

# Dell EMC SC Series: Microsoft Hyper-V Best Practices

## Abstract

This document provides best practices for configuring Microsoft® Hyper-V® to perform optimally with Dell EMC™ SC Series storage.

June 2019

## Revisions

Date	Description
June 2009	Initial release
August 2009	Updated for Windows Server 2008 R2
September 2011	Updated for Windows Server 2008 R2 SP1 and SCOS 5.5
October 2012	Updated for Windows Server 2012, Enterprise Manager 6.2, and SCOS 6.2
December 2012	Added support for virtual Fibre Channel with SCOS 6.3 and Enterprise Manager 6.3
October 2013	Updated for Windows Server 2012 R2; applied new document template
October 2016	Updated for Windows Server 2016
June 2019	Updated for Windows Server 2019; applied new document template

## Acknowledgements

Author: Marty Glaser

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

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

Copyright © 2009–2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [6/10/2019] [Best Practices] [CML1009]

# Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents .....	3
Executive summary.....	5
<b>1 Introduction.....</b>	<b>6</b>
1.1 SC Series.....	6
1.2 Microsoft Server Hyper-V .....	6
1.3 Supported versions.....	7
1.4 Best practices overview.....	8
1.5 General best practices for Hyper-V .....	8
<b>2 Optimize Hyper-V for SC Series.....</b>	<b>9</b>
2.1 Hyper-V integration services .....	9
2.2 Hyper-V guest VM generations .....	11
2.2.1 Convert VMs to a newer generation .....	12
2.3 Virtual hard disks .....	13
2.3.1 Virtual hard disk format.....	13
2.3.2 Virtual hard disk type .....	14
2.3.3 Virtual hard disks and thin provisioning with SC Series storage .....	16
2.3.4 Overprovisioning with dynamic virtual hard disks.....	17
2.4 Present SC Series storage to Hyper-V.....	17
2.5 Transport options.....	18
2.5.1 SC Series and front-end SAS support for Hyper-V .....	18
2.5.2 Multiple transports .....	18
2.6 MPIO best practices .....	19
2.7 Guest VMs and in-guest iSCSI and virtual Fibre Channel disks .....	19
2.8 Guest VMs and direct attached storage .....	20
2.9 Guest VMs and pass-through disks.....	21
2.10 SC Series arrays and cluster server objects .....	22
2.11 SC Series LUN limits for larger Hyper-V clusters.....	23
2.12 Volume design considerations for SC Series .....	24
2.13 Offloaded data transfer.....	24
2.14 Disable automount.....	25
2.15 Placement of page files .....	25
2.16 Placement of Active Directory domain controllers.....	26

2.17	SC Series data reduction and Hyper-V .....	26
<b>3</b>	<b>SC Series snapshots and Hyper-V .....</b>	<b>27</b>
3.1	SC Series Replay Manager support for Hyper-V .....	27
3.2	Use SC Series snapshots to recover guest VMs .....	27
3.2.1	Recover a guest VM on a standalone Hyper-V host .....	28
3.2.2	Recover a guest VM on a cluster shared volume.....	29
3.3	Change a cluster shared volume disk ID with Diskpart .....	29
3.4	Use SC Series snapshots to create a test environment.....	31
3.5	Leverage SC Series to create gold images .....	32
3.5.1	Gold images and preserving balanced SC Series controllers .....	34
3.6	SC Series snapshots and Hyper-V VM migration .....	36
<b>4</b>	<b>Data Progression and Hyper-V .....</b>	<b>37</b>
4.1	Tuning Data Progression settings for Hyper-V.....	37
4.1.1	Data Progression with archival data .....	38
4.1.2	Data copies and migrations .....	38
<b>5</b>	<b>Disk space recovery with Hyper-V .....</b>	<b>40</b>
5.1	SC Series support for Trim/Unmap with Hyper-V.....	40
5.2	Space recovery with 2008 R2 Hyper-V .....	40
<b>6</b>	<b>Boot-from-SAN for Hyper-V.....</b>	<b>42</b>
6.1	Configure Hyper-V hosts to boot-from-SAN .....	42
<b>7</b>	<b>PowerShell integration .....</b>	<b>43</b>
7.1	Importance of PowerShell .....	43
7.2	PowerShell automation with Hyper-V and SC Series.....	43
7.3	Best practices for PowerShell.....	44
<b>8</b>	<b>Business continuity with Hyper-V and SC Series.....</b>	<b>45</b>
8.1	Cost/risk analysis.....	45
8.2	Disaster recovery and disaster avoidance .....	46
8.3	Live Volume with Auto Failover for Microsoft .....	46
8.4	Replay Manager for Hyper-V.....	46
<b>A</b>	<b>Technical support and resources .....</b>	<b>48</b>
A.1	Related resources.....	48

## Executive summary

Dell EMC™ SC Series storage provides a powerful and complete set of storage integrations, and management and monitoring tools for Microsoft® Windows Server® environments. This document provides best practice guidance for deploying and optimizing the Windows Server Hyper-V® role with SC Series arrays.

The documentation at [Dell.com/support](https://Dell.com/support) for specific SC Series arrays and SCOS versions serves as the primary reference material for optimal configuration of SC Series for Windows Server and Hyper-V. These resources include deployment guides, owner's manuals, administrator's guides, installation guides, and release notes.

For SC Series best practices for Windows Server that are not specific to the Hyper-V role, see these best practices guides: [Dell EMC SC Series and Microsoft Windows Server](#) and [Dell EMC SC Series and Microsoft MPIO](#).

See appendix A for additional resources including demo videos and reference architecture white papers in support of application workloads running on SC Series storage and Hyper-V.

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

# 1 Introduction

Microsoft Windows Server Hyper-V and SC Series storage are feature-rich solutions that together present a diverse range of configuration options to solve key business objectives such as storage capacity, workload optimization, performance, and resiliency.

## 1.1 SC Series

The SC Series family includes midrange storage appliances with many robust features including 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 storage is designed with redundancies to avoid downtime for events such as component failures, maintenance, upgrades, and expansion. SC Series arrays provide an efficient scalable platform for the ultimate experience in performance, adaptability, and efficiency.



Figure 1 Front and rear view of the SC7020F all-flash array

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. SC Series storage is well suited to provide a solid, proven storage solution for Hyper-V environments to meet all these business needs.

SC Series arrays support storage area network (SAN) configurations when equipped with Fibre Channel (FC) or iSCSI front-end ports. SC Series arrays also support direct-attached storage (DAS) configurations when select models are equipped from the factory with SAS front-end ports. For more information about SC Series DAS configuration support for Hyper-V, see the guide: [Dell EMC SC Series Storage with SAS Front-end Support for Microsoft Hyper-V](#).

To learn more about specific SC Series arrays and features, visit the [Dell EMC SC Series storage solutions](#) website.

## 1.2 Microsoft Server Hyper-V

The Windows Server platform leverages the Hyper-V role to provide virtualization technology. Hyper-V is one of many optional roles offered with Windows Server and is installed using the **Add Roles and Features** wizard in Server Manager or with PowerShell®. The Hyper-V role is not installed by default.

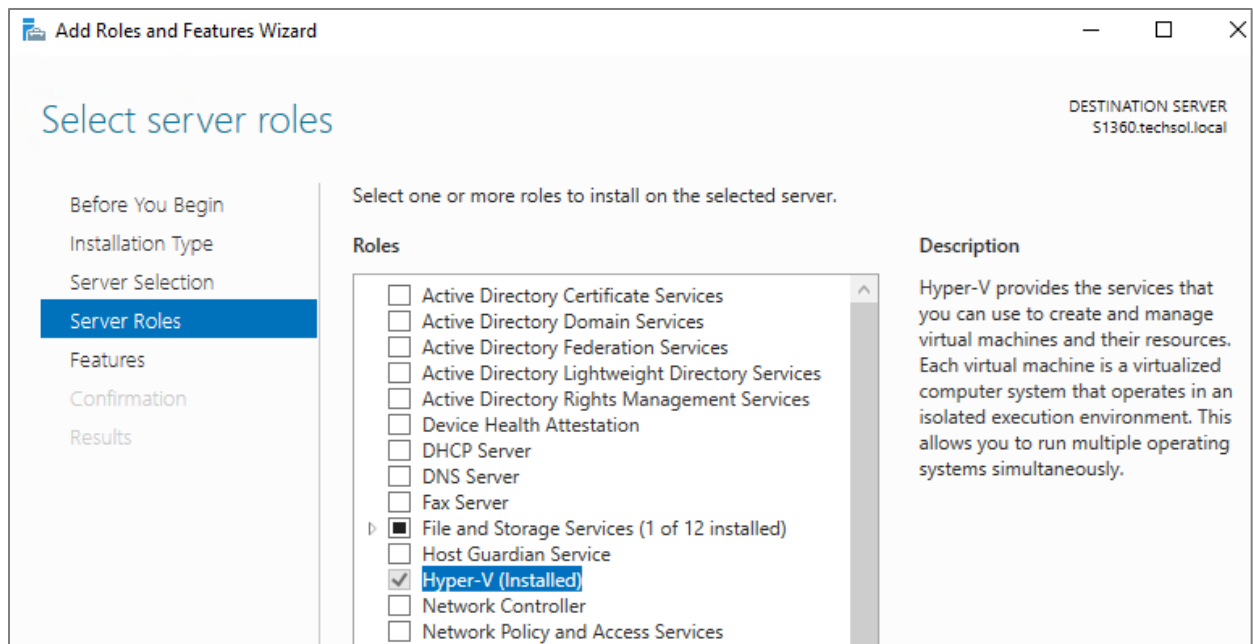


Figure 2 Install the Hyper-V role with the Add Roles and Features Wizard

Initially offered with Windows Server 2008, Hyper-V has matured with each release to include many new features and enhancements. It has evolved to become a mature, robust, proven virtualization platform. In simplest terms, it is a layer of software that presents physical host server hardware resources in an optimized and virtualized manner to guest virtual machines (VMs) and their workloads. Hyper-V hosts (referred to as nodes when clustered) greatly enhance utilization of physical hardware (such as 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, Microsoft System Center Virtual Machine Manager (SCVMM), Windows Admin Center (WAC), and PowerShell, offer administrators great control and flexibility for managing host and VM resources.

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

## 1.3 Supported versions

SC Series storage has supported Hyper-V since the release of Windows Server 2008. Depending on your version of SCOS, Windows Server 2008 R2 Hyper-V and newer is supported. With the release of SCOS 7.4, support for Hyper-V with SC Series is extended to Windows Server 2019 and Windows Server 2019 Hyper-V.

---

**Note:** Not all versions of Windows Server Hyper-V are supported with all SCOS versions. Consult your SCOS documentation to confirm version support.

---



---

**Note:** Microsoft has announced that extended support for Windows Server 2008 R2 will end in January 2020. Customers running Windows Server 2008 R2 Hyper-V should plan to migrate to a newer version before Microsoft extended support ends.

---

## 1.4 Best practices overview

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

As a result, tuning is often unnecessary (and therefore discouraged) unless a specific design, situation, or workload is known to benefit from a different configuration. One of the purposes of a best-practices document is to call attention to situations where default settings or configurations may not be optimal.

Some common best practice objectives include the following:

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

It is important to remember that best practices are baselines that may not be ideal for every environment. Some notable exceptions include the following:

- Legacy systems that are performing well and have not reached their life expectancy may not adhere to current best practices. Dell EMC recommends upgrading to the latest technologies and adopting current best practices at key opportunities such as upgrading or replacing infrastructure.
- A test or development environment that is not business critical may use a less-resilient design or lower-tier hardware to reduce cost and complexity.

---

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

---

## 1.5 General best practices for Hyper-V

There are many general best practices for Hyper-V not specific to storage that are not discussed in detail in this document. See resources such as the [Microsoft Documentation Library](#) for guidance on general Hyper-V best practices.

Common best practices tuning steps for Hyper-V include the following:

- Minimize or disable unnecessary hardware devices and services to free up host CPU cycles that can be used by other VMs (this also helps to reduce power consumption).
- Schedule tasks such as periodic maintenance, backups, malware scans, and patching to run after hours, and stagger start times when such operations overlap and are CPU or I/O intensive.
- Tune application workloads to reduce or eliminate unnecessary processes or activity.
- Leverage Microsoft PowerShell or other scripting tools to automate step-intensive repeatable tasks to ensure consistency and avoid mistakes due to human error. This can also reduce administration time.



## 2 Optimize Hyper-V for SC Series

SC Series storage is an excellent choice for external storage for stand-alone or clustered Windows Servers including servers configured with the Hyper-V role. Core SC Series features such as thin provisioning, Data Progression, data reduction, snapshots (Replays), and replication work seamlessly in the background regardless of the platform or OS. In most cases, the default settings for these features are optimal for Windows Server and Hyper-V. This document will point out additional configuration or tuning steps needed to enhance performance, utilization or up-time.

### 2.1 Hyper-V integration services

Guest integration services are a package of virtualization-aware drivers that are installed on a guest VM to optimize the guest VM virtual hardware for interaction with the physical host hardware and storage. Installing these drivers is typically the first step for optimizing VM performance. If a VM is not performing as expected (due to CPU, disk I/O, or network performance), verify that the VM integration services are current.

Installing and updating integration services is a commonly overlooked step to ensure overall stability and optimal performance of guest VMs. Although newer Windows-based OSs and some enterprise-class Linux-based OSs come with integration services out of the box, updates may still be required. New versions of integration services may become available as the physical Hyper-V hosts are patched and updated.

With earlier versions of Hyper-V (2012 R2 and prior), during the configuration and deployment of a new VM, the configuration process does not prompt the user to install or update integration services. In addition, the process to install integration services with older versions of Hyper-V (2012 R2 and prior) is a bit obscure and is explained in this section. With Windows Server 2016 and Windows Server 2019 Hyper-V, integration services are updated automatically (in the case of Windows VMs) as a part of Windows Updates, requiring less administration time to ensure Windows VMs stay current.

One common issue occurs when VMs are migrated from an older physical host or cluster to a newer host or cluster (for example, from Windows Server 2008 R2 Hyper-V to Windows Server 2012/R2 Hyper-V). The integration services do not get updated automatically, and degraded performance may be encountered as a result. This may erroneously lead an administrator to suspect the storage array as the cause of the problem.

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 for the VM as shown in Figure 3.

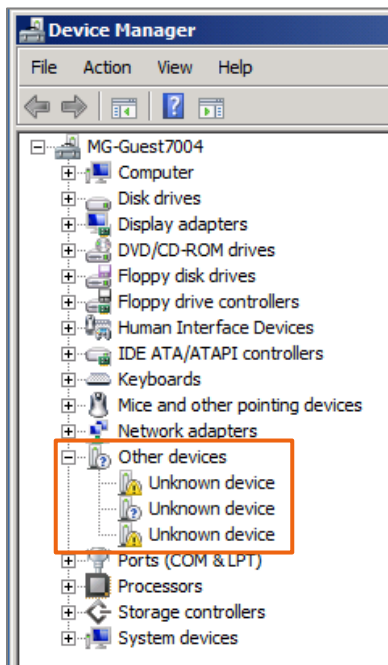


Figure 3 Unknown devices listed for a guest VM indicates missing or outdated integration services

For versions of Hyper-V prior to 2016, use Hyper-V Manager to connect to a VM. Under the **Action** menu, mount the **Integration Services Setup Disk** (an ISO file) as shown in Figure 4, and follow the prompts in the guest VM console to complete the installation.

---

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

---

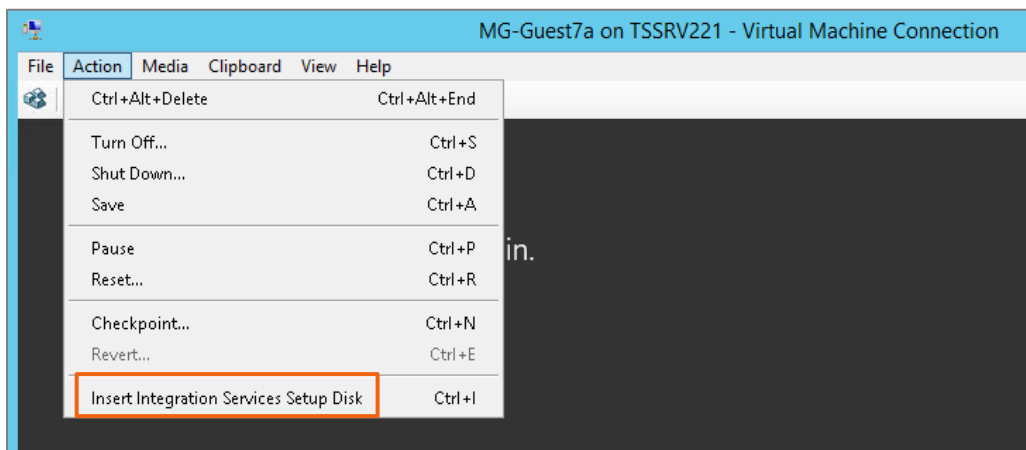


Figure 4 Mount Integration Services Setup Disk

To verify the version of integration services for a VM, click the **Summary** tab in **Failover Cluster Manager**.

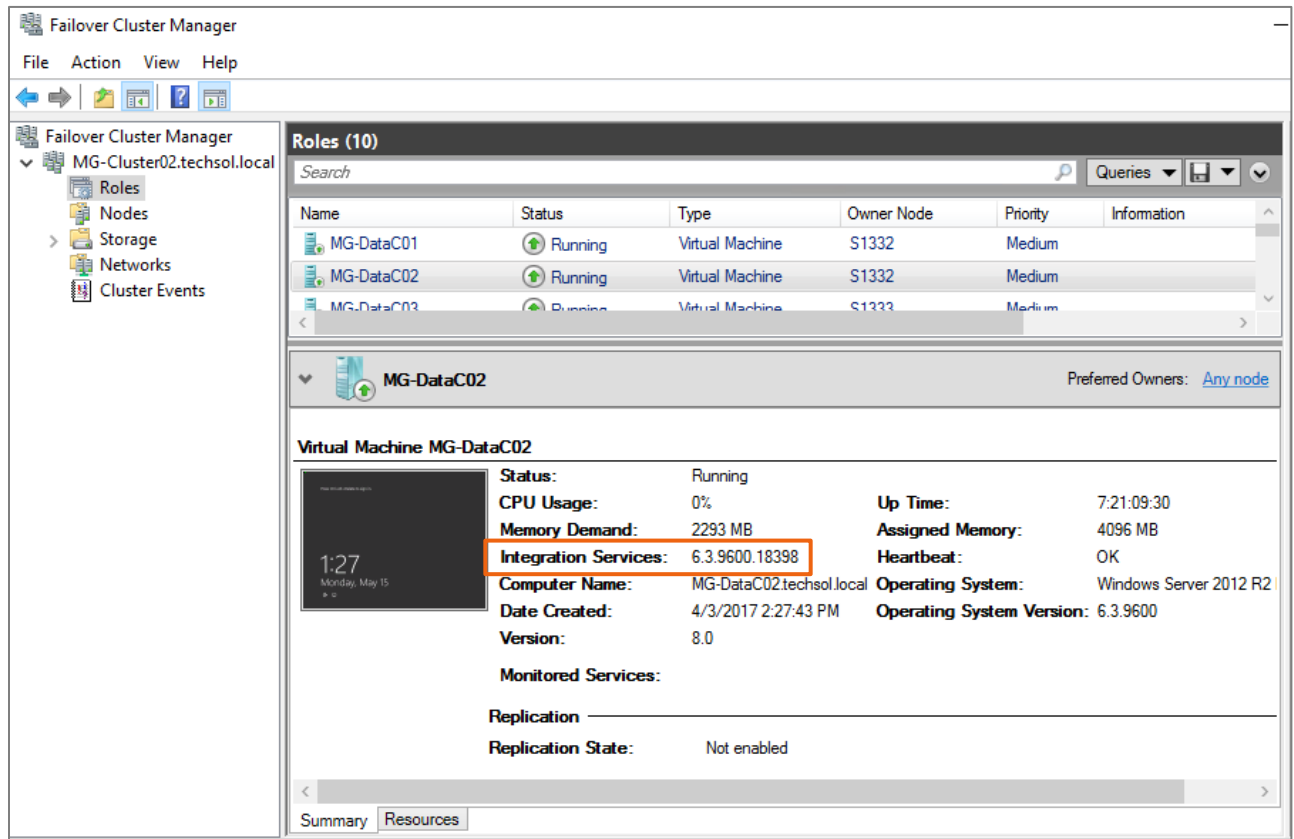


Figure 5 Verify integration services version with Failover Cluster Manager

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

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

## 2.2 Hyper-V guest VM generations

When Windows Server 2012 R2 Hyper-V was released, Microsoft designated all existing VMs as generation 1 to differentiate them from a new classification of VMs that could be created as generation 2. From the perspective of SC Series storage, either generation of VM is supported. The best practices recommendation from Microsoft and Dell EMC is to configure new guests as generation 2 if the workload supports it.

Generation 2 guests use Unified Extensible Firmware Interface (UEFI) when booting instead of a legacy BIOS. UEFI provides better security and better interoperability between the 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 the dependency on virtual IDE for the boot disk. Generation 1 VMs

require the boot disk to use a virtual IDE disk controller. Generation 2 guests instead use virtual SCSI controllers for all disks. Virtual IDE is not a supported option with generation 2 VMs.

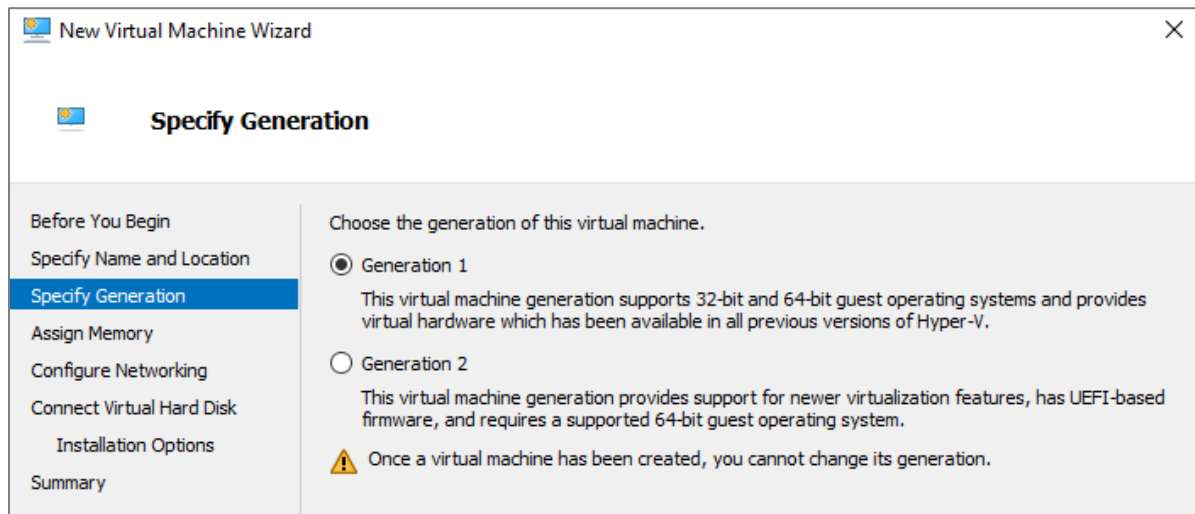


Figure 6 Specify a guest as generation 1 or generation 2

### 2.2.1 Convert VMs to a newer generation

The warning message in the wizard in Figure 6 indicates that the VM generation cannot be changed once a VM has been created. However, third-party tools are available to convert VMs. Dell EMC does not endorse any specific methods for converting VMs (use at your own risk).

---

**Tip:** Leverage SC Series snapshots to create a test environment for a production VM workload where VM conversion can be attempted without affecting the production environment. See section 3 for more information on snapshots.

---

The recommended method is to create a new generation 2 VM; migrate roles, features, workloads, and data; and retire the generation 1 VM.

## 2.3 Virtual hard disks

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

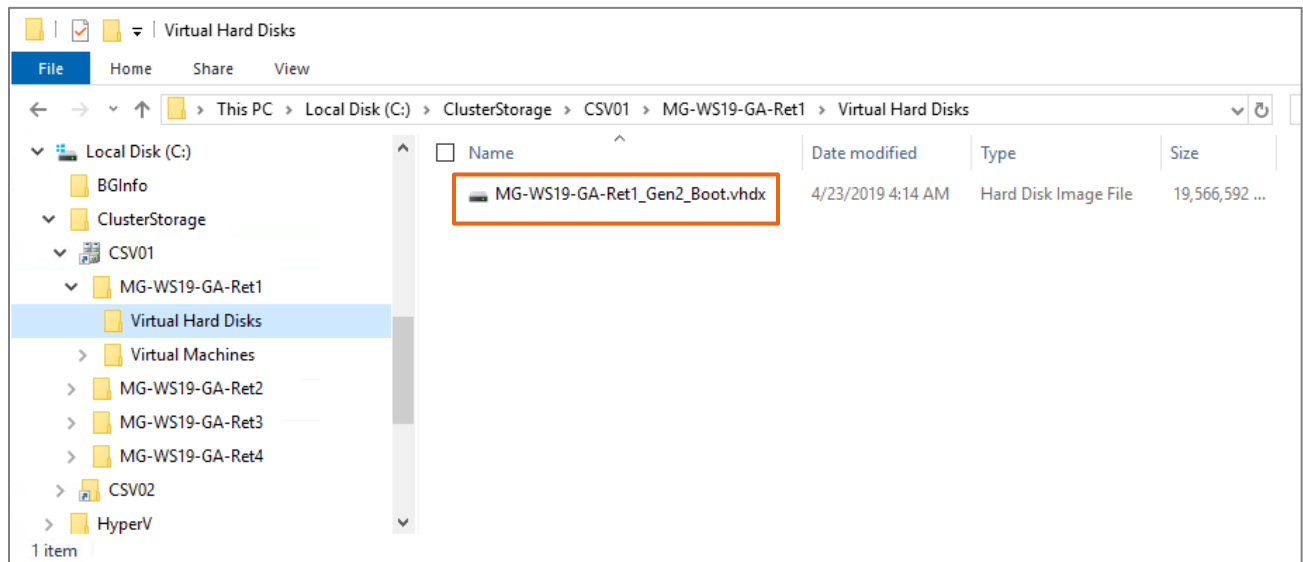


Figure 7 Virtual hard disk file (vhdx) on a cluster shared volume

### 2.3.1 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, but is limited to a maximum size of 2,048 GB.
- **VHDX** is supported with Windows Server 2012 Hyper-V and newer. The VHDX format offers better resiliency 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 Windows 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.

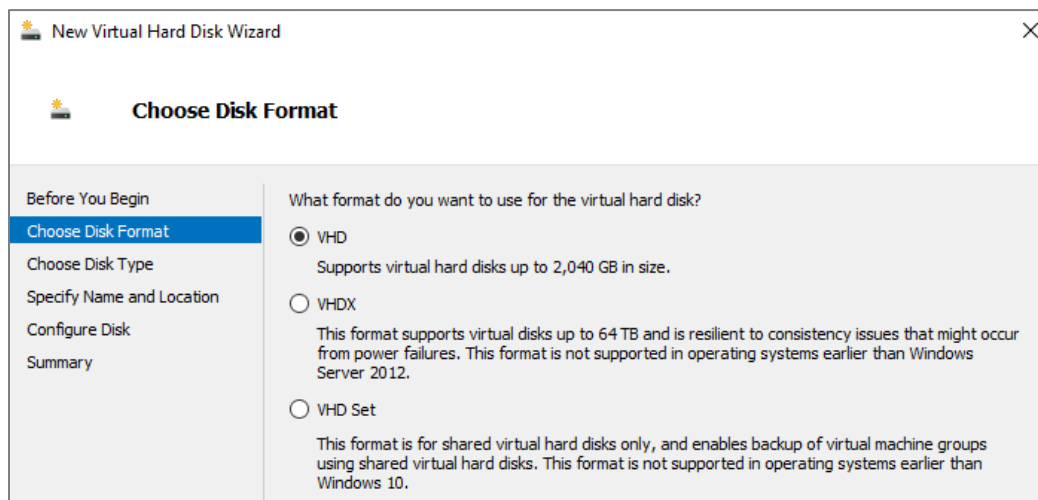


Figure 8 Different formats available for virtual hard disks

### 2.3.2 Virtual hard disk type

In addition to the format, a VHD can be designated as fixed, dynamically expanding, or differencing. Since SC Series storage leverages thin provisioning, only the data that is written to a virtual hard disk, regardless of the disk type, will consume space on the storage array. As a result, determining the best disk type is more a function of workload requirements rather than how it will impact storage utilization on SC Series storage.

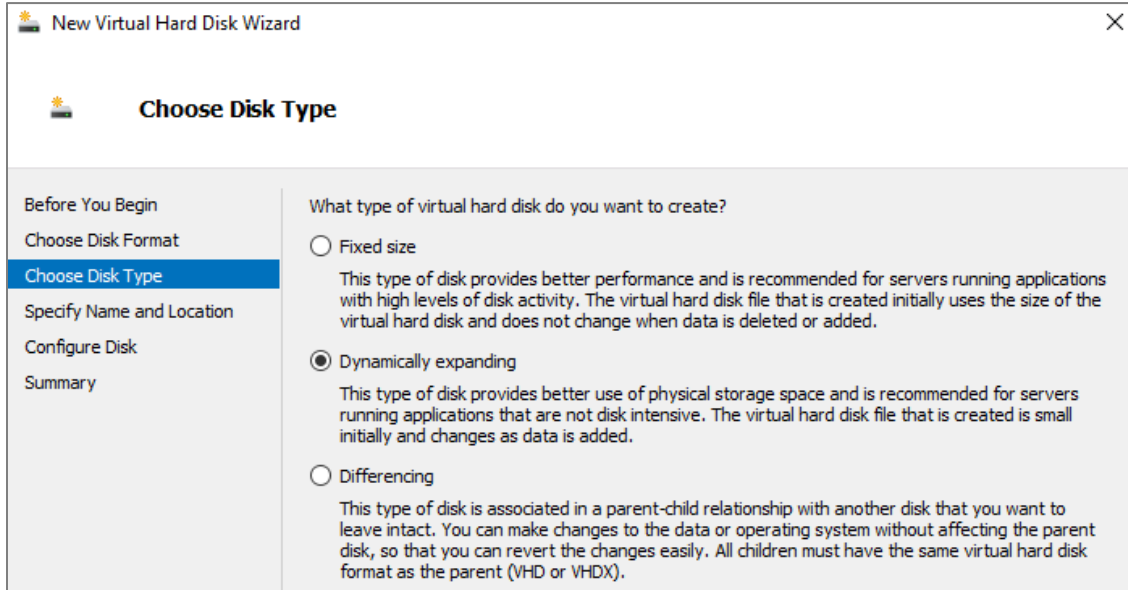


Figure 9 Select a virtual hard disk type

The dynamically expanding disk type will work well for most workloads. For workloads that generate 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 the amount of data on the virtual disk (plus a little extra for metadata) 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.

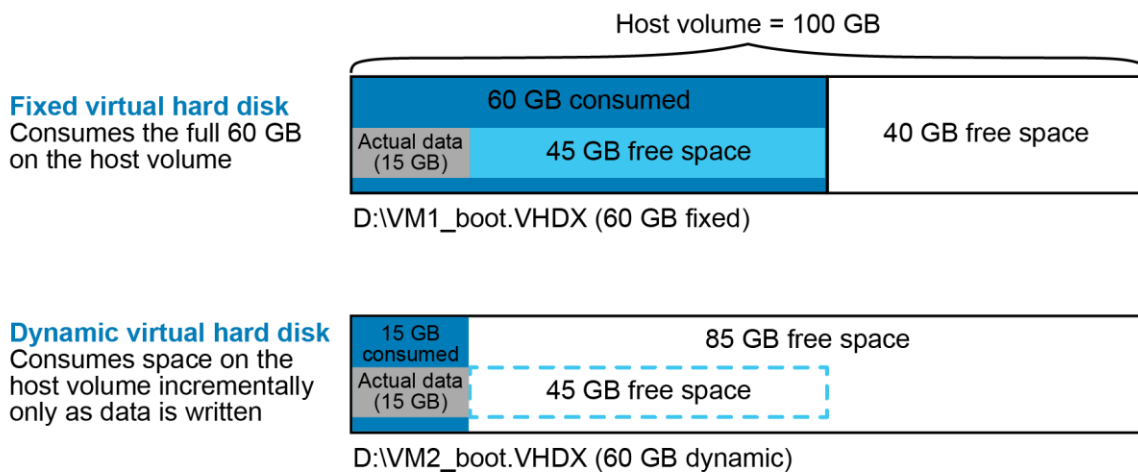


Figure 10 Fixed and dynamic virtual hard disk comparison

SC Series storage supports any of the VHD types. From the perspective of the host server, there are some best practice performance and management considerations to keep in mind when choosing the right kind of VHD type for your environment.

- Fixed-size VHDs:
  - Recommended for workloads with 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 negligible.
  - 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 VHDs may require significant time due to lack of native Offloaded Data Transfer (ODX) support. With Windows 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 VHDs:
  - Recommended for most workloads, except in cases of extremely high disk I/O.
  - When initially formatted, they consume very little space on the host, and expand only as new data is written to them by the guest VM or workload.
  - As they expand, they require a small amount of additional CPU and I/O. 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.
  - Require less time to copy than fixed VHDs.
  - Allows the physical storage on a host server or cluster to be over over-provisioned. It is important to configure alerting to avoid running physical storage out of space when using dynamically expanding VHDs.
- 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.
  - Are typically practical for limited use cases such as a virtual desktop infrastructure (VDI) deployment.
  - All children must use the same VHD format as the parent.
  - Reads of unchanged data reference the differencing VHD.
  - 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 a child VHD, which by default will be written to the highest performing tier and RAID level on the SC Series for maximum performance.

A native Hyper-V based checkpoint (snapshot) of a Hyper-V guest VM creates a differencing virtual hard disk (avhdx) that freezes changed data since the last snapshot. Each additional checkpoint of the VM creates another differencing virtual hard disk, maintained in a hierarchical chain.

- Use of native Hyper-V based checkpoints of a Hyper-V guest VM can impact read I/O performance because data is spread across the virtual hard disk and one or more differencing disks (this injects latency).

- Longer chains of differencing virtual hard disks are more likely to negatively impact read performance. It is therefore a best practice to keep native Hyper-V based checkpoints to a minimum if they are used.
- Administrators can leverage array-based SC Series snapshots to replicate data to other SC Series arrays for archive or recovery of Hyper-V guest VMs and workloads to avoid the use of native Hyper-V checkpoints.
  - Use SC Series Replay Manager to leverage VSS to achieve application consistency when protecting guest VMs. For more information on Replay Manager for Hyper-V, see section 3.1.

### 2.3.3 Virtual hard disks and thin provisioning with SC Series storage

Disk space utilization on SC Series storage is optimized regardless of the type of virtual hard disk used due to the advantages of thin provisioning. For all virtual hard disk types, only the actual data written by a guest VM or workload will consume space on SC Series arrays.

The example below illustrates a 100 GB SC Series (SAN or DAS) volume presented to a Hyper-V host that contains two 60 GB virtual hard disks. The volume is overprovisioned in this case to demonstrate behavior, but not as a general best practice. One virtual hard disk is fixed, and the other is dynamic. Each virtual hard disk contains 15 GB of actual data. From the perspective of the host server, 75 GB of space is consumed and can be described as follows:

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

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

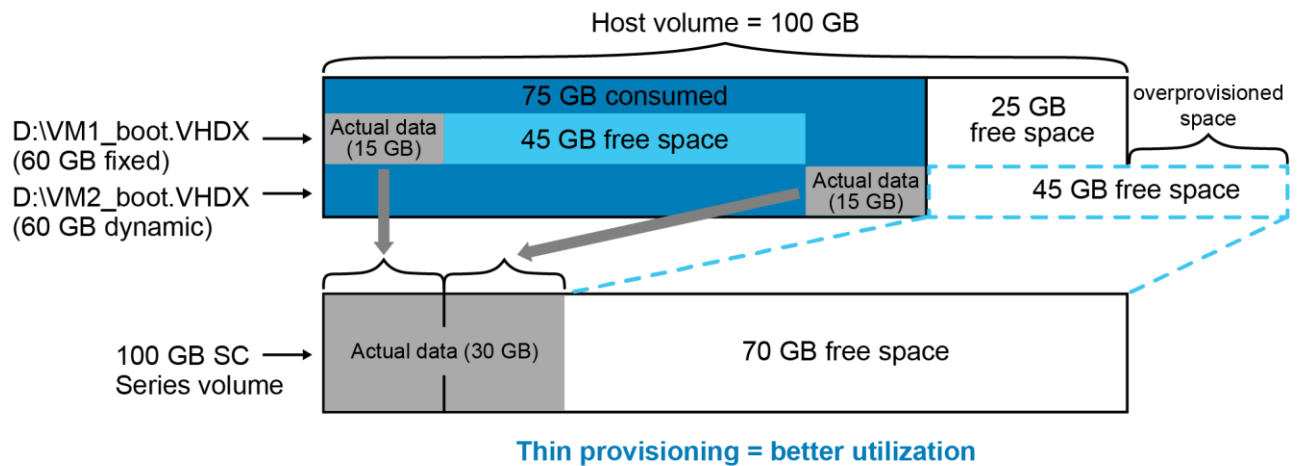


Figure 11 Thin provisioning with SC Series storage

Compare this to how SC Series storage reports storage utilization on the same volume:

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

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



### 2.3.4 Overprovisioning with dynamic virtual hard disks

When using 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 array running out of space. See Figure 11 for an example. If the dynamic disk used by VM2 on the host volume expanded far enough, it would fill up the host volume and negatively impact 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 move the virtual hard disk for VM1 or VM2 elsewhere to free up space on the 100 GB host volume or expand the host volume. In either case, it may be difficult to identify the root cause of the problem, and resolution may require a service interruption.

To mitigate risks, consider the following best practice recommendations:

- Create a Hyper-V volume on the physical host that is large enough so that current and future expanding dynamic virtual hard disks will not fill the host volume to capacity. Creating larger Hyper-V host volumes does not negatively impact space efficiency on SC Series because of the benefits of thin provisioning.
- Hyper-V based checkpoints (snapshots) create differencing virtual hard disks on the same physical volume. Allow adequate overhead on the host volume for the extra space consumed by the differencing virtual hard disks.
- At the host level, set up monitoring on overprovisioned volumes so that if a percent-full threshold is exceeded (such as 90 percent), an alert is generated with enough lead time to allow for remediation.
- At the SC Series level, configure thresholds and alerting so that warnings are generated before a storage tier or storage pool reaches capacity.
  - If tier 1 fills to capacity, new writes are forced into a lower performing tier (if there is capacity in a lower tier), resulting in degraded performance.
  - If all storage tiers in a storage pool fill to near capacity, the SC Series array will enter conservation mode. Assistance from Dell EMC Support may be required to recover from conservation mode.

## 2.4 Present SC Series storage to Hyper-V

There are several ways to present SC Series SAN or DAS storage volumes as LUNs to Hyper-V hosts, nodes, and VMs.

- Present SC Series volumes as LUN 0 to physical Hyper-V hosts or nodes that will boot-from-SAN.
  - Requires a FC or iSCSI adapter that supports boot-from-SAN
  - Boot-from-DAS (hosts with SAS front-end adapters) is not supported

See the [Dell EMC SC Series and Microsoft MPIO](#) best practices guide for details on how to configure boot-from-SAN.

- Present SC Series volumes as data volumes to physical Hyper-V hosts and clusters:
  - Support for Fibre Channel, iSCSI, and SAS-FE
  - Support for mixed transports
  - Leverage server cluster objects on SC Series storage to map data volumes to multiple nodes at the same time to support clustering and cluster shared volumes

Present SC Series volumes directly to guest VMs:

- In-guest iSCSI
- Virtual Fibre Channel (vFC) with Windows Server 2012 Hyper-V and newer
  - > SC Series Replay Manager offers limited support for protecting guest VMs using vFC.
- Present SC Series volumes indirectly to guest VMs as pass-through disks in Hyper-V
  - Use of pass-through disks is a legacy configuration introduced with Windows Server 2008 Hyper-V.
  - While still supported, use of pass-through disks is discouraged by Dell EMC and Microsoft with Windows Server 2012 Hyper-V and newer

## 2.5 Transport options

SC Series storage and Hyper-V support iSCSI, Fibre Channel, or front-end SAS, and the configuration will typically include multipath I/O (MPIO) as a best practice 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, 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.

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

### 2.5.1 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, select SC Series arrays can be configured from the factory to support front-end SAS connectivity. Front-end SAS may work well for small Hyper-V environments deployed in edge cases such as a branch office or other remote location where avoiding the cost and overhead of maintaining infrastructure to support Fibre Channel or iSCSI is desirable.

For more information about SC Series DAS configuration support for Hyper-V, see the guide: [Dell EMC SC Series Storage with SAS Front-end Support for Microsoft Hyper-V](#).

### 2.5.2 Multiple transports

Although SC Series storage arrays support multiple transports, there is limited Microsoft host support for configuring multiple transports on the same host. Except for a few valid use cases, a single transport (iSCSI, FC or SAS-FE) should be used for each host server. In a Hyper-V cluster environment, all nodes should be configured to use a single common transport.

For more information on using multiple transports, see the [Dell EMC SC Series and Microsoft MPIO](#) best practices guide.

## 2.6 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 Series 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 VMs) to avoid service outages when performing routine SAN maintenance. This includes hosts that are configured to use single path.

For more information on MPIO for Hyper-V environments, see the [Dell EMC SC Series and Microsoft MPIO best practices guide](#).

## 2.7 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. If a use case requires direct-attached storage, SC Series storage supports in-guest iSCSI and virtual Fibre Channel to present block storage volumes directly to guest VMs.

- **In-guest iSCSI:** Simply use the iSCSI initiator software in the guest VM and configure the VM to use iSCSI targets on the SC Series. Native iSCSI support is provided with Windows Server 2008 and above.

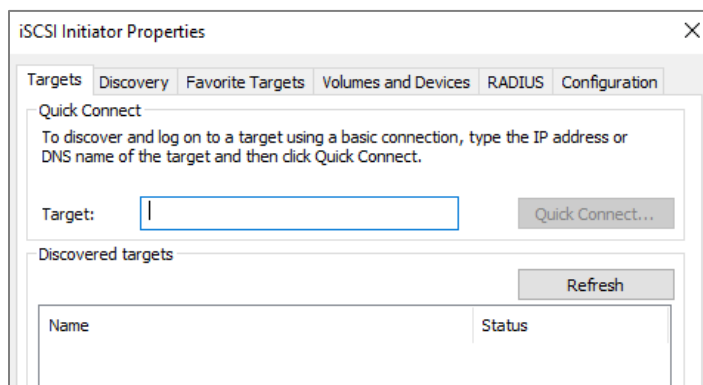


Figure 12 iSCSI Initiator Properties

- **Virtual Fibre Channel:** With Windows Server 2012 Hyper-V and newer, Windows Server guest VMs (2008 R2 and newer) support virtual Fibre Channel (vFC) 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 an available option.
  - Virtual FC requires that all the components in the fabric (HBAs, switches, SC Series controllers) support N\_Port ID Virtualization (NPIV).
  - The setup is more complicated than in-guest iSCSI. This [SC Series Virtual Fibre Channel for Hyper-V](#) demo video provides helpful configuration guidance.
  - SC Series Replay Manager supports protecting Hyper-V guest VMs with in-guest iSCSI and vFC volumes (some limitations exist). See the [Dell EMC SC Series Replay Manager for Hyper-V](#) best practice guide for more information.

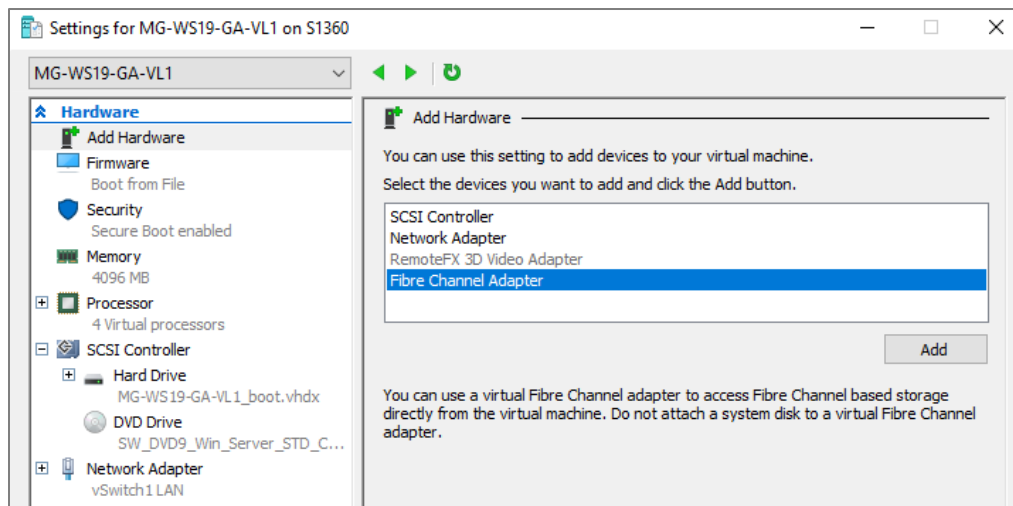


Figure 13 Use VM Settings in Hyper-V Manager to add a virtual Fibre Channel adapter to a guest VM

## 2.8 Guest VMs and direct attached storage

Although SC Series arrays support in-guest 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 the following:

- **Application-consistent backups with SC Series Replay Manager and VSS:** When the Replay Manager agent is installed directly on a guest VM to order to obtain consistent backups of SQL Server or Microsoft Exchange data, the protected data must reside on in-guest iSCSI volumes (SC Series Replay Manager offers limited support for virtual FC volumes).
- **SC Series Replay Manager recovery:** Recovery volumes presented to a guest VM by Replay Manager for restore or recovery operations use in-guest iSCSI.
- **High I/O demand:** Situations where a workload has very high I/O requirements, and the performance gain (even if small) over using a virtual hard disk is beneficial. Direct-attached disks bypass the host server file system. This reduces host CPU overhead for managing guest VM I/O. For many workloads, there will be no notable difference in performance between a direct-attached or virtual hard disk.
- **High availability:** 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.
- **I/O isolation:** When needing to troubleshoot I/O performance on a 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.
- **Data isolation:** When there is a need to create a custom SC Series storage profile, snapshot profile, or replication profile for a specific 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 storage profile(s).
- **Large capacity volumes:** When a single data volume presented to a guest VM will (or may) exceed the maximum size for a VHD (2 TB) or VHDX (64 TB).

Limitations when using direct-attached storage for guest VMs include:

- **Checkpoints:** The ability to perform native Hyper-V checkpoints is lost. However, the ability to leverage SC Series snapshots is unaffected.
- **Simplicity:** More complicated and therefore requires more management overhead. This is especially true for vFC.
- **Mobility:** VM Mobility is reduced due to creating a physical hardware layer dependency.
- **Compatibility:** SC Series Replay Manager offers limited support for virtual Fibre Channel.
- **Zoning requirements:** Virtual FC requires that the host servers and SC Series use soft zoning (by WWN) instead of hard zoning (by switch port).
- **Ease of management:** SC Series cluster server objects are not supported with guest VM clusters using Virtual FC adapters.

**Note:** Legacy environments that use direct-attached disks solely for guest VM clustering (HA) should consider switching to shared virtual hard disks as they modernize their infrastructure.

## 2.9 Guest VMs and pass-through disks

A block-based pass-through disk is a special type of Hyper-V disk that is mapped to a Hyper-V host or cluster, and then is passed through directly to a Hyper-V guest VM. The Hyper-V host or cluster has visibility to a pass-through disk but does not have I/O access. The Hyper-V host will show the disk in a reserved state as only the guest VM has I/O access.

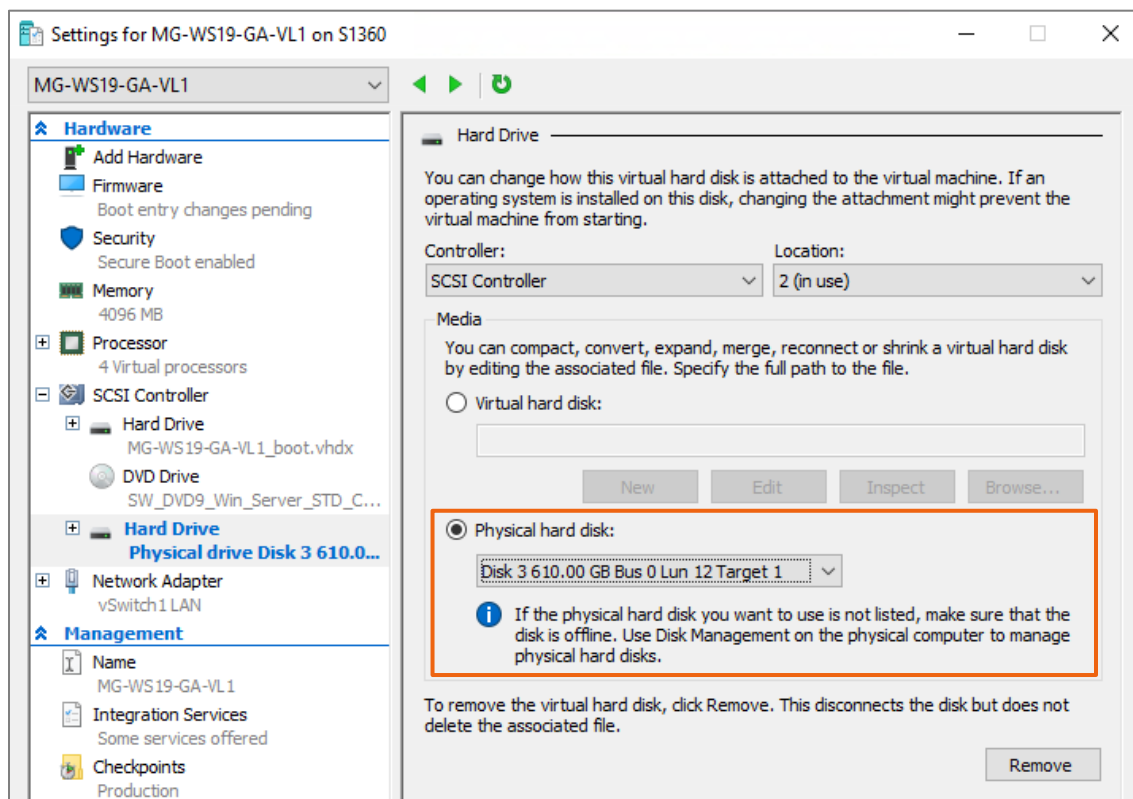


Figure 14 Add an SC Series volume to a guest VM as a pass-through disk

Although SC Series supports pass-through disks, the use of pass-through disks is a legacy design that is discouraged unless there is a specific use case that requires it. They are no longer necessary because of the

feature enhancements offered with newer releases of Hyper-V (generation 2 guest VMs, VHDX format, and shared VHDs in Windows Server 2016 Hyper-V and newer). Use cases for pass-through disks are like those for direct-attached storage (see section 2.8).

Although SC Series Replay Manager supports protecting pass-through disks, use of in-guest iSCSI is recommended, since in-guest iSCSI is required to perform a recovery. See the [Dell EMC SC Series Replay Manager for Hyper-V](#) best practice guide for more information.

Limitations when using pass-through disks includes the list in section 2.8 for direct-attached storage, along with the following:

- **Support for differencing disks is lost:** The use of a pass-through disk as a boot volume on a guest VM prevents the use of a differencing disk. However, SC Series View Volumes (created from a gold image volume) can still be used to maximizing SAN space utilization.
- **Difficult to manage:** The use of pass-through disks becomes unmanageable and impractical at larger scale.
- **LUN number limits:** 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. See section 2.11 for more details.

## 2.10 SC Series arrays and cluster server objects

When mapping shared SC Series volumes (quorum disks, cluster disks, or cluster shared volumes) to multiple Hyper-V nodes, make sure that the volume is mapped to all nodes in the cluster using a consistent LUN number.

Leverage cluster server objects on the SC Series array to simplify the task of mapping a new volume to many nodes at the same time. This can be a significant time saver in larger environments and can help reduce the risk of user error.

---

**Note:** The use of SC Series cluster server objects is not supported when a guest VM cluster is configured to use vFC adapters.

---

Some advantages of using SC Series cluster server objects include the following:

- **Fast:** A new volume is mapped to all nodes in the cluster in one operation.
- **Consistent:** The volume is mapped to each node using a consistent LUN number.
- **Accurate:** Reduces the chance of user error or inconsistencies when mapping volumes.
- **Efficient:** Saves time, particularly as the number of nodes and volumes increases. For example, existing cluster volumes are automatically mapped to new cluster node when a new node is added to the SC Series cluster object.

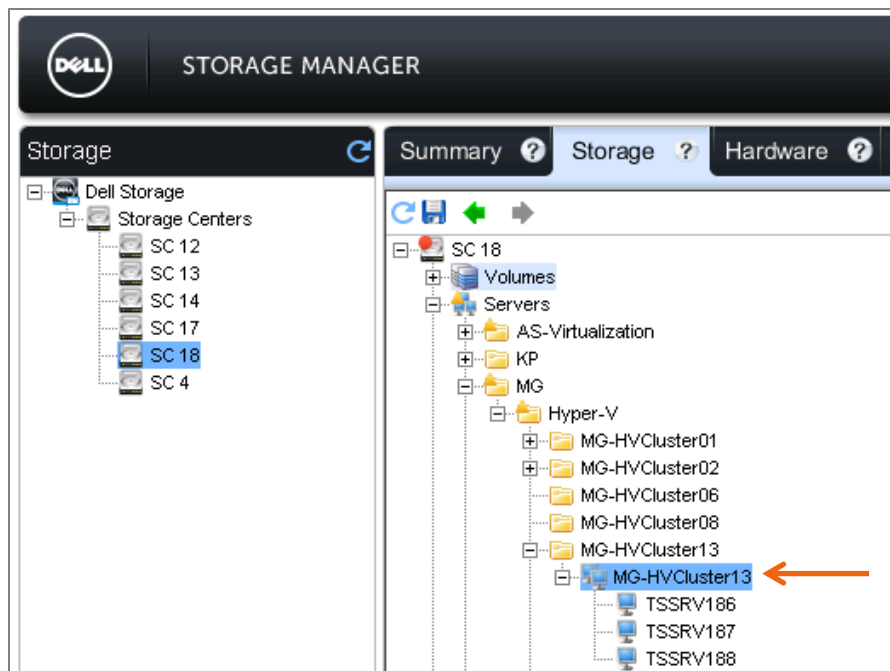


Figure 15 Map storage to a cluster object to ensure consistent LUN numbers on all nodes

## 2.11 SC Series 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 storage 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 volumes, temporary volumes, or SAN maintenance. SAN maintenance might include operations such as expiring snapshots or SC Series Replay Manager restore operations that require LUN numbers on a temporary basis when presenting View Volumes from snapshots 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:

- **Many-to-1:** Use a many VMs-per-CSV strategy when using virtual hard disks.
- **Direct-attached:** Use direct-attached storage (iSCSI or virtual Fibre Channel) to present SAN volumes directly to guest VMs instead of pass-through disks, if using direct attached/pass-through is required.
- **Increase number of data paths:** Add physical FC or iSCSI cards to the cluster nodes to expand the functional limit. This will not increase the 254 LUNs-per-SC Series cluster server limit.
- **Smaller cluster size:** If using pass-through disks, create smaller Hyper-V clusters with fewer nodes and fewer guest VMs instead of larger clusters with more guest VMs.

## 2.12 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 use cases.

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

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

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

- **I/O isolation:** Easier to isolate and monitor disk I/O patterns for a Hyper-V guest VM.
- **Quicker recovery:** 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.
- **Granular control over replicated data:** If SC Series volumes are replicated to a second location, an administrator has more granular control over what data gets replicated (avoid replicating unnecessary data).
- **Agility:** It is often quicker to move a guest VM from one host or cluster to another by remapping the volume rather than copying large virtual hard disk files from one volume to another over the network.

Other strategies might include placing all boot virtual hard disks on a common CSV, and data volumes on other CSVs. Workload vendors may specify an optimal volume configuration to spread I/O over several volumes.

## 2.13 Offloaded data transfer

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 supported in Hyper-V environments. SC Series controllers running SCOS 6.3.1 and newer support ODX with Windows Server 2012 and newer. ODX is enabled by default and it is a best practice to leave it enabled, unless there is a need to obtain performance benchmarks or troubleshoot an ODX issue.

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, see the [Dell EMC SC Series and Microsoft Windows Server](#) best practices guide.



## 2.14 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.

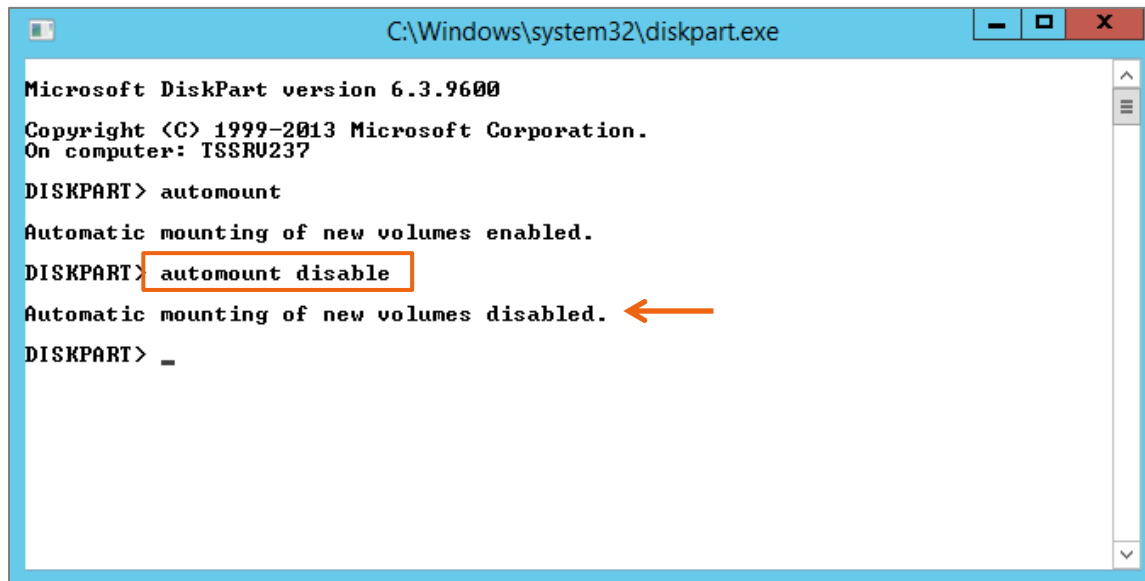


Figure 16 Disable the automount feature

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.

## 2.15 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 specific workload.

With SC Series storage, 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-from-SAN volumes which will conserve space in the disk pool.
- Volumes or virtual hard disks dedicated to page files typically do not require snapshot protection, and therefore do not need to be replicated to a remote SC Series array for DR protection. This is especially beneficial in cases where there is limited bandwidth for replication of volumes and snapshots to other SC Series arrays.
- 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.

## 2.16 Placement of Active Directory domain controllers

It is a best practice to avoid placing all the Microsoft Active Directory® (AD) domain controllers for a domain on the same Hyper-V cluster (as VMs) if the cluster service requires AD authentication to start.

If the cluster goes off line (along with all the domain controller VMs), the cluster service will not be able to start.

To protect against this situation in SC Series environments, it is a best practice to configure at least one domain controller on a physical server with local storage (along with other critical services) so that regardless of the state of the external storage array or storage fabric, critical services such as AD, DNS, and DHCP will be continuously available if the network is also functional.

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

- Place virtualized domain controllers on standalone Hyper-V hosts or on individual cluster nodes if there is an AD dependency for cluster service authentication.
- Use Hyper-V Replica (2012 and newer) to ensure that domain controller VMs can be recovered on another host.
- Leverage Windows Server 2016 Hyper-V and newer which does not have an AD dependency to authenticate cluster services.

## 2.17 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 array.

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](#) white paper.

## 3 SC Series snapshots and Hyper-V

SC Series snapshots are space-efficient, meaning they consist of pointers to frozen data blocks, and therefore consume no additional space unless, for example, a View Volume is created from a snapshot, mapped to a host and new data is written.

SC Series snapshots can be taken of volumes mapped as LUNs to a Hyper-V environment regardless of content. This applies to boot-from-SAN volumes, data volumes, cluster shared volumes (CSV), pass-through disks, in-guest iSCSI volumes, and vFC volumes. These volumes along with their snapshot histories can also be replicated to other SC Series arrays for DR or archive purposes.

SC Series 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.

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 crash-consistent. When recovering a server using a crash consistent snapshot, it is like 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 or loss 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.

### 3.1 SC Series Replay Manager support for Hyper-V

Replay Manager can be used to obtain application-consistent backups (Replays) of Hyper-V guest VMs. 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.

To learn more about Replay Manager for Hyper-V, see the [Dell EMC SC Series Replay Manager](#) best practices guide and [demo video](#). Many other Replay Manager documents and videos are found at [SC Series technical documents and videos](#).

### 3.2 Use SC Series snapshots to recover guest VMs

A Hyper-V guest VM can be recovered to a previous point in time by using crash-consistent SC Series snapshots of the underlying host volume containing the virtual hard disk. 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 arrays is used. This section provides guidance and best practices for several different recovery options using snapshots.

### 3.2.1 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 virtual 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 array so that crash-consistent snapshots occur at the same exact 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 host 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 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 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 may 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 need to 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 hardware configuration (such as number of virtual CPUs, RAM, virtual networks, and IP addresses) in case importing a VM configuration is not supported (Hyper-V 2008 R2) or fails.

### 3.2.2 Recover a 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 like 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 as the source volume.

With standalone Windows or Hyper-V servers, disk ID conflicts are avoided because standalone servers can automatically detect duplicate disk IDs and change them dynamically without user intervention.

However, host servers are not able dynamically change conflicting disk IDs when disks are configured as CSVs when the disks are mapped to two or more nodes concurrently.

When attempting to map a View Volume of a snapshot 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.

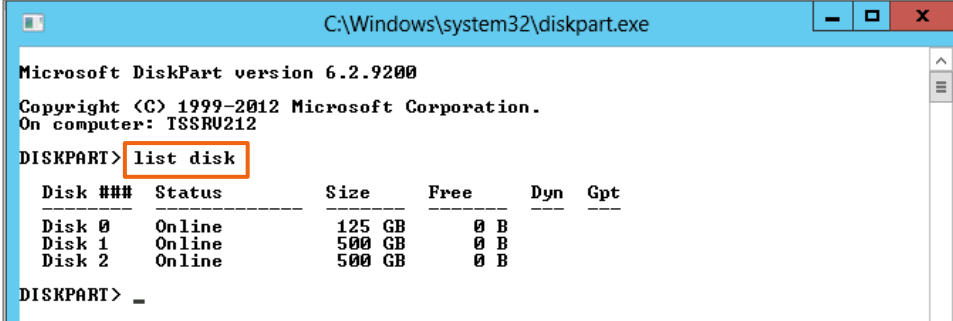
There are a couple of ways of working around the duplicate disk ID issue as detailed below.

- **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.

### 3.3 Change a cluster shared volume 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).



```

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		

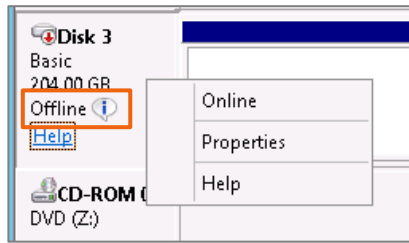
```

DISKPART> _

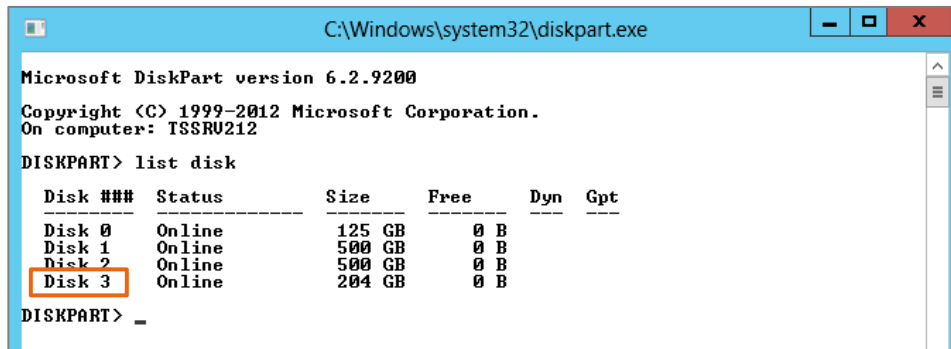
```

6. Use the Dell Storage Manager Client to map a 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**.
14. For **<newid>** provide a new ID.
  - a. 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).

15. Type **uniqueid disk** again and press **Enter** to verify the new ID.

The screenshot shows a Windows command prompt window titled "C:\Windows\system32\diskpart.exe" running Microsoft DiskPart version 6.2.9200. The user has entered the command "list disk", which displays a table of disks. Then, they entered "select disk 3", and the prompt indicates "Disk 3 is now the selected disk." Next, they entered "uniqueid disk", which returned "Disk ID: 12341256". Finally, they entered "uniqueid disk ID=1234abcd", which returned "Disk ID: 1234ABCD" with an arrow pointing to the new ID. The prompt ends with a dash "\_".

```

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> _

```

16. Now that the View Volume has a new signature, it can be unmapped from the standalone host server and mapped to the cluster without causing a disk ID conflict.
17. Recover the guest VM.

### 3.4 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 host servers 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 the original VM continues operation, and the VM copy is configured so that it is isolated from the original VM.

## 3.5 Leverage SC Series to create gold images

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

- Host servers (when using boot-from-SAN)
- Guest VMs that use pass-through disks for boot disks (although use of pass-through disk is not recommended)
- Guest VMs that boot from a sysprepped virtual hard disk as the gold image source

Using gold images provides the following benefits:

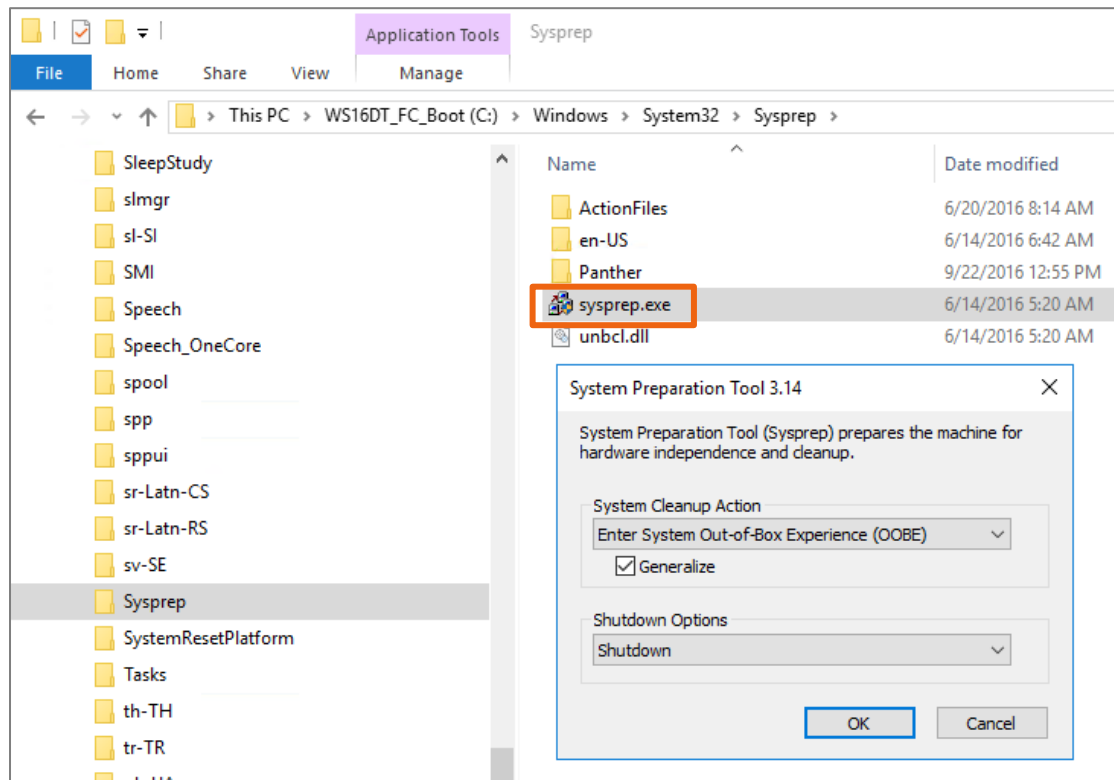
- Faster server provisioning with minimal reconfiguration.
- Better SAN utilization. When a host or VM is provisioned from a gold image, only new data consumes SAN space. Data that has not changed is read from the gold image source volume.
- Disk tiering and Data Progression optimize data placement for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

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

1. Create and map an SC Series volume to a host server that is configured to boot-from-SAN. For more information on boot-from-SAN see the [Dell EMC SC Series and Microsoft MPIO](#) best practices guide.
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), this snapshot can then be used to create an updated gold image source without having to stage the OS from scratch.



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



- 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.
- Using this 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 3.3 for more information on changing a disk ID.
- Boot the host server and allow the initial boot process to complete.
- Customize the server configuration as needed. Leverage PowerShell to automate the workflow if desired.

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

- Create and map a VHD to a VM as the boot volume.
- Follow the steps in the previous section to stage the VM and then run Sysprep.
- After running Sysprep, use Hyper-V Manager to delete the guest VM. This will delete the guest VM configuration files but will preserve the boot virtual hard disk intact. The virtual hard disk is the only file needed for this sysprepped VHD to serve as a gold image for provisioning new VMs.

---

**Note:** 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.

---

- Copy the gold (sysprepped) VHD file in safe location.

5. Create a new VM. Make a copy of the gold VHD and place it in the desired location to serve as the boot volume for the VM.
6. Rename the VHD to reflect the name of the VM or its purpose, and then attach the VHD to the new VM as the boot volume.
7. Power on the VM and customize the VM as needed.

---

**Note:** Although this is not as space efficient, it is a very quick and easy way to provision new VMs. To learn more about how to use SCVMM to provision VMs from a gold image this is space efficient, see the [Dell EMC SC Series Storage and SMI-S Integration with Microsoft SCVMM](#) configuration guide.

---

### 3.5.1 Gold images and preserving balanced SC Series controllers

When a new volume on an SC Series storage appliance consisting of two controller heads is created and mapped to a host server, cluster, or VM, volume ownership is assigned to one controller head or the other as the primary owner. As additional volumes are created and mapped to servers from the SC Series array, 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 head of their choice when mapping them to servers, but this is usually not required unless there is a storage controller imbalance.

During SC Series maintenance that requires staggered controller head 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 that can inadvertently cause the controller heads to become imbalanced. 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 a large number new hosts are deployed from the same gold image, this can result in one controller head 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, array 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.

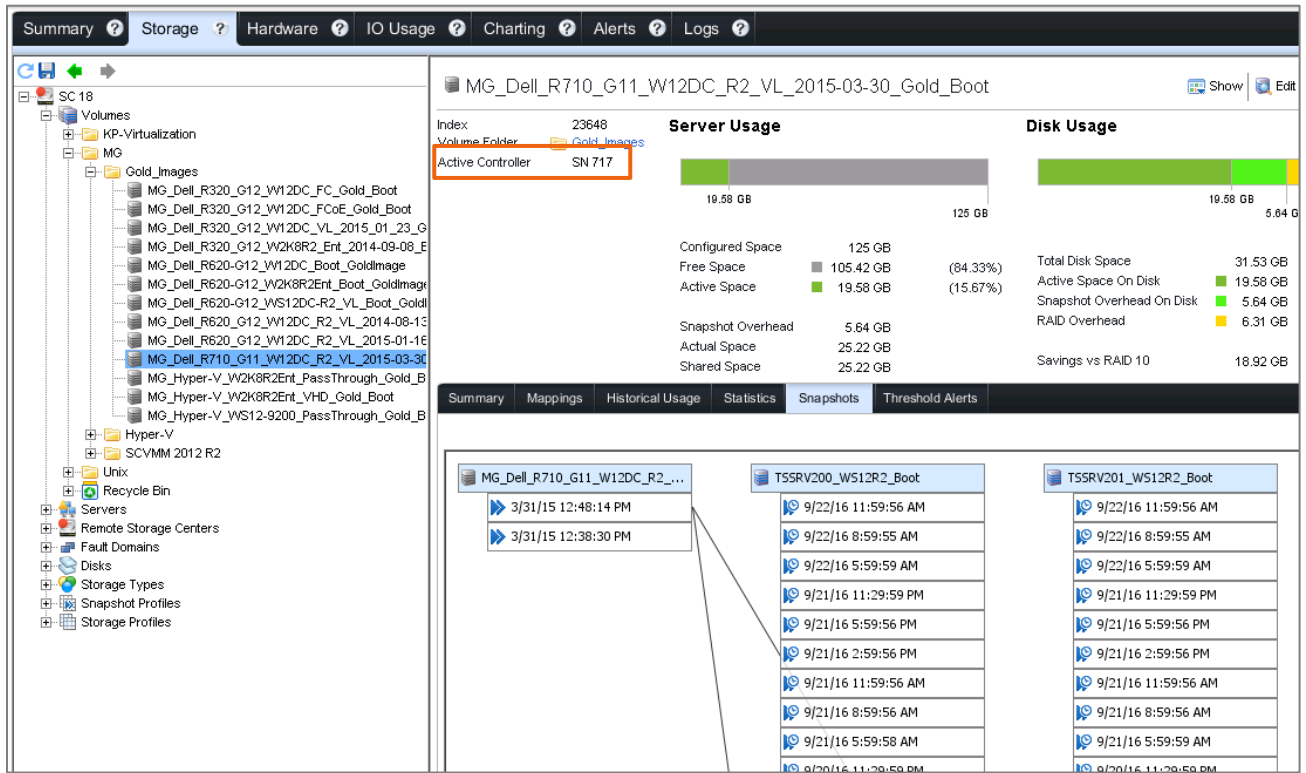


Figure 17 Verify controller ownership for a gold image and associated View Volumes

If a gold image will be used to deploy many new host servers, avoid causing a controller imbalance by creating an additional gold image source that is owned by the other controller head in the pair, in this example, SN 716. All new host volumes provisioned from it will then be owned by SN 716.

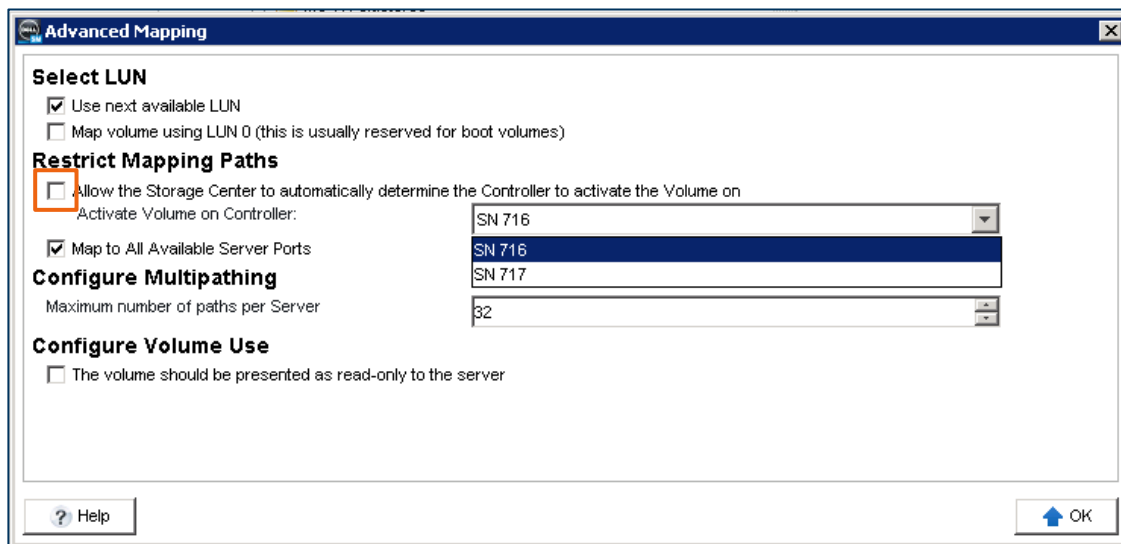


Figure 18 Specify controller ownership for a new gold image volume

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

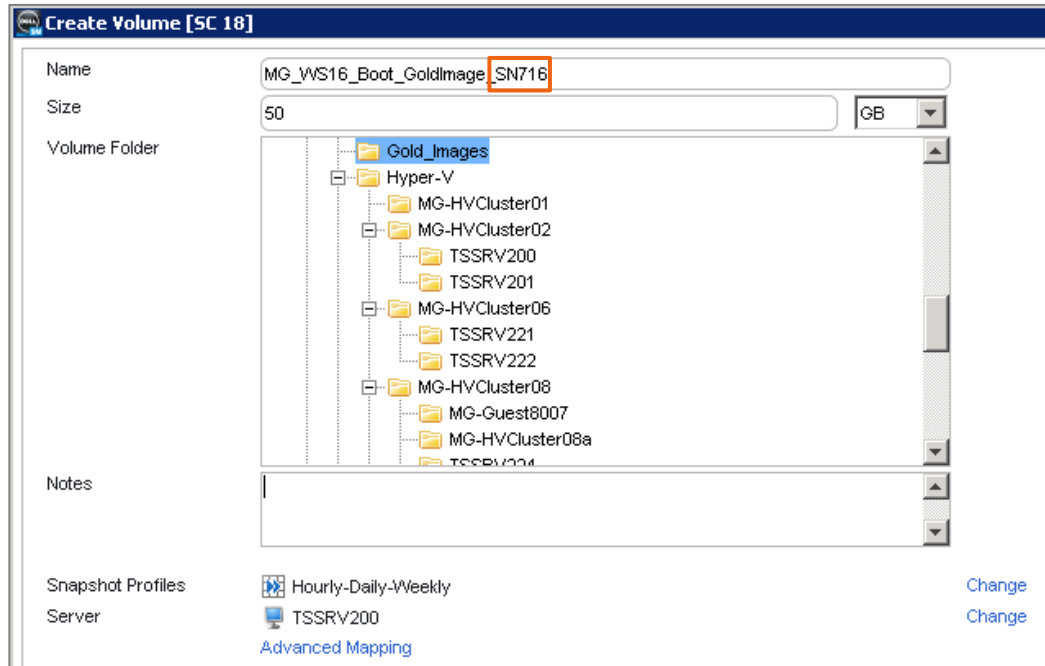


Figure 19 Add the controller serial number to the volume name

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

### 3.6 SC Series snapshots and Hyper-V VM migration

Microsoft provides native tools to move or migrate VMs with Windows Server 2012 Hyper-V and newer, so there are fewer use cases for using SAN-based snapshots to move VMs. When a guest VM is live migrated from one node to another node within the same Hyper-V cluster configuration, no data needs to be copied or moved 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 array to simply unmap the host volumes containing the VM configuration files 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 a large amount 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, or SC Series Federation features. For more information on replication between SC Series including Live Volume with automatic failover for Microsoft, see the [Dell EMC SC Series Synchronous Replication and Live Volume](#) guide.

## 4 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.

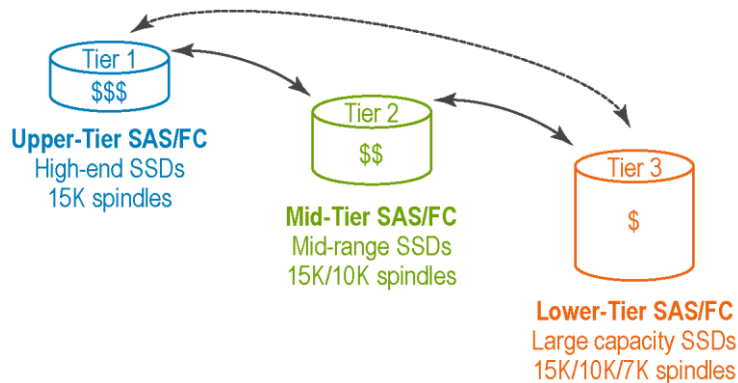


Figure 20 SC Series Data Progression

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

In most environments, about 80% of all data is inactive or archival in nature and therefore, providing a lower tier comprised of large capacity media (spinning or SSD) for Data Progression for long term storage 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 move inactive data to the lowest tier over time. Conversely, if data in a lower tier begins to experience frequent activity, Data Progression will automatically move it to a higher performing tier.

### 4.1 Tuning Data Progression settings for Hyper-V

Since Data Progression is platform-agnostic, there are no extra steps required 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 most Hyper-V environments.




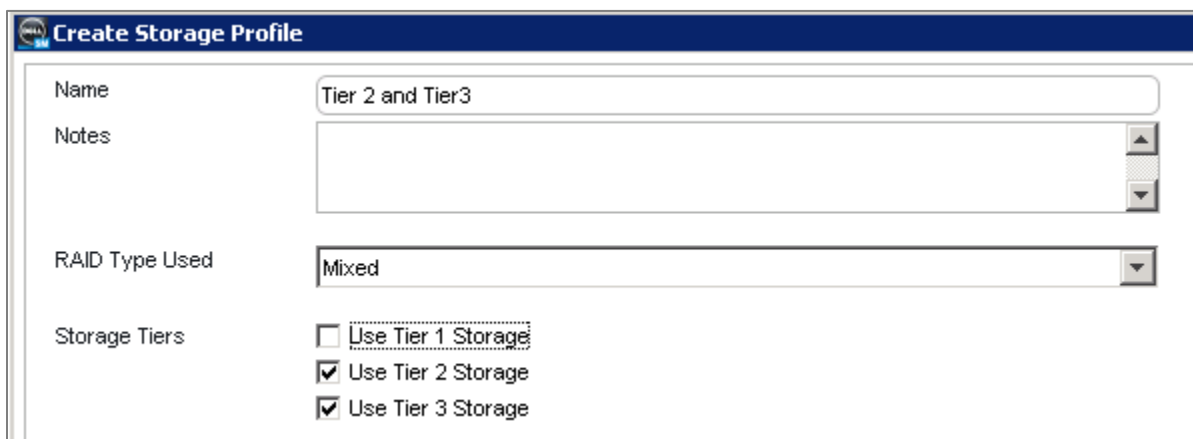
Figure 21 Default storage profiles

Choosing a different storage profile for a volume might be advisable for some Hyper-V configurations. Consider the examples in sections 4.1.1 and 4.1.2.

### 4.1.1 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 once written. With this design, tier 3 would need to ingest data while maintaining adequate application performance.
- **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 kept out of tier 1, and Data Progression will eventually move the data to from tier 2 to tier 3 (about 12 days).



The screenshot shows a 'Create Storage Profile' window with the following configuration:

- Name:** Tier 2 and Tier3
- Notes:** (Empty text area)
- RAID Type Used:** Mixed
- Storage Tiers:**
  - Use Tier 1 Storage
  - Use Tier 2 Storage
  - Use Tier 3 Storage

Figure 22 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.

### 4.1.2 Data copies and migrations

When copying or moving data from one location to another, it is possible to inadvertently consume all 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 a 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.

## 5 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 can result 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 [Dell EMC SC Series and Microsoft Windows Server](#) best practices guide.

### 5.1 SC Series support for Trim/Unmap with Hyper-V

Hyper-V environments also support TRIM/Unmap natively given the following conditions:

- The Windows Server OS (for hosts or clustered nodes) must be version 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 given the following conditions:
  - 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 VM OS is Windows Server 2012 or newer.
  - The guest VM is a generation 2 VM.
  - The guest VM OS formats the virtual hard disk (fixed or dynamic) as a basic disk NTFS volume.

---

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

---

### 5.2 Space recovery with 2008 R2 Hyper-V

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

- Recovery is possible from physical volumes (including boot-from-SAN disks, and direct-attached and pass-through disks).
- Disks must be basic disks formatted as NTFS volumes. Other formats such as ReFS are supported with SC Series but not for disk space recovery.

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

- Virtual hard disks (dynamic or fixed). If free space recovery is highly desirable for a guest VM scenario, then present storage as a 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 (virtual IDE will not support disk space recovery).
  - 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 an NTFS volume, disk space recovery is not supported on cluster shared volumes with Windows Server 2008 R2 Hyper-V hosts.

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

SC Series storage supports boot-from-SAN when hosts are configured with FC or iSCSI cards that also support boot-from SAN. In Microsoft environments, boot-from-SAN is supported with standalone and clustered Hyper-V hosts and nodes. Some of the pros and cons of booting from local disk or from a SAN are detailed below.

### **Boot-from-SAN advantages:**

- SC Series snapshots of boot-from-SAN volumes provide for quick recovery.
- Replicate boot-from-SAN 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-from-SAN volumes can be leveraged to quickly provision new Hyper-V host servers.

### **Boot-from-local-disk advantages:**

- Offline SAN maintenance will not affect the host. Critical roles such as an AD domain controller, DNS and DHCP may need to remain online during offline SAN maintenance or unplanned outages. However, Live Volume (with or without automatic failover) can be used to migrate workloads to another SC array when more than one SC Series array is available.
- The Dell Storage Manager Data Collector/Unisphere™ for SC Series Data Collector can remain online regardless of the state of SC Series array.

### 6.1 Configure Hyper-V hosts to boot-from-SAN

To learn more about how to configure Windows Server Hyper-V hosts to boot from SAN, see the [Dell EMC SC Series and Microsoft MPIO](#) best practices guide.

## 7 PowerShell integration

SC Series storage has incorporated Windows PowerShell for many years and now supports a large 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 storage including many examples, see the [Dell Storage PowerShell SDK Cookbook](#).

### 7.1 Importance of PowerShell

With newer versions of Windows Server and Hyper-V, the use of PowerShell may be required for some day-to-day tasks because there may not be a GUI equivalent in Server Manager, Hyper-V Manager, Failover Cluster Manager, or Windows Admin Center (WAC). This is particularly true when servers are running Server Core and a GUI tool is available only if the server is configured with a desktop. While the real power of PowerShell lies in being able to script and automate complex or repeatable processes, many simple administration tasks are easy to perform with PowerShell.

While initial creation and testing of PowerShell scripts may be time consuming, the return on the investment in cost savings and ease of administration can be significant. 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 [Microsoft PowerShell documentation library](#).

### 7.2 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.

Some use cases include the following:

**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 inefficient and time consuming. By leveraging PowerShell, an administrator could script an automated process that would accomplish the same task in a fraction of the time and save the script to use as a foundation for building similar scripts in the future.

**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 to complete properly. With PowerShell, the process can be automated so that an administrator does not have to be available to perform these steps manually after hours.

**Example 4:** Several order-dependent steps that involve the SC Series array 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).

## 7.3 Best practices for PowerShell

GUI interfaces for the most part keep administrators safe by guiding their steps with wizards that provide warning messages that help protect against inadvertently executing commands that are destructive. The trade-off is that GUI functionality is sometimes limited.

PowerShell, on the other hand, generally allows administrators greater control of their environment — 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, resulting in data loss.

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.

Many PowerShell examples specific to SC Series storage are provided in the [Dell Storage PowerShell SDK Cookbook](#) as mentioned above.

## 8 Business continuity with Hyper-V and SC Series

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 can 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 industry-specific regulations.

The disaster recovery scenarios that may be encountered are diverse and may vary by location. Disasters can be small or large. The loss of a single document that impacts one user is a disaster for that user. 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, cost effective, and easy to implement. They address or prevent most 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.

### 8.1 Cost/risk analysis

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

- 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? Examples include the following:
  - 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 subset of data? Is backing up the email server once a day adequate? If so, the mail server has an RPO 24 hours, meaning the potential for up to 24 hours data loss if the mail server is 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 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 risks?

## 8.2 Disaster recovery and disaster avoidance

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

- Events that cause data loss such as malware infection, corruption, accidental deletion, 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 SC Series Live Volume 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 array.

## 8.3 Live Volume with Auto Failover for Microsoft

Live Volume is an optional SC Series feature that has been available for many years. Auto Failover is an enhancement to Live Volume that has supported Microsoft Hyper-V since the release of SCOS version 7.1. Working closely in conjunction with Live Volume, Dell Storage Manager can be used to create predefined DR plans for Hyper-V environments.

To learn more about how LV-AFO can be configured to protect Microsoft clusters and Hyper-V (including stretched clusters), see the [Live Volume with Auto Failover for Support for Microsoft](#) demo video (parts 1–3) and the [Dell EMC SC Series Synchronous Replication and Live Volume](#) solutions guide.

## 8.4 Replay Manager for Hyper-V

SC Series Replay Manager is a GUI-based client/server 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 many PowerShell cmdlets to allow administrators to automate Replay Manager operations that might otherwise require an administrator to be available after hours.

In Microsoft environments, Replay Manager leverages the power of the Microsoft Volume Shadow Copy Service (VSS) to create and manage application-consistent snapshots (Replays) 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 for Hyper-V see the [Dell EMC SC Series Replay Manager and Microsoft Hyper-V](#) best practices guide and the [Dell EMC SC Series Replay Manager 7 for Hyper-V](#) demo video.

## A Technical support and resources

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

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

### A.1 Related resources

See the following referenced or recommended Dell EMC publications and resources:

- [Dell EMC SC Series Storage Solutions](#)
- [Dell EMC SC Series with SAS Front-end support for Microsoft Hyper-V Configuration Guide](#)
- [Dell EMC SC Series Storage and Microsoft Multipath I/O Best Practices Guide](#)
- [Dell EMC SC Series Virtual Fibre Channel for Hyper-V Demo Video](#)
- [Dell EMC SC Series Replay Manager 7 and Hyper-V Best Practices Guide](#)
- [Dell EMC SC Series Replay Manager 7 and Hyper-V Demo Video](#)
- [Dell EMC SC Series and Microsoft Windows Server Best Practices Guide](#)
- [Dell EMC SC Series Data Reduction with Deduplication and Compression White Paper](#)
- [Dell EMC SC Series Synchronous Replication and Live Volume Solutions Guide](#)
- [Dell EMC SC Series Live Volume with Auto Failover Support for Microsoft Demo Video \(3 parts\)](#)
- [Dell EMC SC Series PowerShell SDK Cookbook](#)
- [Dell EMC SC Series and SMI-S Integration with Microsoft SCVMM Configuration Guide](#)
- [Dell EMC SC Series and Microsoft SCVMM Demo Video](#)

Also see the following referenced or recommended Microsoft publications and resources:

- [Microsoft Windows Documentation Library](#)
- [Microsoft Virtualization Documentation Library](#)
- [Microsoft PowerShell Documentation Library](#)