

Dell EMC SC Series: Microsoft Windows Server Best Practices

Abstract

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

June 2019

Revisions

Date	Description
October 2016	Initial release for Windows Server 2016
November 2016	Update to include BitLocker content
February 2017	Update MPIO best practices
November 2017	Update guidance on support for Nano Server with Windows Server 2016
June 2019	Update for Windows Server 2019 and SCOS 7.4; template update

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 © 2016–2019 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. [5/29/2019] [Best Practices] [680-042-007]

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	5
Audience	5
1 Introduction.....	6
1.1 SC Series overview	6
1.2 Windows Server 2019 overview	7
1.3 SC Series support for previous Windows Server versions.....	7
1.4 Advantages of using SC Series storage with Windows Server	8
1.5 Best practices overview.....	8
1.6 General best practices for Windows Server	9
2 Design best practices	10
2.1 Right-size the storage array and fabric	10
2.2 Determine optimal transport and front-end configuration	11
2.3 Transport options.....	11
2.4 Mixed transports	11
2.5 Multipath I/O	13
2.6 Boot-from-SAN	13
2.7 Adjust timeout settings for single path and multipath I/O	14
2.8 Queue depth and timeouts	14
3 Administration best practices	16
3.1 Present SC Series storage to Windows Server.....	16
3.2 SC Series storage profiles.....	16
3.3 Data Progression.....	17
3.4 Data reduction (compression and deduplication)	17
3.5 Overprovisioning.....	17
3.6 Guest VMs with in-guest iSCSI and virtual Fibre Channel.....	18
3.7 SC Series and server clustering.....	18
3.8 Optimize format-disk wait time for large volumes.....	19
3.9 Placement of page files	20
3.10 SC Series snapshots with Windows Server	20
3.11 TRIM/Unmap	23
3.12 Volume maintenance.....	24

3.13	Offloaded data transfer	26
3.14	Resilient file system	26
4	Best practices for specific roles, features, and integrations	28
4.1	Failover Clustering	28
4.2	Microsoft Hyper-V	30
4.3	Cluster shared volumes	30
4.4	SMI-S	31
4.5	Disaster recovery with Live Volume	31
4.6	Storage Spaces	31
4.7	SC Series Replay Manager	31
4.8	Using BitLocker with SC Series volumes	32
4.9	SC Series and PowerShell SDK	34
A	Technical support and additional resources	35
A.1	Related resources	35

Executive summary

This document provides best practice guidance for deploying Microsoft® Windows Server® with Dell EMC™ SC Series™ storage systems.

The documentation found at Dell.com/support for your specific SC Series array serves as the primary reference material for optimal configuration of SC Series for Windows Server. Available resources (which may vary by SC Series model) include owner's manuals, administrator's guides, installation guides, and release notes.

This document contains supplemental information to optimize your SC Series configuration for Windows Server.

Audience

This document is intended for Dell EMC customers, partners, and employees who desire to learn more about best practices when configuring Windows Server with SC Series storage systems. It is assumed the reader has working knowledge of SC Series storage and Windows Server.

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

1 Introduction

Microsoft Windows Server and Dell EMC SC Series storage are feature-rich solutions that together present administrators with a diverse range of configuration options to solve key business objectives such as storage capacity, performance, resiliency, and security. This section provides an overview of SC Series storage, Windows Server, and general best practices.

1.1 SC Series overview

The SC Series storage array is well suited to provide a solid, proven, high-performance storage solution for Windows Server environments. SC Series storage is designed from the ground up with features and redundancies to optimize performance, maximize storage efficiency, and avoid downtime due to component failures, maintenance, upgrades, and expansion.



Figure 1 SC7020F all-flash storage array

SC Series storage solutions offer enterprise-grade features at mid-range prices including:

- Hybrid and all-flash configurations
- Thin provisioning
- Automated sub-LUN tiering
- Intelligent block-level data placement, migration, and optimization between disk tiers
- Data encryption with self-encrypting drives
- Intelligent data reduction (compression and deduplication)
- Sub-disk RAID levels and distributed sparing
- Synchronous and asynchronous replication
- Live Volume with automatic failover for disaster recovery protection
- Federation with Live Migrate and Volume Advisor
- Cluster server support
- MPIO support for DAS and SAN configurations with SAS, Fibre Channel, and iSCSI
- Data-consistent snapshots (Replays) of Microsoft workloads with Replay Manager
- Easy-to-use management tools for configuration, monitoring, and reporting
- PowerShell integration
- LDAP/AD user authentication
- CloudIQ web-based analytics and monitoring

SC Series arrays provide a robust platform for the ultimate experience in performance, adaptability, and efficiency.

Note: Not all features are supported with all SC Series models. To learn more about a specific SC Series array and available features, visit the [SC Series](#) product page.

1.2 Windows Server 2019 overview

Microsoft Windows Server 2019 is the latest release of the Windows Server operating system. Windows Server 2019 is designed to be the operating system of choice for the modern data center and the cloud. It offers many improvements and enhancements to features that were introduced in previous Windows Server versions, along with many new features. To learn more about specific Windows Server features, visit the [Microsoft Windows IT Pro Center](#).

Windows Server 2019 is supported with SC Series SAN and DAS configurations with the release of SCOS 7.4.

1.2.1 Long-term servicing channel and semi-annual channel

After the release of Windows Server 2016, Microsoft adopted a new model for future Windows Server releases that uses two separate but parallel tracks. Windows Server 2019 is the next-generation release of the long-term servicing channel (LTSC) track which offers five years of mainstream plus five years of extended support for ten years of total support. Windows Server 2019 certification testing with SC Series storage was completed with the October 2018 generally available (GA) LTSC release.

In parallel with the LTSC track is the new semi-annual channel (SAC) track, which offers a core-only release of Windows Server every six months. SAC versions are supported by Microsoft for only 18 months. They allow Microsoft to release new features to customers without a two to three year wait cycle for the next LTSC release. New features in SAC versions roll up for inclusion in the next LTSC release. SAC versions of Windows Server are primarily intended for virtualized test and development environments as opposed to long-term production deployments. SC Series certification testing is conducted with LTSC versions of Windows Server.

1.2.2 Windows Admin Center

In 2018, Microsoft released Windows Admin Center (WAC) as a new locally-installed, browser-accessible HTML5-based tool for consolidated and simplified management of Windows Server environments. It consolidates the most commonly used in-box and remote management tools for core and With Desktop installs of Windows Server to one location. Because WAC is an extensible platform, third parties have the option to develop integrations for their own products or solutions.

Administration and monitoring of SC Series storage from WAC is not currently supported. For more information on WAC, visit the [Microsoft Windows IP Pro Center](#).

1.3 SC Series support for previous Windows Server versions

Several LTSC Windows Server versions that are supported with SC Series are in various phases of mainstream or extended Microsoft support: Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019. Microsoft extended support for Windows Server 2008 R2 is scheduled to end in January 2020.

SCOS 7.4 supports Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019. To verify Windows Server version compatibility with your version of SCOS, consult the applicable SC Series documentation for your array.

The best practice guidance in this document applies to versions of Windows Server currently in mainstream or extended supported by Microsoft (Windows Server 2008 R2 through Windows Server 2019) unless otherwise noted.

Note: Nano Server was originally released as an additional installation option (in addition to Core and With Desktop) for Windows Server 2016. Microsoft support for Nano Server as a physical host or virtual machine (VM) OS was discontinued about a year after its initial release by Microsoft. As of the 1709 (September 2017) semi-annual channel release of Windows Server, Nano Server is supported by Microsoft as a container OS only.

Generally, when a Windows Server OS approaches end of life (end of support), hardware vendors including Dell EMC will also discontinue support for the OS with new product releases. Reasons include the following:

- Testing and certification resources are limited.
- It is much more difficult to obtain hardware specific patches and bug fixes from Microsoft as a Windows Server OS enters extended support phase and approaches end of life.
- It encourages customers to modernize.

1.4 Advantages of using SC Series storage with Windows Server

While Microsoft has continued to make improvements to Storage Spaces Direct (S2D), Storage Replica, and hyper-converged infrastructure (HCI) offerings, the rich feature set and management tools offered with a mature SAN or DAS product such as SC Series storage makes it a compelling choice when architecting the storage component of any Windows Server solution. Microsoft often showcases the IOPS performance of an optimally-configured S2D or HCI configuration. However, raw IOPS potential is only part of the equation when architecting a solution, just as raw horsepower is only part of the equation when deciding what sportscar to purchase.

SAN and DAS features and management tools including intelligent snapshots, managed replication, monitoring, alerting, and DR can be equally or even more important to server and storage administrators than raw IOPS potential. While native S2D or HCI Microsoft solutions may offer similar basic features as a SAN or DAS solution, the rich set of SC features and management tools listed in section 1.1 presents administrators with a much more powerful and complete solution for managing, monitoring, reporting, replicating, and protecting their storage investment and data.

1.5 Best practices overview

Best practices are typically based on and developed from the collective wisdom and experience of many users over time, and this learning is built into the design of next-generation products. With mature technologies such as Windows Server and SC Series storage arrays, best practices are already factored in to the default configurations, settings, and recommendations.

Because default settings typically incorporate best practices, tuning is often unnecessary (and discouraged) unless a specific design, situation, or workload is known to benefit from a different configuration. One of the purposes of this best-practices documents is to call attention to situations where using a default setting or configuration may not be optimal.

Some common goals of best practices include:

- Minimize complexity and administrative overhead
- Optimize the performance of a workload
- Maximize security
- Ensure resiliency and recoverability
- 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:

- A legacy system that is performing well that has not reached its life expectancy may not adhere to current best practices. Often the best course of action is to run legacy configurations until they reach their life expectancy because it is too disruptive or costly to make changes outside of a normal hardware progression or upgrade cycle. Dell EMC recommends upgrading to the latest technologies and incorporating current best practices at key opportunities such as upgrading or replacing infrastructure.
- A common best practices tradeoff is to implement a less-resilient design on lower-tier hardware (to save cost and reduce complexity) in a test or development environment that is not business critical.

Note: While following the best practices in this document is strongly recommended by Dell EMC, some recommendations may not apply to all environments. For questions about the applicability of these guidelines in your environment, contact your Dell EMC representative.

1.6 General best practices for Windows Server

In most cases, Windows Server will perform optimally with default out-of-the-box settings. Tuning may be required depending on the type of server roles, features, and hosted workloads. This paper provides additional guidance for SAN and DAS configurations with SC Series storage. See resources such the [Microsoft Windows IT Pro Center](#) for guidance on general Windows Server best practices that are not specific to external storage.

Some of the most common design, installation, configuration and tuning best practices for Windows Server include the following:

- Keep the design simple.
- Ensure redundancies with core design elements to eliminate single points of failure.
- Leverage Server Core to minimize the attack surface of a server and reduce administrative overhead.
- Leverage Windows Admin Center to centrally manage hosts and clusters.
- Ensure that all hosts and VMs are patched regularly.
- Provide means for adequate malware protection.
- Ensure that essential data is protected with back-ups that meet recovery time objectives (RTO) and recovery point objectives (RPO).
- Minimize or disable unnecessary hardware devices and services to free up host CPU cycles and memory that can be used by other processes (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 human error (this can also help reduce administration time).

2 Design best practices

This section provides guidance on sizing and configuration options for SC Series storage and Windows Server.

2.1 Right-size the storage array and fabric

Optimizing performance is a process of identifying and mitigating design limitations that cause bottlenecks — the point at which performance begins to be impacted under load because a capacity threshold is reached somewhere within the overall design. The goal is to maintain a balanced configuration across the design that allows the workload to operate at or near peak efficiency.

Before deploying a new SC Series storage array, it is important to consider the environmental design factors that impact storage capacity and performance so that new or expanded storage is right-sized for the environment. If the SC Series array will be deployed to support an existing Windows Server workload, metrics such as storage capacity and I/O demands might already be understood. If the environment is new, these factors need to be determined to correctly size the storage array and the storage fabric.

Many common short- and long-term problems can be avoided by making sure the storage part of the solution will provide the right capacity and performance in the present and future. Scalability is a key design consideration. For example, Windows Server clusters can start small with two nodes, and expand one node at a time, up to a maximum of 64 nodes per cluster with newer operating systems. Storage including SC Series arrays can start with a small number of drives and expand capacity and I/O performance over time by adding expansion enclosures with more drives, along with additional front-end ports on supported models, as workload demands increase. With [SC Series Live Migrate and Live Volume](#), up to 10 similar or dissimilar arrays can be clustered as part of a federation to support horizontal and vertical scaling.

One common mistake made when sizing a storage array is assuming that total disk capacity translates to disk performance. Installing a small number of large-capacity spinning drives in an array does not automatically translate to high performance just because there is a lot of available storage capacity. There must be enough of the right kind of drives to support the I/O demands of a workload in addition to raw storage capacity. With SC Series, customers can choose between spinning, hybrid, and all-flash arrays depending on the needs of the workload.

Another common mistake is incorrectly sizing (or failing to anticipate the I/O demand on) the storage fabric. A high-performance storage array will not perform to expectations (and may appear to be the performance-inhibiting culprit erroneously) if the speed or capacity of the FC or SCSI fabric (HBAs, NICs, switches) severely limits I/O throughput because the fabric design is significantly undersized for the storage array. If a storage array or storage federation is expected to generate a certain level of IOPS, make sure the fabric is right-sized to support this level of I/O also.

Where available, customers can confidently use the configuration guidance in Dell EMC storage reference architecture white papers as good baselines to right-size the design of their environments. See [SC Series technical documents and videos](#) for a list of reference architecture white papers.

Work with your Dell EMC representative to complete a performance evaluation if there are questions about right-sizing an SC Series storage solution or fabric for your environment and workload.

2.2 Determine optimal transport and front-end configuration

Depending on the model, the SC Series is configurable as direct-attached storage (DAS) or as part of a storage area network (SAN). DAS is supported with select SC Series models that are configured from the factory with SAS front-end ports.

The most common configuration for SC Series is a SAN configuration leveraging iSCSI, Fibre Channel (FC), or both types of front-end ports. SAN configurations offer customers great flexibility and scale potential when designing their environment.

Consider the following best practice recommendations:

- If an environment is likely to scale beyond four physical hosts or nodes attached to the same SC Series array, choose a SAN configuration (FC or iSCSI).
- DAS configurations may work well for edge cases such as a remote office or branch office that requires a small number of hosts.
- With DAS, the server hosts must be located near the SC Series array (this distance is governed by the length of the front-end SAS cables).
- For SAN configurations, choosing the type of transport is often a function of what is already in place in the environment or according to customer preference.
- In cases where the infrastructure to support an FC or iSCSI SAN is already in place, customers can continue to maximize the return on this investment.

2.3 Transport options

In the early days of SAN technology, Fibre Channel was the dominant transport. Improvements with iSCSI (speed and dependability) in recent years have made it an equally good choice in most cases for a storage fabric if the fabric is designed correctly.

Note: Although FC and iSCSI are considered transports in the purest sense, the term transport is often used interchangeably with protocol in the storage industry.

Windows Server hosts and SC Series arrays support FC and iSCSI in a SAN configuration.

Windows Server hosts and SC Series arrays support front-end SAS (up to 4 hosts per SC array) in a DAS configuration.

Typically, an environment such as a data center is configured to use a preferred transport when it is built and will be part of the infrastructure's core design. When deploying Windows Server hosts to existing environments, the existing transport is typically used.

Deciding which transport to use in an existing or new environment is usually based on customer preference. Influencing factors include size of the environment, cost of the hardware, and existing familiarity with the technology.

2.4 Mixed transports

The use of mixed transports in a SAN environment (access to FC and iSCSI fabrics concurrently) is fully supported by SC Series storage. It is not uncommon, especially in larger environments, for a customer to configure SC Series storage to use mixed transports. This might be required to support diverse platforms that

are collocated that have different transport requirements. It is also common during a transition period if an environment is migrated from one transport to another.

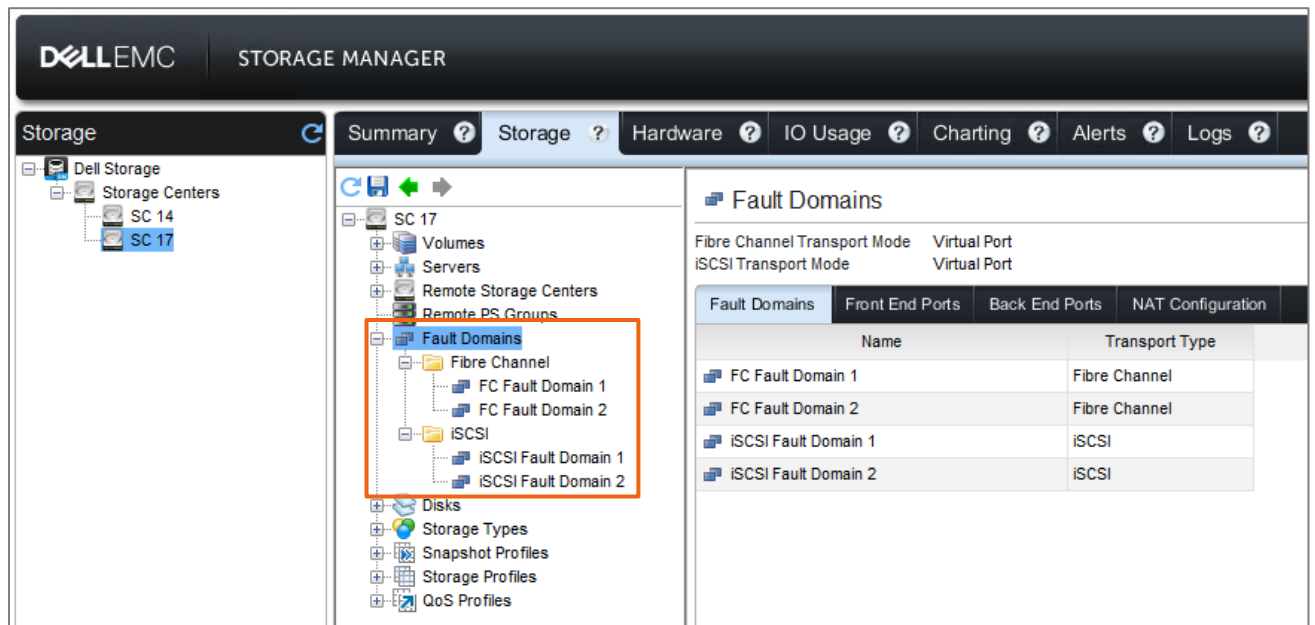


Figure 2 SC Series storage supports mixed transports

Configuring any given volume (mapped to a host server as a LUN) to use multiple transports generally should be avoided because it increases design complexity unnecessarily and can introduce unpredictable service-affecting I/O behavior in path failure scenarios.

For example, a Windows Server (2012 R2 and newer) OS will detect the presence of FC and iSCSI paths when a LUN with multiple transports is mapped to it, but during discovery, the host will choose a primary transport based on an enumeration algorithm which typically results in the host selecting the Fibre Channel paths to pass I/O and will ignore the iSCSI paths. The problem is, if all paths for the primary transport go down, the host server will not automatically start sending I/O using the alternate transport without a manual disk re-scan or host reboot which can obviously be service-impacting.

Because there is limited Microsoft support for using mixed transports for any given LUN, it is not recommended as a best practice. In addition, in a cluster environment, all nodes should be configured to use a uniform transport (FC, iSCSI, or SAS).

There are some use cases where using mixed transports on the same Windows host may be necessary, such as when migrating the overall environment from one type of transport to another, and both transports need to be available to a host during a transition period. If multiple transports must be used, use a single transport for any given LUN.

- Map a LUN to a Host using FC paths only: Acceptable
- Map another LUN to the same Host using iSCSI paths only: Acceptable
- Map a LUN to a Host using FC and iSCSI paths concurrently: Not recommended

2.4.1 Limit volumes to one transport

When multiple transports are available to a host, when mapping a volume to that host, limit the transport to iSCSI or FC by unchecking the **Map to All Available Server Ports** option under **Advanced Mapping** in the **Create Volume** wizard. Select either the FC or the iSCSI paths and complete the wizard.

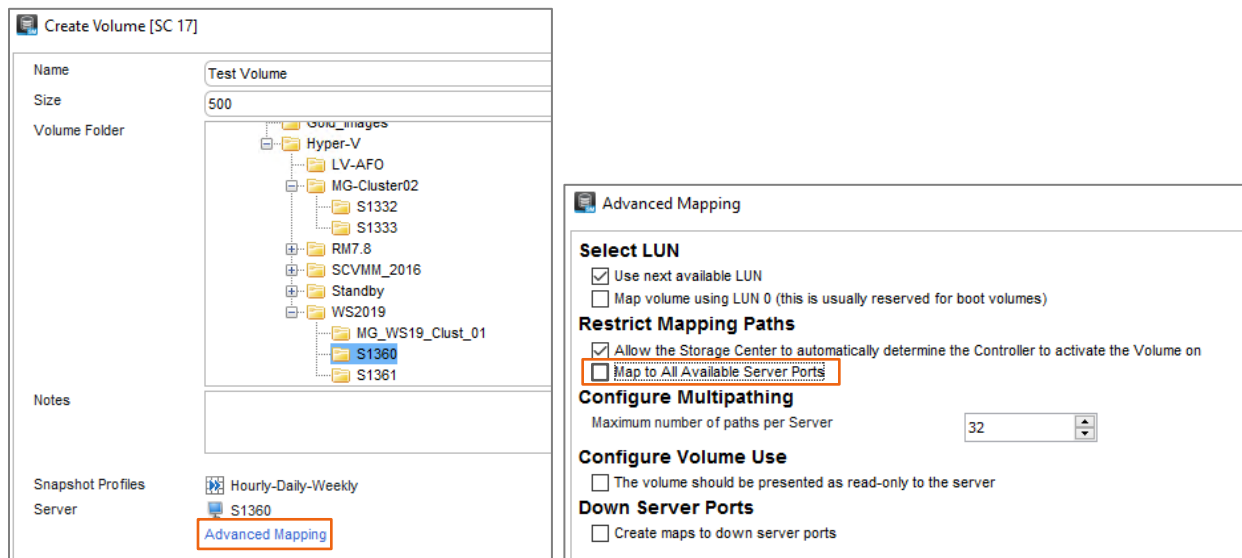


Figure 3 Limit I/O paths to one transport

2.5 Multipath I/O

Regardless of the transport chosen, it is a best practice to ensure redundant paths for each SC Series volume mapped to a Windows Server or cluster as a LUN.

The main purpose of MPIO is to provide redundant paths for a server to access external storage. With multiple paths, if one path goes down, another path is available to provide connectivity to prevent a service outage. MPIO also allows for load-balancing over multiple paths (round robin is typical) to improve performance when there is more than one active-optimized path available.

For more information on how to install and configure MPIO and adjust timeout settings, see the [Dell EMC SC Series Storage and Microsoft Multipath I/O](#) best practices guide.

Note: See section 2.7 for important guidance on adjusting timeout values for single path and MPIO Windows Server hosts.

2.6 Boot-from-SAN

SC Series storage supports MPIO boot-from-SAN with Windows Server with iSCSI or FC HBAs. This is supported in the sense that the HBA natively supports boot-from-SAN in its firmware or BIOS, and the HBA is listed in the [Dell EMC Storage Compatibility Matrix](#) for SC Series storage.

For more information on how to configure boot-from-SAN for Windows Server with MPIO, consult the product documentation for the HBA and the [Dell EMC SC Series Storage and Microsoft Multipath I/O](#) best practice guide.

2.7 Adjust timeout settings for single path and multipath I/O

The registry changes to timeout settings specified in the [Dell EMC SC Series Storage and Microsoft Multipath I/O](#) best practices guide must be applied to all Windows Servers that access storage on SC Series arrays. This includes all FC, iSCSI, and SAS front-end configurations that use single path or multipath.

This is necessary because some default timeout settings in Windows Server are too brief to accommodate brief pauses to I/O during routine maintenance such as an SCOS upgrade, when volume ownership may change from one controller head to another. Leaving the default Windows timeout settings unchanged may result in unintended service interruptions.

2.8 Queue depth and timeouts

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

2.8.1 When to change queue depth

One question that is commonly asked is when to change queue depth settings for Windows Server hosts. On a Windows Server, queue depth is a function of the Microsoft storport.sys driver and the vendor-specific miniport driver for the FC HBA, iSCSI NIC, or CNA.

In many cases, there is no need to change the default queue depth, unless there is a specific use case where changing the queue depth is known to improve performance. For example, if a storage array is connected to a small number of Windows Server hosts with large block sequential read application workloads, increasing the queue depth setting may be very beneficial. However, if the storage array has many hosts all competing for a few target ports, increasing the queue depth on a few hosts might overdrive the target ports and negatively impact the performance of all connected hosts.

While increasing the queue depth can sometimes increase performance significantly for specific workloads, if it is set too high, there is an increased risk of overdriving the target ports on the storage array. Generally, if transactions are being queued and performance is being impacted, and increasing the queue depth results in saturation of the target ports, then increasing the number of target ports, if supported (front-end ports on the SC Series array) to spread out I/O can be an effective remediation.

2.8.2 Vendor-specific HBA and CNA queue depth settings

It is important to understand the firmware and miniport driver registry settings for your host server FC HBA, iSCSI NIC, or CNA adapter and how these settings affect queue depth. In the case of QLogic® FC HBAs for example, the execution throttle setting can be adjusted to control queue depth.

Note: It is very important to verify that each host server HBA, NIC, or CNA adapter used for FC, iSCSI, or SAS connectivity to SC Series storage is supported. View the [Dell EMC Storage Compatibility Matrix](#) for more information. It is equally important to install the latest firmware or boot code along with the most recent Windows Server driver for each host server adapter. Unnecessary troubleshooting is avoided by following this guidance.

See the documentation for your FC HBA, iSCSI NIC, or CNA for direction on adjusting firmware or registry settings to modify queue depth. For example, see this [QLogic support article](#) which explains in detail the relationship between the Windows Server storport.sys driver and the vendor-specific miniport driver, and how to modify parameters that affect queue depth by editing registry keys or firmware settings.

Note: Changes to firmware or registry settings that affect queue depth should be evaluated in a test environment prior to implementation on hosts with production workloads.

3 Administration best practices

This section covers SC Series administration best practices for Windows Server.

3.1 Present SC Series storage to Windows Server

There are several ways to present SC Series storage volumes as LUNs to Windows Server hosts in single path or MPIO (recommended) configurations.

- Physical Windows Server hosts support FC, iSCSI, or SAS connectivity
 - FC and iSCSI in a SAN configuration
 - SAS in a direct-attached storage (DAS) configuration
- VMs support in-guest iSCSI and virtual FC in a SAN configuration

For more information about cabling guidance, see the documentation specific to your SC Series array (such as the Administrator's Guide, User's Guide, or Deployment Guide).

3.2 SC Series storage profiles

Storage profiles define the RAID level used to protect data on a volume and the tiers where that data is stored. The information in the profile is used by Data Progression algorithms when moving existing data pages, as well for optimal placement of new SC Series data pages.

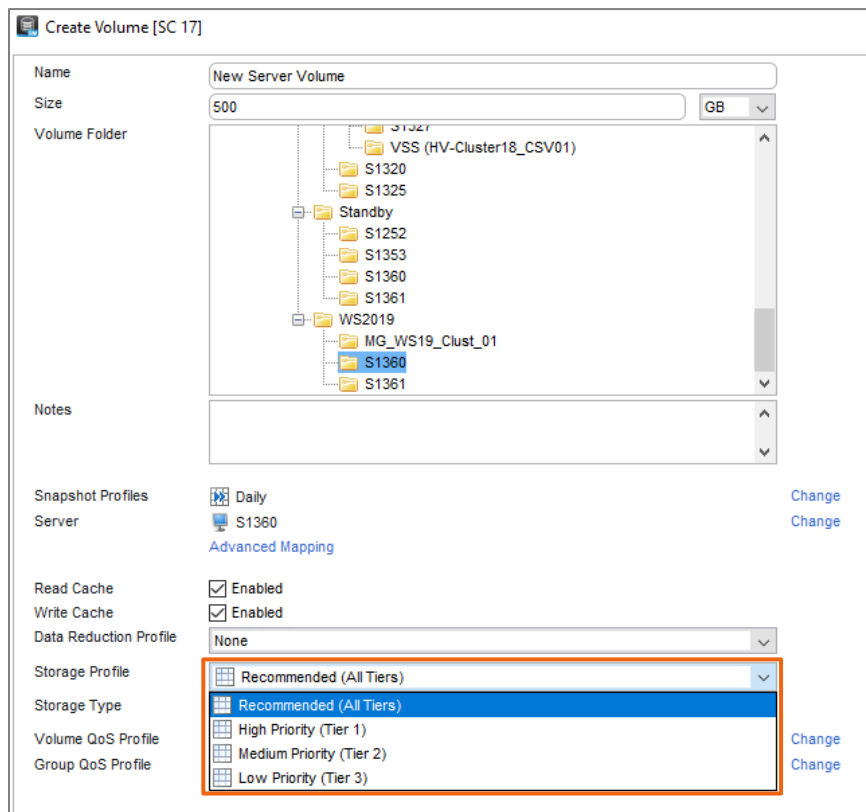


Figure 4 Storage profile selection

For most Windows Server environments, when creating a new volume, the default storage profile **Recommended (All Tiers)** is the best choice to provide good I/O performance and protection for the volume.

It is recommended to use this storage profile first and evaluate its suitability before attempting to change the storage profile.

Using a different storage profile can be very beneficial in some cases. For example, if a data volume on a Windows host contains archival data that will rarely be accessed, a lower priority storage profile may be the best choice, so the data is placed in a more cost-effective tier and RAID level by default without needing to be moved there over time by Data Progression.

3.3 Data Progression

Data Progression is a core SC Series feature with tiered storage. With Data Progression, data is automatically and intelligently placed in the optimal storage tier and RAID level based upon usage and performance metrics. Data Progression runs daily at 7:00 p.m. by default. The start time can be changed if desired.

Data Progression runs seamlessly in the background and there are no special considerations for Windows Server hosts, other than leveraging SC Series storage profiles for initial data placement.

Note: As a Data Progression best practice, apply at least one snapshot profile to each volume with a snapshot expiration of at least 25 hours. The Data Progression process will not work optimally for a volume unless the volume has at least one unexpired snapshot available at the time Data Progression runs.

3.4 Data reduction (compression and deduplication)

Released with SCOS 7.0, data reduction features take advantage of both compression and deduplication for reducing the data footprint of a volume, which can be very beneficial to Microsoft Windows environments. With SC Series, data reduction typically occurs post-process during the daily Data Progression cycle. Depending on the data reduction policy set on a volume, data is compressed and/or deduplicated to achieve maximum storage space efficiency.

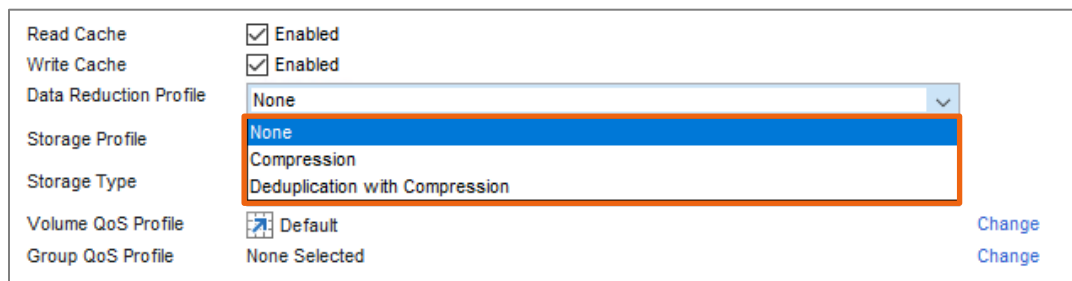


Figure 5 Data reduction options

Like Data Progression, data reduction runs seamlessly in the background. When to enable host-based compression and/or deduplication in addition to SC Series data reduction is a common question, and this is addressed in the [Dell Storage Center OS 7.0 Data Reduction with Deduplication and Compression](#) guide.

3.5 Overprovisioning

Because storage is virtualized and thinly-provisioned with SC Series, it is possible to create one or more host volumes with a max size (individually or cumulatively) that exceed the amount of physical storage space

available on the array. Overprovisioning is possible because only actual data consumes space on the array. This translates to much better storage utilization.

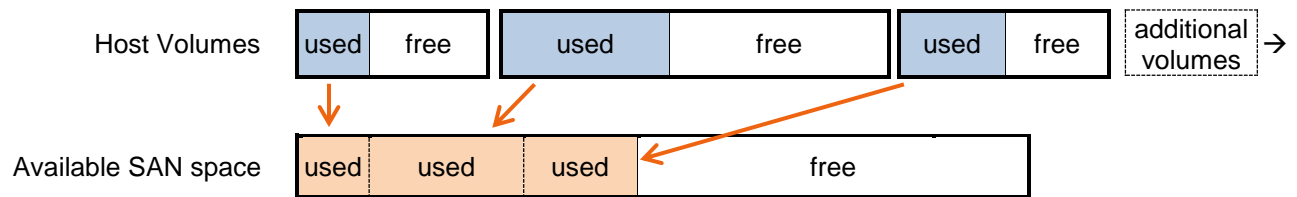


Figure 6 Better storage utilization with thin provisioning

The risk with overprovisioning is running the array out of physical space inadvertently, even though individual volumes from the perspective of host servers might still report free capacity.

To mitigate risks when an SC Series array is overprovisioned, consider these best practice recommendations:

- Set up SC Series monitoring with threshold alerts so that if a percent-full threshold is exceeded for any given disk tier or disk pool (such as 90 percent), an alert is generated with enough lead time to allow for remediation
- Configure notification settings so administrators receive SC Series alerts automatically

Although it is much easier with modern Windows Server OS versions to expand boot-from-SAN and data volumes by using native Windows tools and without down time, it is still generally preferable to size a volume initially so that it will not need to be expanded later. Avoiding disk expansion is particularly helpful with cluster volumes mapped to multiple hosts, and cluster shared volumes (CSV) that hold dynamically expanding virtual hard disks.

3.6 Guest VMs with in-guest iSCSI and virtual Fibre Channel

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

Guest VMs also support virtual Fibre Channel (vFC) with SC Series storage if all the components in the fabric support N_Port ID Virtualization (NPIV). Support for NPIV is typical with modern FC hardware.

For more information, see the [Dell SC Series Storage and Microsoft Hyper-V](#) and [Dell SC Series Storage and Microsoft Multipath I/O](#) best practices guides.

3.7 SC Series and server clustering

When mapping shared volumes (quorum disks, cluster disks, or cluster shared volumes) to multiple hosts, make sure that the volume is mapped to all nodes in the cluster using a consistent LUN number. Leverage host groups on the SC Series array to simplify the task of mapping a consistent LUN number to multiple hosts.

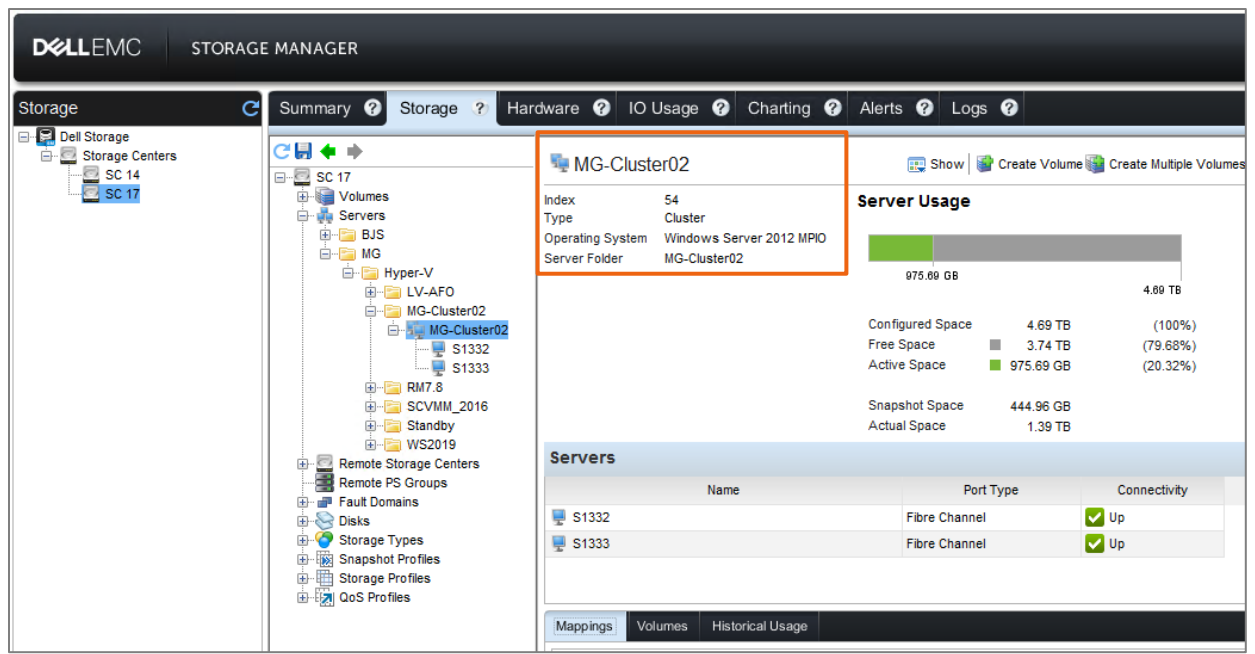


Figure 7 Cluster object on SC Series with two nodes

3.8 Optimize format-disk wait time for large volumes

Formatting a large SAN volume mapped to a Windows host should complete in a few seconds. This should always be true with modern Windows Server operating systems. If long format wait times are experienced for larger volumes (minutes instead of seconds) with older versions of Windows Server, disable the file system **Delete Notify** attribute on the host by completing the following steps:

1. Access a command prompt on the host server with elevated (administrator) rights.
2. To verify the state of the attribute, type **fsutil behavior query disableddeletenotify** and press **Enter**. A result of zero for NTFS or ReFS means the attribute is enabled.
3. To disable the attribute, type **fsutil behavior set disableddeletenotify NTFS 1** or **ReFS 1** and press **Enter**.

The result should display an attribute value of one. To test the result, map a large temporary volume (several TB) from the SC Series array to the host and format the volume. It should complete in a few seconds.

```

Administrator: Cmd Prompt

C:\Windows\system32>fsutil behavior query disableddeletenotify
NTFS DisableDeleteNotify = 0 (Disabled)
ReFS DisableDeleteNotify = 0 (Disabled)

C:\Windows\system32>fsutil behavior set disableddeletenotify NTFS 1
NTFS DisableDeleteNotify = 1 (Enabled)

C:\Windows\system32>fsutil behavior query disableddeletenotify
NTFS DisableDeleteNotify = 1 (Enabled)
ReFS DisableDeleteNotify = 0 (Disabled)

C:\Windows\system32>

```

Figure 8 Modify disableddeletenotify attribute

3.9 Placement of page files

Windows Servers and VMs typically place the page file on the boot volume by default, and automatically manage page file size 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. Ultimately, each customer will need to decide on the best strategy as they consider variables that are unique to their environment.

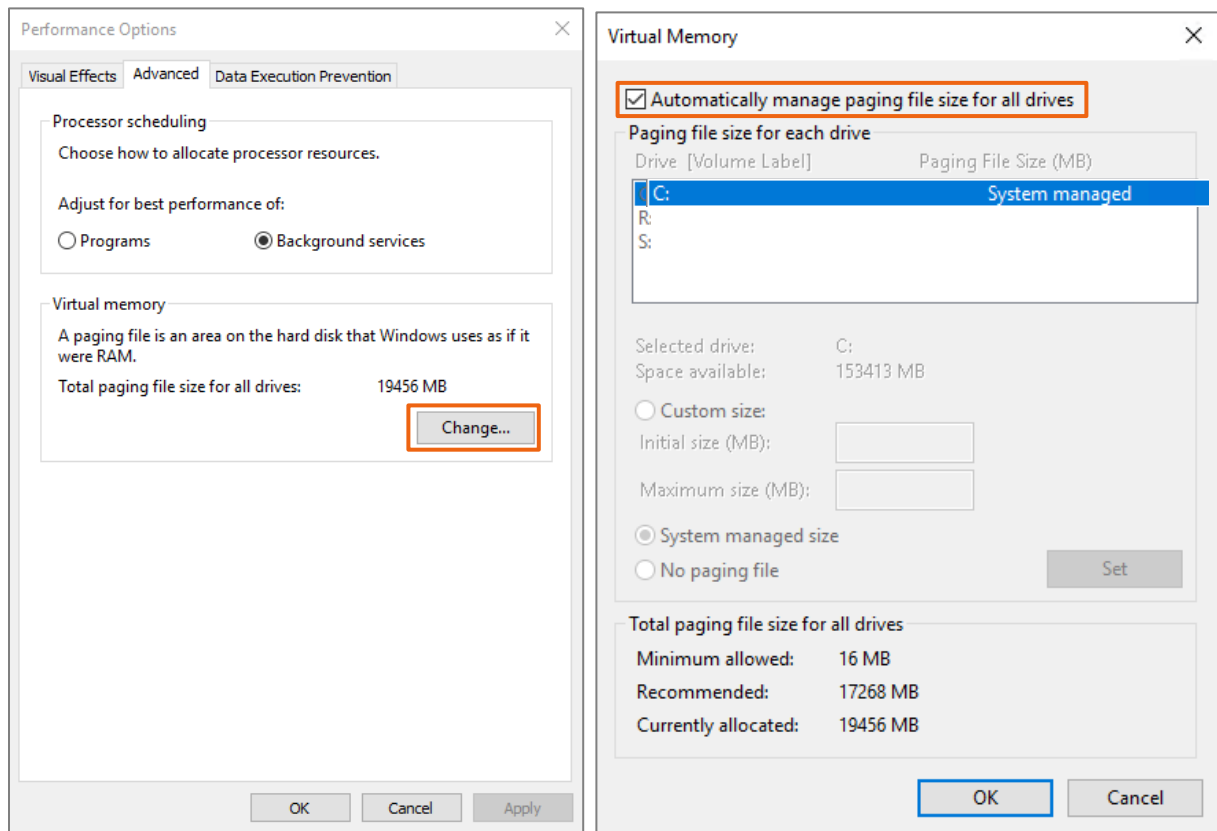


Figure 9 Default Windows Server page file settings

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 if leveraging boot-from-SAN. The following reasons may not be sufficiently advantageous by themselves to justify changing 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.

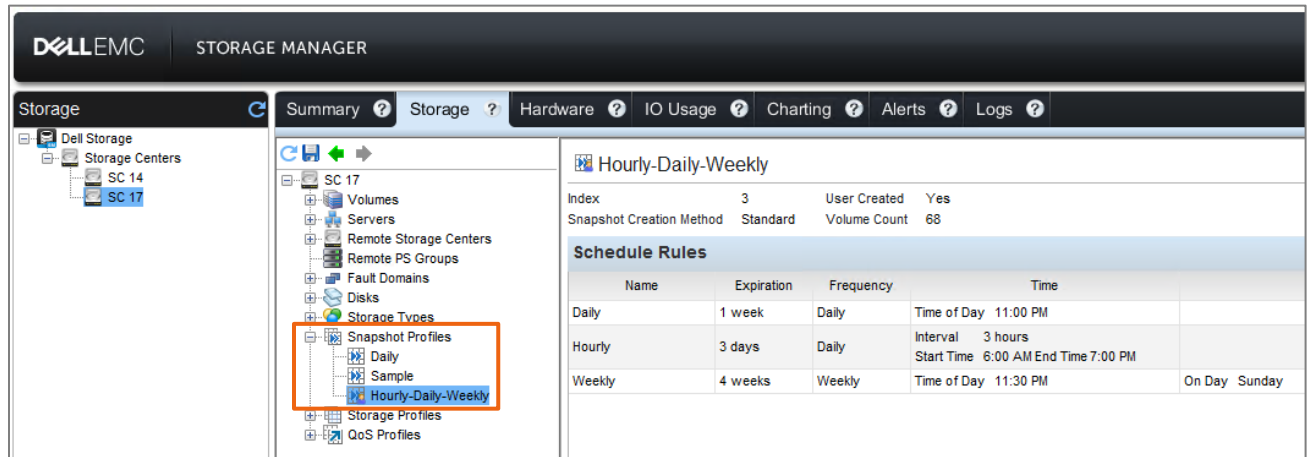
- Moving the page file to a separate dedicated SAN volume reduces the amount of data that is changing on the system (boot) volume. This can help reduce the size of SC Series snapshots of boot volumes which will conserve space in the disk pool.
- Volumes dedicated to page files typically do not require snapshot protection, and therefore do not need to be replicated to a remote site as part of a DR plan. This is especially beneficial in cases where there is limited bandwidth for replication of volumes and snapshots to another SC Series array.

3.10 SC Series snapshots with Windows Server

SC Series snapshots can be used to protect and replicate Server boot-from-SAN or data volumes and associated workloads. Snapshots are space-efficient, meaning they consist only of pointers to frozen blocks

representing only the data that has changed since the last snapshot. Active volumes will generally have larger snapshots because more data is changing.

Choose a built-in snapshot profile or create a custom profile to suit your needs. Figure 10 shows a custom **Hourly-Daily-Weekly** profile with schedule rules that expire hourly snapshots more frequently than weekly snapshots so only snapshots that are needed longer term are retained longer term.

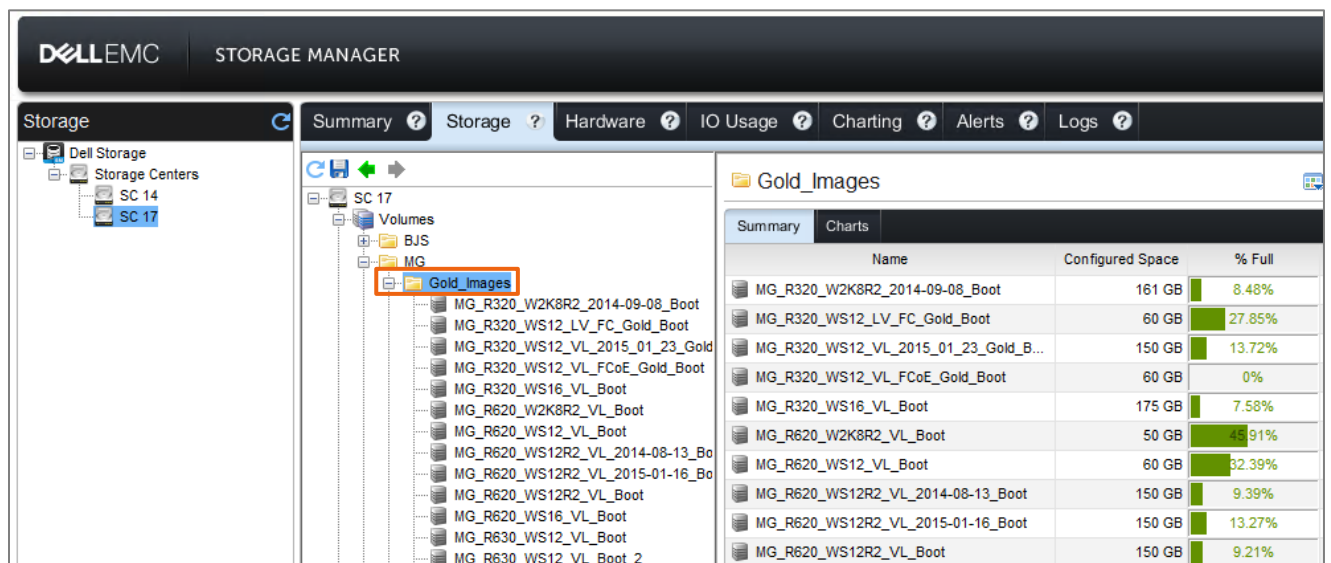


Name	Expiration	Frequency	Time
Daily	1 week	Daily	Time of Day 11:00 PM
Hourly	3 days	Daily	Interval 3 hours Start Time 6:00 AM End Time 7:00 PM
Weekly	4 weeks	Weekly	Time of Day 11:30 PM On Day Sunday

Figure 10 Built-in and custom snapshot profiles

SC Series snapshots allow administrators to do the following in Windows Server environments:

- Create View Volumes that can be used to provision an isolated test environment that matches the production environment
- Recover a server to a previous point in time
- Provision new servers from a gold image, saving admin time and conserving SAN space



Name	Configured Space	% Full
MG_R320_W2K8R2_2014-09-08_Boot	161 GB	8.48%
MG_R320_WS12_LV_FC_Gold_Boot	60 GB	27.85%
MG_R320_WS12_VL_2015_01_23_Gold_B...	150 GB	13.72%
MG_R320_WS12_VL_FCoE_Gold_Boot	60 GB	0%
MG_R320_WS16_VL_Boot	175 GB	7.58%
MG_R620_W2K8R2_VL_Boot	50 GB	48.91%
MG_R620_WS12_VL_Boot	60 GB	32.39%
MG_R620_WS12R2_VL_2014-08-13_Bo	150 GB	9.39%
MG_R620_WS12R2_VL_2015-01-16_Bo	150 GB	13.27%
MG_R620_WS12R2_VL_Boot	150 GB	9.21%

Figure 11 Create gold images to deploy servers from View Volumes for quick provisioning

SC Series snapshots can be taken of volumes mapped as LUNs to a Windows Server environment regardless of content or transport used. This applies to data volumes, cluster volumes, cluster shared volumes (CSV), pass-through disks, and in-guest iSCSI or virtual FC volumes. These volumes can also be replicated to another SC Series array for DR or other purposes.

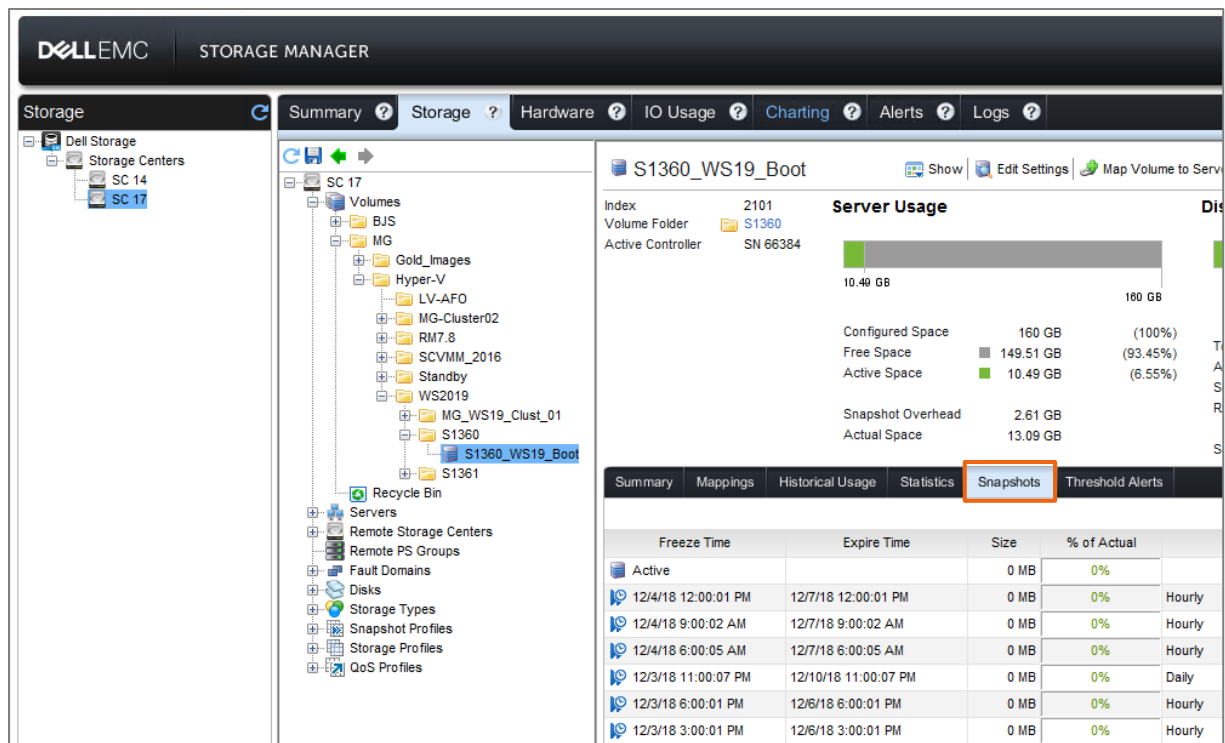


Figure 12 SC Series volume snapshots

For guidance on how to configure Windows Server and Hyper-V with Live Volume, see the document [Dell EMC SC Series Storage: Synchronous Replication and Live Volume](#).

3.10.1 Crash-consistent and application-consistent snapshots

When a snapshot is created of a volume on an SC Series array according to a schedule (or on demand) without regard to the state of the server, its applications, or data, it is assumed to be crash-consistent.

When recovering a server (or data) using a crash-consistent snapshot, it is like recovering 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 complications, whether the cause is an unexpected power outage, or the server is being recovered to a previous point in time due to an event such as a malware infection.

An exception to this is when the environment hosts a transactional workload such as Microsoft Exchange or SQL Server®. With transactional workloads, the risk of data loss or corruption is higher when attempting to recover to a crash-consistent state, especially when databases span multiple volumes.

Crash-consistent snapshots should not be relied upon for recovery when application consistency is critical for recovery of the workload.

There are several ways to ensure application-consistency. For example, an administrator can power off a server and its workload, or pause an application temporarily, before taking snapshots. However, manual processes may not be very practical and risk introducing user error. This can be mitigated in cases where automation can be used to script processes, so they can, for example, run unattended after business hours. See section 4.9 for more information on leveraging the PowerShell SDK toolkit with SC Series storage.

With SC Series, Replay Manager can be used to leverage Microsoft Volume Shadow Copy Service (VSS) to automate the process of obtaining application-consistent snapshots in Microsoft environments. For more information on Replay Manager, see section 4.7.

3.10.2 Recovery with SC Series snapshots

Server volumes, including boot-from-SAN volumes, can be recovered to a previous point by creating View Volumes from crash-consistent snapshots. Snapshots can also be used to create duplicate instances of servers or server data in an isolated environment at the same or a different location when replication between SC Series arrays is used.

3.10.3 Create a test environment with SC Series snapshots

SC Series snapshots can be used to quickly create test or development environments that mirror a production environment. When volumes are replicated to another location, this makes it very easy to do this at a different location. PowerShell integration can help automate a recurring workflow that involves View Volumes. See section 4.9 for more information on the PowerShell SDK offered with SC Series.

Note: Duplicate servers or data created from View Volumes should be completely isolated to avoid the risk of data corruption or server name/IP conflicts.

3.11 TRIM/Unmap

Windows Server 2012 and newer automatically identifies thin-provisioned LUNs and will reclaim disk space in real time when data is deleted or removed from a volume. TRIM allows this space to be automatically reclaimed by the SAN. Reclaimed space is returned to the disk pool on the SC array as free space, so it is available for other volumes. Without TRIM, this reclaimed space would show as free in Windows, but would continue to show as allocated on the SC array. TRIM is enabled by default in Windows Server 2012 and newer.

Dell EMC introduced support for TRIM/Unmap in SCOS 6.3.1. With the TRIM/Unmap feature, the DSM server agent is no longer required to recover deleted disk space from server volumes and return it to the disk pool to be used elsewhere. The server agent can still be installed on Windows Server 2012 server and newer, but the disk space recovery feature in the agent will be disabled by default because it is no longer needed.

TRIM/Unmap is supported with the following types of physical volumes and disks:

- SC Series volumes mapped to physical Windows Server 2012 and newer hosts using iSCSI, Fibre Channel, or SAS
- SC Series volumes mapped to guest VMs as pass-through or direct-attached disks using iSCSI, SAS, or virtual Fibre Channel
 - SAN volumes must be basic disks and formatted as NTFS
 - Other formats such as FAT and ReFS do not support TRIM/Unmap
 - SC Series volumes mapped to Windows Server 2012 and newer Hyper-V nodes as cluster shared volumes (CSVs)
 - CSVs must be basic disks and formatted as NTFS
 - CSVs formatted with ReFS do not support TRIM/Unmap

TRIM/Unmap is supported with the following types of virtual volumes and disks:

- The virtual hard disk format must be VHDX (dynamic or fixed)
 - TRIM/Unmap is not supported with the **VHD** virtual hard disk format
 - The guest VM OS must support TRIM/Unmap
 - When the guest VM OS is Windows Server 2012 or later, from the perspective of the guest, the VHDX must be a basic disk, formatted as NTFS
 - TRIM/Unmap is not supported on a VHDX when the guest VM OS is Windows Server 2008 R2 or earlier

3.12 Volume maintenance

Windows Server 2012 and newer automatically performs routine maintenance on volumes through the Optimize Drives application. To open the **Optimize Drives** application, navigate to **Control Panel > Administrative Tools > Defragment and Optimize Drives**. Alternately, right-click a volume, select the **Tools** tab, and click the **Optimize** button.

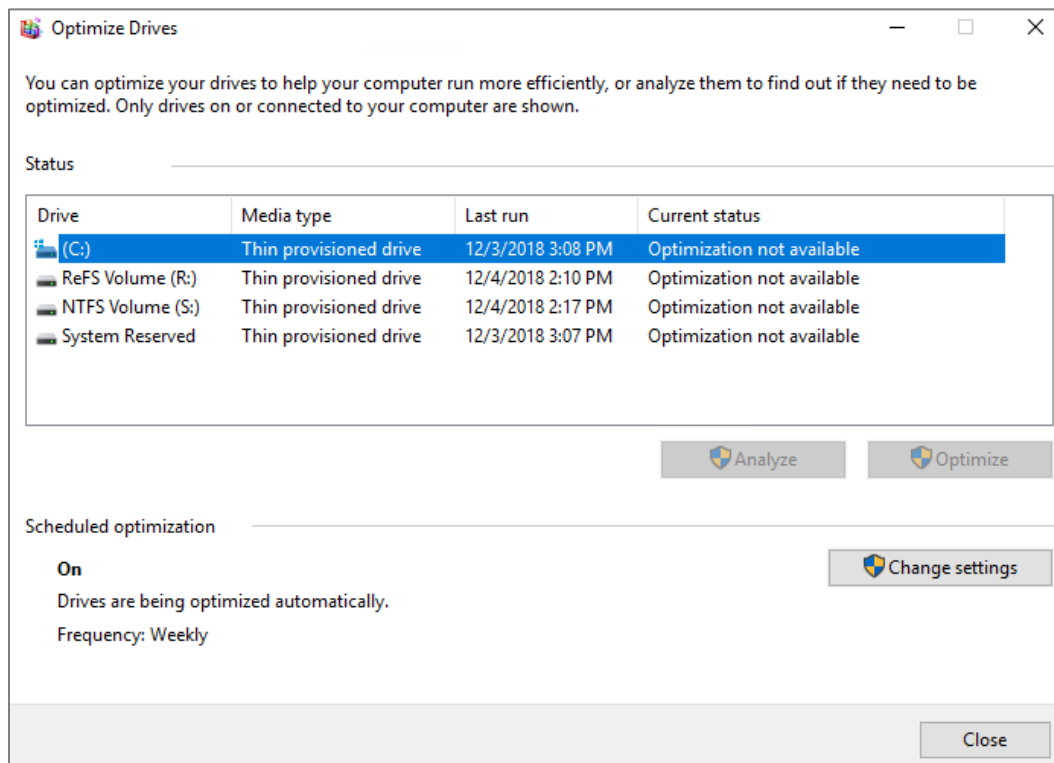


Figure 13 Optimize Drives application

By default, drive optimization runs weekly on all drives (volumes) on the system. Automatic scheduling of drive optimization can be set to run on specific volumes on a daily, weekly, or monthly basis.

On thinly-provisioned LUNs, the first thing drive optimizer attempts is slab consolidation and slab analysis (also known as defragmentation). A slab refers to an allocation of data on the underlying LUN. On SC Series storage, this is known as a data page. Windows will not run slab consolidation and analysis on thinly-provisioned LUNs with a slab size of less than 8 MB. Since the page size on an SC Series array is 512 KB, 2 MB (default), or 4 MB, slab consolidation and analysis will not occur on an SC Series LUN.

Note: As a best practice, do not run defragmentation on SC Series volumes mapped to Windows hosts because it can have a negative impact on the performance and effectiveness of Data Progression and replication.

As a last step, drive optimizer will run reTRIM on the volume. Windows Server 2012 (with SCOS 6.7 and newer) will automatically reclaim freed-up space on an SC Series NTFS volume using TRIM/Unmap. TRIM commands are processed asynchronously by the file system to minimize any performance impact on the underlying hardware. When space is freed up on a volume, the file system queues the TRIM requests to be processed. In cases where the file system receives multiple TRIM requests, the file system may reach a maximum queue size, and some TRIM requests could be dropped. Drive optimizer issues a reTRIM request on the volume to reclaim any space that was missed by dropped TRIM commands. The reTRIM request is processed in a way that limits the number of TRIM requests in the queue and avoids any dropped requests.

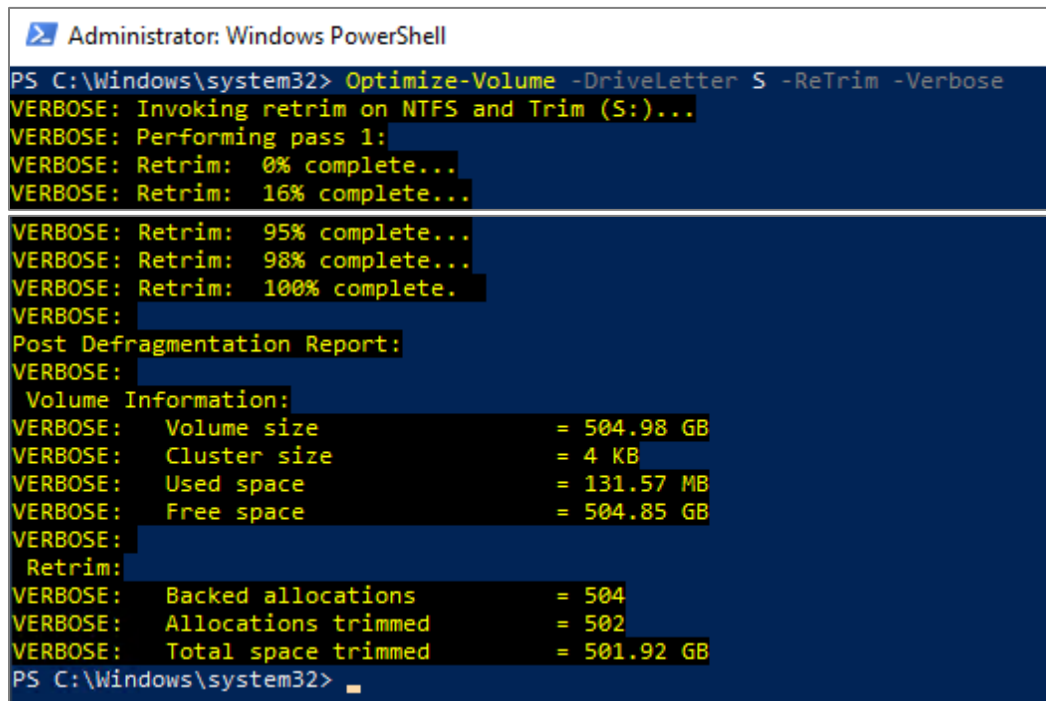
Note: TRIM and reTRIM are supported on SC Series volumes formatted as NTFS (Server 2012 and newer).

Because drive optimizer will periodically help to reclaim additional space on SC Series LUNs, leaving this feature enabled is recommended.

To run the reTRIM process manually on SC Series LUNs, use the Optimize-Volume PowerShell command. In the following example, a ReTrim command is issued on an SC Series volume that is mapped to a Windows Server 2019 host as S:\. Use the `-Verbose` switch to output detailed information to the screen while the command is running.

Example:

```
Optimize-Volume -DriveLetter S -ReTrim -Verbose
```



```
Administrator: Windows PowerShell
PS C:\Windows\system32> Optimize-Volume -DriveLetter S -ReTrim -Verbose
VERBOSE: Invoking retrim on NTFS and Trim (S:)...
VERBOSE: Performing pass 1:
VERBOSE: Retrim: 0% complete...
VERBOSE: Retrim: 16% complete...
VERBOSE: Retrim: 95% complete...
VERBOSE: Retrim: 98% complete...
VERBOSE: Retrim: 100% complete.
VERBOSE:
Post Defragmentation Report:
VERBOSE:
Volume Information:
VERBOSE: Volume size = 504.98 GB
VERBOSE: Cluster size = 4 KB
VERBOSE: Used space = 131.57 MB
VERBOSE: Free space = 504.85 GB
VERBOSE:
Retrim:
VERBOSE: Backed allocations = 504
VERBOSE: Allocations trimmed = 502
VERBOSE: Total space trimmed = 501.92 GB
PS C:\Windows\system32>
```

Figure 14 Optimize-Volume PowerShell command

3.13 Offloaded Data Transfer

Offloaded Data Transfer (ODX) is a feature developed by Microsoft that offloads data copy and move operations from Windows hosts and network switches directly to the storage array. Offloading data copy and move operations to the SAN reduces the demand for CPU and bandwidth resources on the host servers, NICs, and network switches. In addition, overall copy and move performance for data can be significantly faster with ODX.

ODX is enabled by default on Windows Server 2016 newer and is supported with SCOS 6.3 and newer. ODX should normally be left enabled, except when there is a need to obtain ODX instead of non-ODX performance benchmarks or to perform troubleshooting.

ODX requires the source and destination volumes be formatted as NTFS volumes, and be located on the same SC Series array. When volumes from different SC Series arrays are mapped to a host server, data transfers between these volumes will use standard network copy instead of ODX.

ODX operations can be initiated from a physical server or a virtual machine. The source and destination volumes can be physical disks, VHDs, or SMB shared disks (the share must be hosted on a volume located on the same SC Series array as the source/destination volume).

Within Hyper-V, ODX is used to speed up the virtualization platform layer. This allows Hyper-V to achieve native-like performance when virtual machines read and write to the SC Series array. ODX also allows for rapid deployment of guests.

Additionally, ODX can be utilized when creating a fixed-size virtual hard drive (VHD). Without ODX enabled, Windows will explicitly zero-out all the disk space assigned to the new VHD file. Depending on the size of the VHD file, this can be a slow, time-consuming process. With ODX enabled, Windows issues the SC Series array a command to write all zeros to the blocks that represent the new VHD file. This process takes seconds to complete. Windows reads the newly created VHD file as the full size, but with thin provisioning on the SC Series array, the file is not consuming any actual space until data is written to it.

ODX operations on VMs require the VMs be running Windows Server 2012 or above, or Windows 8 or above, and the VM virtual hard drive(s) must be in the VHDX format. Transferring data between VMs requires that the virtual hard drives of both VMs be housed on volumes hosted on the same SC Series array. VMs can use ODX to transfer data to other guests, physical pass-through volumes, virtual Fibre Channel volumes, and SMB shared disks.

For information on how to enable or disable ODX, and how to establish performance benchmarks, see the [Microsoft Windows IT Pro Center](#).

3.14 Resilient file system

Introduced with the initial release of Windows Server 2012, Resilient file system (ReFS) is a file system that is specifically intended for managing extremely large data volumes. Using a new file system design, ReFS can auto-detect data corruption and automatically perform needed repairs without taking a volume offline. ReFS eliminates the need to run **chkdsk** against large volumes.

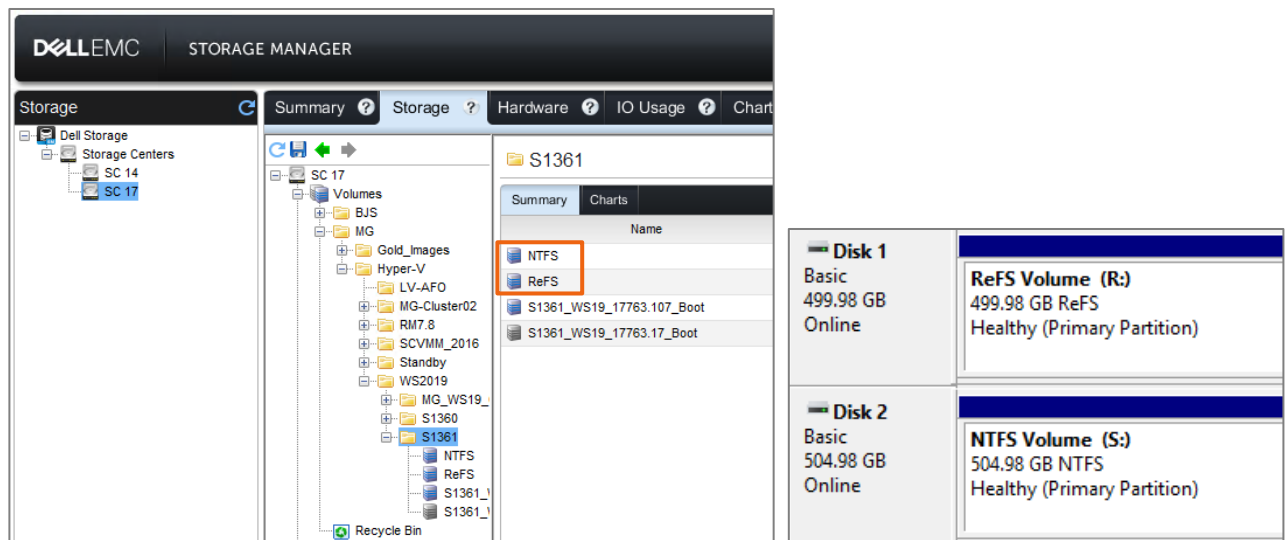


Figure 15 SC Series support for ReFS and NTFS

Microsoft recommends ReFS for very large data volumes. Because of the reduced feature set available from Microsoft with ReFS as compared to NTFS, verify the required features before choosing the format type.

Because SC Series is block-level storage, the type of format chosen for an SC Series volume from the perspective of the OS typically does not matter. ReFS and NTFS work equally well from the standpoint of snapshots, replication, and data reduction. However, ODX and TRIM/Unmap are not supported on SC Series storage with ReFS volumes. If ODX and Trim/Unmap integration are needed, then use NTFS.

For more information about ReFS and NTFS, refer to the [Microsoft Windows IT Pro Center](#).

4 Best practices for specific roles, features, and integrations

This section covers best practices for specific Windows Server roles, features, and integrations as they relate to SC Series storage.

For more information on other roles and features that are not related to storage, see the [Microsoft Windows IT Pro Center](#).

4.1 Failover Clustering

Windows Server Failover Clustering provides the capability to tie multiple servers together with shared storage to offer high availability and scalability for business-critical applications such as Microsoft Exchange, Hyper-V®, Microsoft SQL Server, and File Server. Clustering is designed to maintain data integrity and provide failover support. Windows Server 2016 and newer Failover Clustering can scale up to 64 nodes in a single cluster.

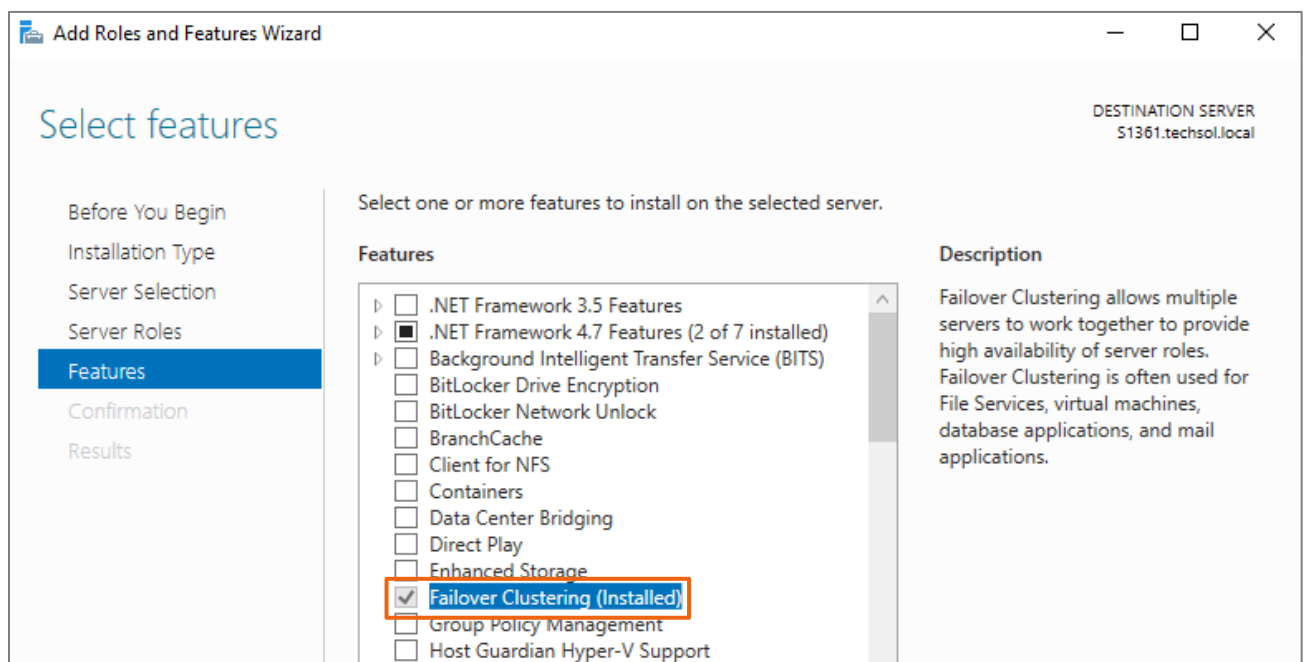


Figure 16 Windows Server Failover Clustering feature

SC Series supports Microsoft server clusters. There is no direct integration between SC Series cluster support and Microsoft Failover Clustering (or Hyper-V), but it is still a best practice to use similar naming. When creating a Windows Server Failover Cluster, also create a matching cluster server object on SC Series using the DSM client.

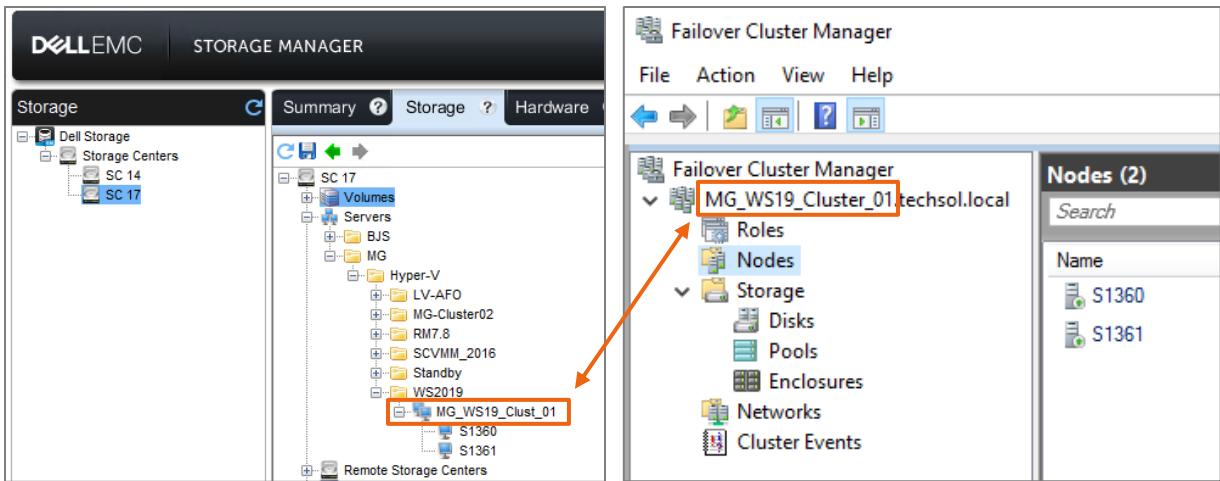


Figure 17 Matching cluster object names on SC Series and Failover Cluster Manager

In Figure 17 the cluster names are the same between the SC Series and Failover Cluster Manager which makes management easier.

When mapping a new SC Series volume, map the volume to the cluster server object using the DSM client. This will ensure uniform mapping of the volume to all nodes in the cluster with a common LUN number.

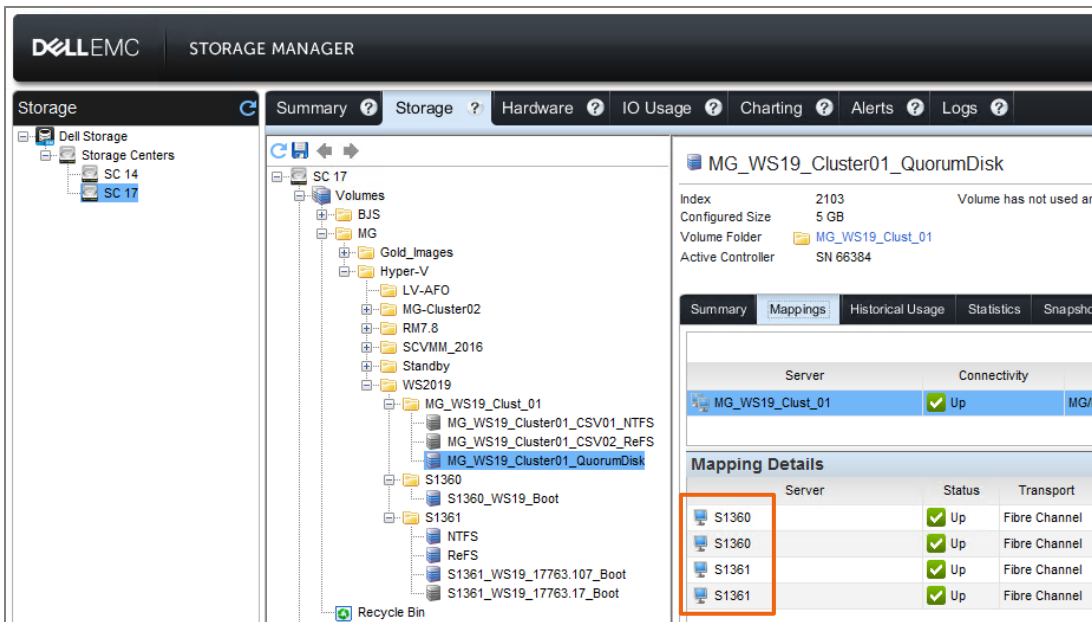


Figure 18 Volume mapped to both nodes in the cluster with MPIO support

To learn more about Windows Failover Clustering, including tools, videos, blogs, and the feature enhancements that have been made available with each new release of Windows Server, see the [Microsoft Windows IT Pro Center](#).

4.2 Microsoft Hyper-V

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

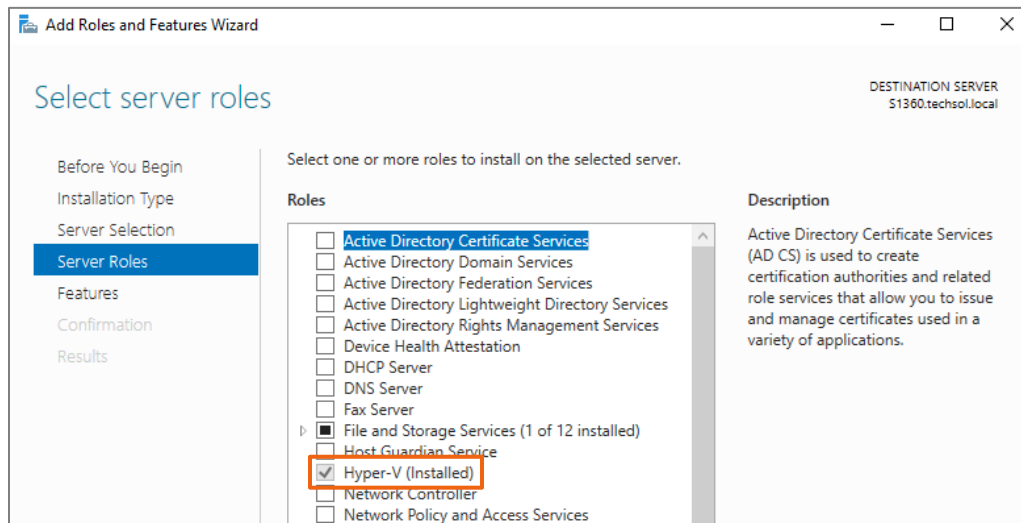


Figure 19 Hyper-V Role

Hyper-V requires Failover Clustering to also be installed on each node in the cluster.

To learn more about Hyper-V, including tools, videos, blogs, and the feature enhancements that have been made available with each new release of Hyper-V, see the [Microsoft Windows IT Pro Center](#).

For guidance with Hyper-V support with SC Series storage, see the [Dell SC Series Storage and Microsoft Hyper-V](#) best practices guide.

See also the SC Series virtualization subsection at [SC Series technical documents and videos](#) for more white papers, videos, and reference architectures.

4.3 Cluster shared volumes

Cluster shared volumes (CSVs) are supported with SC Series storage, and the same principles and best practices that apply to the creation and mapping of regular cluster volumes also apply to CSVs.

Originally introduced in Windows Server 2008 R2 Failover Clustering, CSVs allow all nodes in a cluster to simultaneously have read-write access to the same LUN that is formatted as an NTFS or ReFS volume. Using CSVs, clustered roles can fail over quickly from one node to another node without requiring a change in drive ownership or dismounting and remounting a volume.

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

For guidance with Hyper-V support with SC Series, see the [Dell SC Series Storage and Microsoft Hyper-V](#) best practices guide.

4.4 SMI-S

Microsoft System Center Virtual Machine Manager (SCVMM) 2012 and 2016 support SMI-S integration with SC Series. For more information about this integration, see the [Dell EMC SC Series and SMI-S Integration with Microsoft SCVMM](#) configuration guide.

4.5 Disaster recovery with Live Volume

Windows Server and Server Clusters including Hyper-V are supported with Live Volume, Live Volume Auto-failover, Federation, and Volume Advisor. See the [Dell EMC SC Series Storage: Synchronous Replication and Live Volume](#) guide for more information.

4.6 Storage Spaces

Microsoft Storage Spaces was introduced as a new feature in the initial release of Windows Server 2012. Storage Spaces provided storage virtualization capabilities using readily-available non-enterprise class hard disks. Storage Spaces allowed for the creation of storage pools using two or more heterogeneous physical hard drives. A Storage Space (logical drive) can be created from a storage pool, which can be formatted, partitioned, and used just as a physical disk would be. A Storage Space can be assigned redundancy, such as mirroring or parity. Storage Spaces is not a SAN replacement. Instead, it offers some SAN-like features using disparate disks in a JBOD configuration.

Storage Spaces evolved into Storage Spaces Direct (S2D) with Windows Server 2016 and newer as part of Microsoft's Windows Server software-defined strategy. Storage Spaces and Storage Spaces Direct are not compatible with SC Series storage.

4.7 SC Series Replay Manager

Dell Storage Replay Manager is a GUI-based client-server backup and recovery application that leverages the Microsoft Volume Shadow Copy Service (VSS) to obtain application consistent replays (snapshots) for specific workloads. These include Windows Server, Hyper-V guest VMs, Microsoft SQL Server, Microsoft Exchange, and VMware VMs.

To learn more about Replay Manager see the [Replay Manager Administrators Guide](#). For additional guidance with Replay Manager and Hyper-V, see the [Dell EMC SC Series Replay Manager 7 and Hyper-V](#) best practices guide.

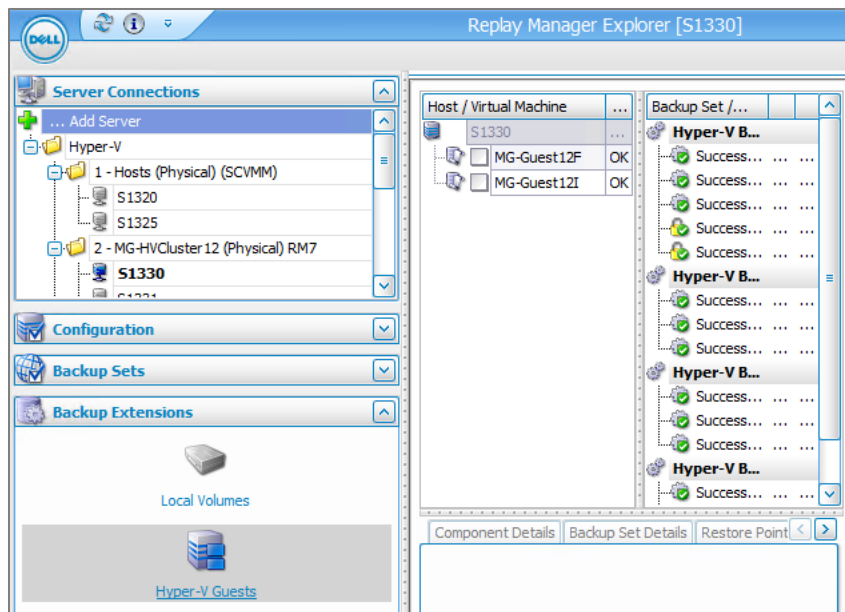


Figure 20 Replay Manager Explorer with Hyper-V Guest backup extension

4.8 Using BitLocker with SC Series volumes

Windows Server versions that support Microsoft BitLocker® have the native capability to encrypt volumes with the BitLocker drive encryption utility. With BitLocker drive encryption, users can encrypt system and data volumes to secure data residing on the volume. BitLocker makes use of the on-board Trusted Platform Module (TPM), or a USB drive if TPM is not available to encrypt the system volume. Details on deploying BitLocker can be found at the [Microsoft Windows IT Pro Center](#).

The BitLocker encryption process only encrypts the data contained on the volume. Any free space on the volume is not encrypted. All blocks that comprise the data on the volume are changed, effectively creating a new version of the data on the volume. As a result, the active snapshot on an SC Series volume will grow to be the size of the total amount of data on the volume, regardless of the state of any historical snapshots on the volume.

Keep in mind that the encryption process will result in re-writes of all the data on the volume to tier 1 by default. On large volumes, this could cause tier 1 to become full. To minimize impact to the system, verify there is enough space in tier 1 prior to applying encryption to accommodate the data in the volume. Do not let the amount of data used exceed the space available on the system, otherwise the system could enter conservation mode. If necessary, the encryption process can be prevented from consuming faster disks by isolating the volume on a lower tier of storage by changing the storage profile on the volume.

In the following example, a snapshot was taken of an SC Series volume before BitLocker encryption was applied. After the snapshot was taken, more data was added to the volume, increasing the total amount of data on the volume to 38 GB. Approximately half of the data is stored in the active snapshot, and the other half is stored in the historical snapshot.

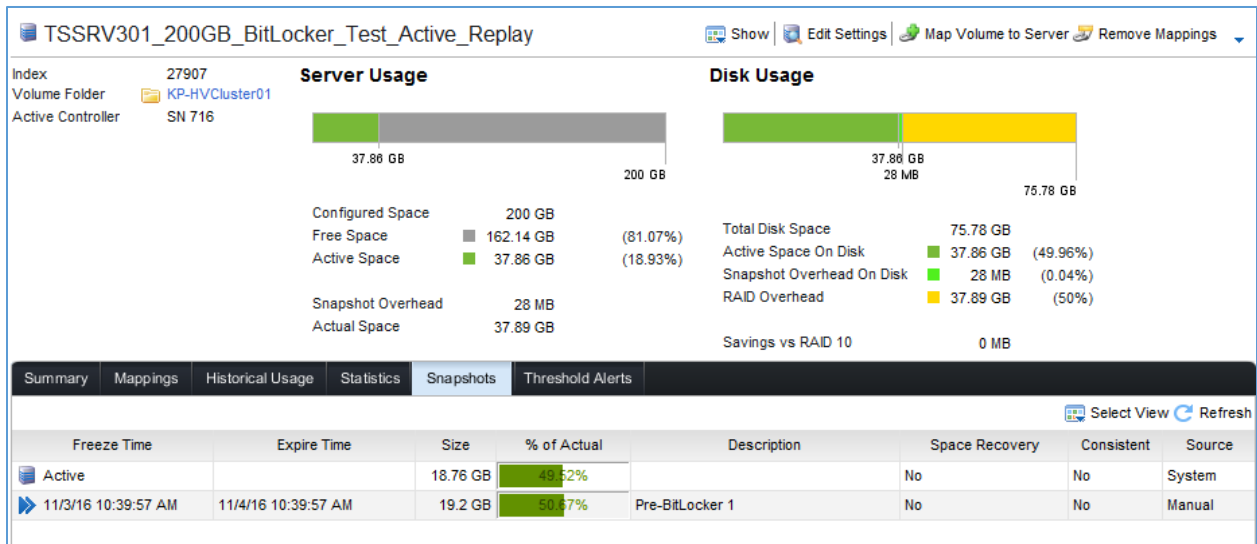


Figure 21 SC Series volume before encryption

After the BitLocker encryption process is run, the active snapshot increases to the full size of the data contained on the volume, plus an additional 500MB overhead required by BitLocker.

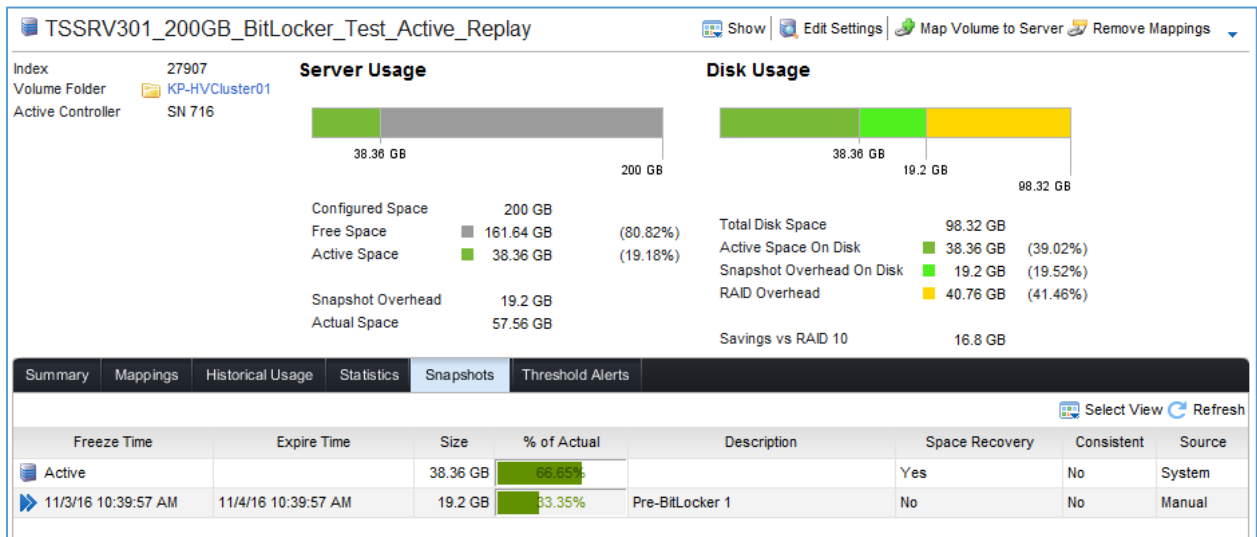


Figure 22 SC Series volume after encryption

At this point, any pre-existing snapshots on the volume are no longer valid because the data contained in those snapshots is not encrypted. Unless there is a specific need for them, it is recommended to expire any pre-existing snapshots. Once the snapshots are expired and Data Progression is run, the space used by the snapshots is returned to the SC Series volume.

Note: If BitLocker encryption is applied to an SC Series volume without any snapshots, the active snapshot size will only increase by the overhead space required by BitLocker (approximately 500 MB).

Running BitLocker on Windows Server VMs has similar results compared to running it on physical servers. Refer to the [Dell SC Series Storage and Microsoft Hyper-V](#) and [Dell EMC SC Series Best Practices with VMware vSphere 5.x-6.x](#), best practices guides for more information on configuring Windows Server VMs on each platform.

4.9 SC Series and PowerShell SDK

The SC Series PowerShell SDK command set provides many cmdlets for running storage tasks. With the SDK, administrators can create scripts that automate processes that involve both Windows Server and SC Series storage.

To learn more about the SC Series PowerShell SDK including many cmdlet examples, refer to the [Dell Storage PowerShell SDK Cookbook](#).

A Technical support and additional 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

The following SC Series publications and additional resources are available at [Dell.com/support](https://dell.com/support):

- Administrator's Guides
- Deployment Guides
- CLI Guides
- Owner's Manuals
- Hardware Compatibility List
- Support Matrix

Additionally, see the following referenced or recommended publications and articles:

Table 1 Referenced or recommended resources

Vendor	Resource
Dell EMC	Dell EMC SC Series technical documents and videos
Dell EMC	Dell EMC SC Series Storage and Microsoft Multipath I/O
Dell EMC	Dell SC Series Storage and Microsoft Hyper-V
Dell EMC	Dell EMC Series Storage and SMI-S Integration with Microsoft SCVMM
Dell EMC	Dell EMC SC Series Replay Manager 7 and Hyper-V
Dell EMC	Dell Storage Center OS 7.0 Data Reduction with Deduplication and Compression
Dell EMC	Dell Storage PowerShell SDK Cookbook
Dell EMC	Dell EMC SC Series Storage: Synchronous Replication and Live Volume
Dell EMC	Dell EMC Storage Compatibility Matrix
Microsoft	Microsoft Windows IT Pro Center
Microsoft	Microsoft TechNet
Microsoft	Microsoft Developer Network (MSDN)
QLogic	Firmware's Execution Throttle and the Windows Queue Depth explained