dell TechCenter as mentioned above.

9 Business continuity

A good business continuity strategy will always incorporate disaster recovery and disaster avoidance planning. At a high level, a disaster recovery plan is a process where a company ensures they are able to recover as quickly as possible from data loss or from an interruption or failure that prevents access to data. It is a very important part of overall IT strategy and in some cases is governed by regulations specific to particular industries.

The disaster recovery scenarios that may be encountered are diverse, but may vary by location. Disasters can be small or large. The loss of a single document that impacts one user is a disaster for that user. A site failure might impact many users and jeopardize the future of the business if not resolved quickly. For the most part, the essential elements of disaster recovery are now commonplace, reliable, fairly inexpensive, easy to implement, and are able to address or prevent the majority of events that are most likely to occur. These protections and safeguards might include moving key workloads to a cloud provider, tape backups with off-site storage, on-line backups, disk-to-disk backups, network and physical security measures, malware protection, redundant hardware and internet connections, SAN-based snapshots with remote replication, and battery backups or generators.

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

9.1 Cost versus risk analysis

The most robust disaster recovery solutions might also be cost prohibitive. Each customer must weigh the costs versus the risks and determine what level of DR protection is necessary for them. Questions that might be asked as part of a cost versus risk analysis include:

- What regulations apply to my industry?
- What are the terms of any service level agreements (SLAs) for business continuity that must be honored?
- What applications and data are the most mission critical to the business or our customers?
- What is the Recovery Time Objective (RTO) for each application or service? How long can something be down before the business impact becomes too great? For example:
 - Practice Management System: 30 minutes
 - Messaging system: 4 hours
 - Research and development Server: 2 days
- What is the Recovery Point Objective (RPO)? How much data loss is acceptable for a particular subset of data? Is backing up the email server once a day adequate? If so, the mail server has an RPO of 24 hours, meaning that up to 24 hours of mail data may be lost in the event the mail server has to be recovered from the last backup.
- What types of events are most likely to occur, factoring in the geographic location? A coastal location may be subject to hurricanes. A location on a fault line may be subject to earthquakes. A location in a low lying area may be prone to flooding.
- Is an alternate site far enough away so that the same event does not impact both locations?

- How much will it cost (hardware, software, and staff) to design, implement and support the desired protections? Is that cost justified given the risk?

9.2 Disaster recovery and disaster avoidance

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

- Events that cause data loss such as malware infection, corruption, accidental deletion, sabotage or hardware failure of disks or disk arrays.
- Events that interrupt the ability to access data within or between sites, such as a network or power failure (but no data is lost).
- Events that cause both loss of data and loss of access to a site, typically caused by more significant and destructive events such as a fire or natural disaster.

Disaster avoidance implies having enough lead time to proactively react to an impending event, such as an approaching hurricane, in a way that avoids or minimizes down time. This is the strategy commonly used when doing system maintenance. An administrator may move a critical workload to an alternate location using Live Volume with Automatic Failover before site maintenance at the main location causes an outage there.

A good business continuity plan will include both disaster recovery and disaster avoidance strategies that leverage a combination of manual and automatic processes to address a wide range of possible scenarios.

- A manual process might be required to restore lost data from a backup or snapshot, or to bring a Hyper-V guest VM on line at an alternate site. PowerShell might be used to automate manual steps.
- An automatic process kicks in on its own, such as when Live Volume with automatic failover moves the primary Live Volumes to a secondary SC Series.

9.3 Live Volume with Auto Failover for Microsoft

Live Volume is an optional feature that has been available with SC Series storage for many years, Auto Failover is a recent enhancement to Live Volume that first became available with the Storage Center Operating System (SCOS) version 6.7 and the feature Live Volume with Auto Failover (LV-AFO) was initially supported in VMware environments only. With the release of SCOS version 7.1, LV-AFO support was extended to Microsoft environments in addition to VMware.

To learn more about how LV-AFO can be configured to protect Microsoft clusters and Hyper-V (including stretched clusters), see the 3-part demo video names and the [Synchronous Replication and Live Volume Solutions Guide](#) and a 3-part demo video [Live Volume with Auto Failover Support for Microsoft](#) on Dell TechCenter.

9.3.1 Predefined DR plans

Working closely in conjunction with Live Volume, Dell Storage Manager can be used to create predefined DR plans for Hyper-V environments. More details about predefined DR plans can be found in the references in the preceding section of this document. In addition, see the [DR Plan for Hyper-V](#) demo video on Dell TechCenter.

9.4 Replay Manager for Hyper-V

Replay Manager 7 for Dell SC Series is a GUI-based data protection application that includes extensions to support the creation of application-consistent backups for:

- Guest VMs in Hyper-V environments or in VMware environments
- VMware datastores
- Exchange or SQL Server data
- Local volumes on hosts and guests, when the Replay Manager agent is installed locally

Replay Manager also includes a number of PowerShell cmdlets to allow administrators to perform Replay Manager operations that might otherwise require an administrator to be available after hours.

In Microsoft environments, Replay Manager 7 leverages the power of the Microsoft Volume Shadow Copy Service (VSS) to create and manage application-consistent snapshots of the protected hosts or guest VMs and the workload. With application consistency, the host OS or guest VM OS, along with the workload, are gracefully paused before snapshots are taken. This is especially important when protecting a transactional workload such as a database, to help ensure recovery without errors or data corruption.

For more information about Replay Manager 7 and Hyper-V see the [Replay Manager 7 and Microsoft Hyper-V Best Practices Guide](#) and the [Replay Manager 7 for Hyper-V Demo Video](#) on Dell TechCenter.

9.5 Azure Site Recovery (ASR) with SC Series

Azure is the Microsoft cloud-based computing and infrastructure platform. Azure Site Recovery provides administrators with the ability to include synchronous or asynchronous SAN volume replication as part of a fully orchestrated recovery plan to protect the VMs in an on-premises private cloud. Incorporating SAN volume replication is made possible through Dell SMI-S provider integration with System Center Virtual Machine Manager (SCVMM).

The Dell SMI-S provider is included at no additional cost with Dell SC Series. When SMI-S is enabled and configured, SCVMM can manage Dell SC Series controllers directly. In addition, with SMI-S integration, administrators can create SAN copy-capable templates to rapidly provision new guest VMs. For more information about SC Series and SMI-S integration with SCVMM, see the [Dell SC Series and Microsoft SCVMM Best Practices Guide](#) and the [Rapid Provisioning Demo Video](#) on Dell TechCenter.

In the case of ASR, Azure leverages SMI-S (through SCVMM) to incorporate asynchronous or synchronous SAN replication and failover as part of an Azure-based recovery plan.

For more information about how to configure ASR with Dell SC Series, see the [Dell SC Storage with Microsoft Azure Site Recovery Configuration Guide](#) and the [Azure Site Recovery Demo Video](#) on Dell TechCenter.

10 Conclusion

This guide provides administrators with information and best practices for configuring Hyper-V environments to work optimally with SC Series storage.

If you have suggestions or questions about this document, send your feedback to StorageSolutionsFeedback@dell.com.

A Technical support and additional resources

Dell.com/support is focused on meeting customer needs with proven services and support.

Dell TechCenter is an online technical community where IT professionals have access to numerous resources for Dell software, hardware and services.

Storage Solutions Technical Documents on Dell TechCenter provide expertise that helps to ensure customer success on Dell Storage platforms.

A.1 Supplemental documentation

Referenced or recommended Dell publications:

- Dell SC Series Storage Solutions
<http://www.dell.com/us/business/p/dell-compellent>
- Dell SCv2000 and SC4020 Arrays with SAS Front-end support for Microsoft
http://en.community.dell.com/techcenter/extras/m/white_papers/20441716
- Dell Storage Center Microsoft Multipath I/O Best Practices Guide
http://en.community.dell.com/techcenter/extras/m/white_papers/20437917
- Dell SC Series Virtual Fibre Channel for Hyper-V Demo Video
http://en.community.dell.com/techcenter/extras/m/white_papers/20440350
- Dell SC Series Replay Manager 7 and Hyper-V Best Practices Guide
http://en.community.dell.com/techcenter/extras/m/white_papers/20437951
- Dell SC Series Replay Manager 7 and Hyper-V Demo Video
http://en.community.dell.com/techcenter/extras/m/white_papers/20440352
- Microsoft Windows Server Best Practices Guide for Dell SC Series
http://en.community.dell.com/techcenter/extras/m/white_papers/20437924
- Dell SCOS 7.0 Data Reduction with Deduplication and Compression Guide
http://en.community.dell.com/techcenter/extras/m/white_papers/20442456
- Dell SC Series Synchronous Replication and Live Volume Solutions Guide
http://en.community.dell.com/techcenter/extras/m/white_papers/20439310
- Dell SC Series Live Volume with Auto Failover Support for Microsoft Demo Video
http://en.community.dell.com/techcenter/extras/m/white_papers/20442845
- Dell SC Series DR Plan for Hyper-V Demo Video
http://en.community.dell.com/techcenter/extras/m/white_papers/20440350

- Dell SC Series PowerShell SDK Cookbook
http://en.community.dell.com/techcenter/extras/m/white_papers/20441893
- Dell SC Series and Microsoft SCVMM Best Practices Guide
http://en.community.dell.com/techcenter/extras/m/white_papers/20438128
- Dell SC Series and Microsoft SCVMM Rapid Provisioning Demo Video
http://en.community.dell.com/techcenter/extras/m/white_papers/20440353

Referenced or recommended Microsoft publications:

- Microsoft TechNet Library
<https://technet.microsoft.com/library/default.aspx>
- Microsoft PowerShell Developer Network (MSDN)
<https://msdn.microsoft.com/powershell>