

Dell EMC SC Series: Performance and Availability

Dell Storage Center OS (SCOS) 7.3

Abstract

This white paper provides recommended best practice guidelines for installing and configuring Dell EMC™ SC Series systems for best performance.

November 2019

Revisions

Date	Description
December 2018	Initial release
November 2019	vVols branding update

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.

© 2018–2019 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. Published in the USA. [11/13/2019] [Best Practices] [4024-BP-SC]

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

Table of contents

Revisions.....	2
Executive summary.....	5
Audience	5
1 System configuration	6
1.1 Essential guidelines	6
1.1.1 Maximize flash capacity.....	6
1.1.2 Distribute workloads across available resources	6
1.1.3 Simplify the configuration	6
1.1.4 Design for resilience	6
1.2 Maintain the latest release of SCOS	7
1.3 SAS cabling and drive placement.....	7
1.4 Distributed spare (new in SCOS 7.3)	7
1.5 RAID spare ratio of distributed spare	7
1.6 Front-end connectivity	8
1.6.1 Fibre Channel	8
1.6.2 iSCSI.....	8
1.7 SAS front end	8
1.7.1 FCoE (SC8000 and SC9000)	9
1.7.2 NAS	9
1.8 Hardware capability guidelines.....	9
1.8.1 Flex Ports.....	9
1.9 Sizing an SC Series array for performance requirements.....	9
2 Storage configuration	11
2.1 Disk folders and volumes	11
2.2 Usable capacity	11
2.3 RAID protection	11
2.4 Dual redundancy requirements	12
2.5 Storage profiles.....	12
2.6 All-flash array.....	12
2.7 Hybrid array	13
2.8 Multi-tier system designs	13
2.9 Different capacity drives within a tier	13
3 SC Series features	14
3.1 Data Progression	14

3.2	Replication	14
3.2.1	Replication topology	14
3.3	Live Migrate and Live Volume	14
3.4	Snapshots.....	15
3.5	Deduplication and Compression.....	15
3.6	Volume Quality of Service	16
4	Application considerations.....	17
4.1	VMware.....	17
4.1.1	Virtual Volumes.....	17
4.2	Transactional applications	17
4.3	Sequential applications.....	17
5	Conclusion	18
A	Technical support and resources	19
A.1	Related resources.....	19

Executive summary

This best practices document delivers guidance to customers using Dell EMC™ SC Series storage systems in a mixed-business environment. This paper focuses on system performance and maximizing the usability of automated storage features while avoiding mismatches of technology.

These guidelines are intended to cover most use cases. They are strongly recommended by Dell EMC but are not strictly required. Some exceptions are addressed in this guide. However, less common edge cases are not covered by these general guidelines and are addressed in specific white papers. See [SC Series technical documents and videos](#) for additional documentation.

This guide presents example configuration choices to demonstrate how a certain recommendation might be applied. These examples are intended to show one possible configuration option and do not necessarily represent a preferred configuration.

If you have questions about the applicability of this document's guidelines in your environment, contact your Dell EMC representative to discuss the appropriateness of the recommendations.

Audience

This document is intended for Dell EMC customers, partners, and employees who are installing or configuring SC Series storage systems. Some familiarity with SC Series storage systems is assumed.

1 System configuration

SC Series arrays are virtualized, automated, intelligent storage systems designed with optimal combinations of flash drives and HDDs at a competitive cost. This guide introduces specific configuration recommendations that maximize the performance and efficiency of the SC Series system.

1.1 Essential guidelines

At the highest level, optimal design follows a few simple rules. The main principles of designing an SC Series system for performance and efficiency include the following:

- Maximize flash capacity
- Distribute workloads across available storage resources
- Simplify the configuration
- Design for resilience
- Maintain the latest release of SCOS

1.1.1 Maximize flash capacity

SC Series arrays can have flash configurations from 0% to 100% of the total capacity. This allows you to target specific price/performance ratios with any mix of SSDs and HDDs, then modify the mix as your needs change. You can add more flash to boost performance or add more spinning disks as cold data volumes grow. The more hot data that is in your environment, the more you will benefit with a larger amount of flash.

SC Series arrays tier storage at both the RAID and the disk level. Using flash provides the opportunity for the lowest latency with the most efficient use of capacity. An SC Series system should be designed to accommodate all the bandwidth and IOPS needs in tier 1. Capacity for this tier simply needs to accommodate the active data. Looking at historical data from thousands of SC Series systems, the space needed in tier 1 varies from 10% to 30% of the total usable space, depending on the amount of active data on the array.

1.1.2 Distribute workloads across available resources

To maximize system performance, eliminate individual bottlenecks by involving all hardware resources. This involves spreading host connections across as many front-end ports as possible, spreading disk drives across back-end SAS chains where possible, and balancing volumes across both controllers.

1.1.3 Simplify the configuration

Designing for simplicity also increases the system's flexibility, and leads to higher, more consistent performance. Standardizing on a single-capacity drive per tier, rather than using multiple drive sizes in a disk folder, can avoid performance issues due to short stripes (see section 2.9 for more details).

When considering the impact of certain configuration options, selecting the **Recommended Storage Profile** is typically the best choice. The Recommended Storage Profile takes advantage of all tiers and provide optimal performance and capacity.

1.1.4 Design for resilience

Hardware can fail, and SC Series storage is architected to continue providing storage services under such conditions. Understanding the hardware limitations of different system components can help you design a solution that provides adequate performance under failure conditions. Hardware capabilities of different components are provided later in this document.

1.2 Maintain the latest release of SCOS

Dell EMC regularly updates Storage Center OS (SCOS) to improve performance, reliability, functionality, and provide new features. Dell EMC recommends running the latest generally available version on your system.

1.3 SAS cabling and drive placement

SC Series systems are installed with resilient cabling so that each controller has two paths to every disk. In general, to achieve the best performance from a system, drives are spread across all SAS chains in the system. Depending on the performance needs for a particular design, you may require one or multiple SAS chains in the system. It is always best to spread disks across all SAS chains to maximize bandwidth.

1.4 Distributed spare (new in SCOS 7.3)

Beginning with SCOS 7.3, spare disk space is distributed across all drives in a drive folder and is designated as spare space. Data from a failed drive is moved to available spare disk space. The distributed spare feature replaces traditional hot-sparing technology used in releases prior to SCOS 7.3. All new systems deployed or upgraded with SCOS 7.3 use the distributed spare model.

Prior to SCOS 7.3, SC Series arrays required dedicated hot spares, such as full SSDs or HDDs reserved for receiving rebuild data from another drive that has failed. These spares stood idle until a drive failed in the system, and they also created a bottleneck during the restore process, since the entire failed drive had to be restored onto a single spare.

The distributed spare feature provides the following benefits over hot sparing:

- Improved I/O efficiency: All drives now participate in normal user transactions.
- Up to 500% faster rebuilds than hot sparing mode: Failed drive data is now restored in parallel to multiple drives.
- Reduced risk of data loss: Faster rebuilds reduce time during which data is unprotected.

All systems deployed with SCOS 7.3 are in distributed spare mode and do not have the option to revert to hot sparing mode. Once the system is upgraded to SCOS 7.3, the UI (Dell EMC Unisphere™ Central, Unisphere, and Dell Storage Manager) will display a message stating that the distributed spare optimizer is currently disabled. The optimizer is a background process that moves data and balances disk resources to yield the best steady-state system performance. When this is turned on for an existing system, data movement consumes back-end resources and may impact host I/O. For this reason, the optimizer is disabled by default upon upgrade to SCOS 7.3. The user may choose when to enable or disable the optimizer. Consider the current I/O activity as well as the amount of drives and capacity that needs to be balanced before enabling the optimizer. In addition, it is recommended to enable the optimizer during a maintenance window if possible.

1.5 RAID spare ratio of distributed spare

The RAID spare ratio is the number of userdata extents per spare extend in a RAID device. Much like the previous hot sparing method, the default for distributed sparing is 20 userdata extents per spare extend. Upon upgrade to SCOS 7.3, the spare ratio is calculated based on the spares and RAID devices that exist at the time of the upgrade. Existing spare ratios will be honored for upgrades to current systems.

1.6 Front-end connectivity

SC Series storage provides multiple options for front-end connectivity with on-board ports directly on the controller and optional host bus adapters (HBAs). This section discusses recommendations for the different types of connectivity.

In general, front-end ports need to be connected and configured symmetrically across the both controllers to facilitate high availability and continued connectivity in case of a controller failure.

For best performance, it is recommended to use all front-end ports that are installed in the system so that the workload is spread across as many resources as possible.

1.6.1 Fibre Channel

When configured for Fibre Channel, 8Gb and 16Gb HBAs are available for all platforms. The SC5020, SC7020, and SC9000 also support 32Gb Fibre Channel HBAs. All Fibre Channel ports can negotiate to lower speeds. Dell EMC recommends single-initiator zoning when creating zone sets (using a single server port per zone). For more information on switch configurations, refer to [Switch configuration guides for SC Series and PS Series SANs](#).

1.6.2 iSCSI

SC Series storage supports iSCSI connections on multiple 1Gb/s and 10Gb/s port options. 10GBase-T ports can auto-negotiate to 1Gb/s speeds. 10Gb/s is recommended for the best performance. Use separate networks and VLANs to segregate iSCSI traffic from normal network traffic. If possible, configure Jumbo frames (MTU 9000) on all ports in the end-to-end network to provide the best performance.

Other suggested iSCSI recommendations include the following:

- Bi-directional flow control enabled on all server and switch ports that handle iSCSI
- Unicast storm control disabled on switches handling iSCSI
- Multicast and broadcast disabled at the switch for any iSCSI VLANs
- Multicast storm control and broadcast storm control enabled when multicast and broadcast cannot be disabled
- Routing disabled between regular network and iSCSI network.
- Not using spanning tree on ports connected directly to end nodes
- Ensuring switches used for iSCSI are of a non-blocking design
- ARP Redirect enabled using QLogic® SANsurfer® for boot from SAN
- ARP Redirect enabled within ESX with hardware iSCSI HBAs to assist in controller failover
- Using 10Gb SFP+ optical or copper, BASE-T, iSCSI DCB, IPv6, or 1Gb Base-T ports

1.7 SAS front end

The SC5020, SC3020, SC4020, and SCv2020 models have the option to directly connect to servers with SAS HBAs through 12Gb SAS. SAS front-end connectivity options are limited to switchless fabrics; Dell EMC does not support SAS switches at this time. Dell EMC recommends that only validated servers and SAS HBAs be used in these configurations. For more information on validated servers and SAS HBAs, see the [Dell EMC Storage Compatibility Matrix](#).

1.7.1 FCoE (SC8000 and SC9000)

SC Series storage supports FCoE on the SC8000 and SC9000 platforms only. See the *Dell EMC Storage Compatibility Matrix* for supported FCoE HBAs, CNAs, and switches.

1.7.2 NAS

SC Series storage supports NAS file shares by attaching the Dell Compellent FS8600 NAS appliance to the SAN. For further information, refer to [FluidFS technical documents and videos](#) or contact your business partner or a Dell EMC specialist.

1.8 Hardware capability guidelines

The following sections describe the expected performance from different SC Series hardware components.

Dell EMC does not recommend continuously operating a hardware resource near its maximum potential. When utilization is below the maximum level, average response times will be better, the system will be able to handle activity bursts without becoming overloaded, and performance may still be maintained during hardware failure scenarios.

For some of the following hardware sections, both IOPS and bandwidth (MB/s) capabilities are shown. Different types of workloads are assumed when determining IOPS capabilities compared to bandwidth capabilities:

- IOPS workloads assume random small-block (<64KB) access containing both read and write operations.
- Bandwidth workloads assume sequential large-block (>64KB) access containing both read and write operations.

This data is a guideline for what most customers can expect as the maximum performance with certain hardware. Note that individual components in your specific system may be able to provide higher or lower numbers, based on differences in workload characteristics and environments.

Most hardware components in the SC Series system can be monitored in Dell Storage Manager using the Charting feature.

1.8.1 Flex Ports

The Flex Port feature on the SCv2000, SCv3000, SC4000, and SC5000 Series arrays allows using a single physical port to function as both a management port and a front-end iSCSI data port. For more information, see the document [SC Series Flex Port Functionality Reference Architecture](#).

1.9 Sizing an SC Series array for performance requirements

An SC Series array can be comprised of any combination of disk drives. The array can be 100% spinning disks, 100% SSDs, or any mix of the two. Typically, I/O operations per second (IOPS) are determined to size a system unless the application is streaming data such as video where bandwidth is used for sizing an array. When using spinning drives, there is a near-constant IOPS capability per drive that is used in determining the design. When designing with SSDs, the IOPS capability of a single drive is very high, but the bandwidth of the drive needs to be considered to ensure that there is enough to meet the needs of the design.

Example: An array needs to provide 20,000 IOPS at a 70/30 read/write mix with an average block size of 32K. The target for designing the array would be 26,000 (14,000 reads + [6000 writes * 2]) IOPS to accommodate the RAID write penalty and future growth. A 15K rpm spinning drive can do approximately 200 IOPS each with a sustained bandwidth of about 200 megabytes per second (MB/s). To meet the 26K IOPS design, you would need 130 x 15K drives for this array. SSDs can provide 100 times or more IOPS, but only have five times the bandwidth of spinning drives. Therefore, the design must be cognizant of the capabilities of the disk drives being used.

Dell EMC and its partners can design the appropriate SC Series configurations to meet your applications' performance and capacity requirements. Contact your Dell EMC representative for assistance with sizing an appropriate solution.

2 Storage configuration

This section provides general recommendations for optimizing the SC Series array to meet your specific needs.

2.1 Disk folders and volumes

In general, it is recommended to use a single disk folder in SC Series arrays because this reduces complexity and increases flexibility. However, it may be appropriate to configure multiple disk folders in some cases:

- Running different workloads (such as video and database applications) with different I/O profiles
- Dedicating resources to meet specific performance goals (for example, using SSD drives for a high-performance application and HDDs for file editing)
- Having separate resources for multi-tenancy (for example, LUNs with VDI user data)

2.2 Usable capacity

SC Series storage allows for multiple RAID types to exist on the same drives. Therefore, calculating the usable capacity on a system is not an exact science because there are many factors to consider, the main ones consisting of the following:

- Storage profiles
- Use of snapshots
- Length of time snapshots are kept
- Amount of data written between snapshots

Analyzing thousands of customers with RAID 10 and RAID 5 in tier 1, the usable capacity averages about 77% efficiency. To better understand how these affect capacity, discuss with your Dell EMC storage specialist or business partner.

2.3 RAID protection

SC Series storage supports multiple levels of RAID-protection schemes to protect user data against drive failure. Disks are associated by tiers and redundancy (for example, tier 1 single redundancy or tier 3 dual redundancy). The redundancy choice defines the types of mirror or parity RAID used for those disks. Single redundancy includes a single drive for protection, such as RAID 5 or RAID 10. Dual redundancy has two drives for protection, such as RAID 6 or RAID 10 dual mirror (RAID 10 DM).

Choose the RAID type that best suits your needs for performance, protection, and cost:

- **RAID 10** provides the highest level of performance from a given set of drive resources, with the lowest CPU requirements. However, it is only 50% efficient in usable drive capacity.
- **RAID 10 DM** uses an additional disk as a mirror and provides better availability than RAID 10 and has a usable capacity of 33%.
- **RAID 5** provides the best usable capacity from a set of drive resources, but at lower overall performance and availability than RAID 10. In terms of capacity, RAID 5-5 is 80% efficient and RAID 5-9 is 89% efficient.
- **RAID 6** provides better availability than RAID 5 and better usable capacity than RAID 10 and RAID 10 DM, but has the lowest performance potential of the RAID types. RAID 6-6 is 67% efficient and RAID 6-10 is 80% efficient.

The disk tier redundancy should be chosen prior to setting up the SC Series system. If you need help determining which type best suits a particular need, contact your Dell EMC storage specialist or business partner. Note that larger drives will default to dual redundancy by the system, and in some case may be forced to use dual redundancy.

2.4 Dual redundancy requirements

Dual redundancy in a tier allows for a dual drive failure within that tier without data loss. Redundancy requirements for each tier are based on the size of the disk drive in the tier, and whether the disks are added to an existing disk folder or a new disk folder. During installation, the desired redundancy level can be selected. Dual redundancy for larger HDD drives will be required. The requirements and recommendations for dual redundancy are shown in Table 1.

Table 1 Dual redundancy requirements based on capacity

Disk type	Recommended capacity	New storage type (required)	Existing storage type (required)
HDD	966 GB to 2.4 TB	2.5 TB or larger	2.79 TB or larger
SSD	1.7 TB or larger	N/A	4.0 TB or larger

For new SC Series array installations, HDD disks of 2.5 TB and above must use a dual redundant RAID level (RAID 10 DM or RAID 6). For existing arrays with 2.79 TB or larger HDDs or 4.0 TB SSD disks, dual redundancy is also required.

Note: For SSDs, these rules may change over time as technology advances. These rules apply to SCOS 7.3 at the current publication date. Refer to the release notes for the appropriate SCOS version.

2.5 Storage profiles

SC Series arrays use **storage profiles** to determine which tier and RAID type that volume data is written to. The storage profile also defines where data will exist after a snapshot is taken and as it ages. Generally, the default storage profile is recommended for volumes. However, there are reasons to use the other available storage profiles, or even create a custom storage profile for volumes that are to be handled differently.

Example: For a volume that needs to remain in tier 1 storage for performance reasons, you can use the **High priority** profile. This profile keeps all volume data in tier 1. New writes will go to mirror RAID and frozen data contained in a snapshot will be progressed to parity RAID.

2.6 All-flash array

All-flash arrays provide the highest level of performance in the SC Series family. Use an all-flash array when the application requires the highest storage performance at the lowest response time. Dell EMC recommends disabling write cache for any volumes that will always be writing data directly into a flash tier. If the flash tier becomes full and writes are redirected to an HDD tier, there will be a performance reduction on those volumes, and in that case, write cache should be enabled on the volume.

Deduplication is only supported on SC Series systems with a minimum number of flash drives attached (see the release notes for the appropriate SCOS release).

Dell EMC recommends that all drives of a similar type be of the same size.

Example: In an all-flash system, use only 1.92 TB flash drives.

2.7 Hybrid array

Hybrid arrays can contain more than one type of HDD (SAS or NL-SAS) as well as flash drives in the same disk folder. Hybrid arrays typically provide greater capacity at a lower cost than all-flash arrays, but also have lower performance and higher response times. Use of hybrid arrays can greatly improve overall performance if properly sized and configured. Consult with your Dell EMC storage specialist or business partner for appropriate implementation.

Dell EMC recommends having enough flash capacity to accommodate 20% to 30% of the data in the volumes that use the flash tier. Storage profiles allow you to select which volumes will use the flash tier and which volumes will only use the HDD tiers.

2.8 Multi-tier system designs

When designing a system with multiple tiers of disk drives, it is a best practice to make tier 1 capable of holding 20% to 30% of the data and provide 100% of the system IOPS needed. Lower tiers should be able to provide enough capacity for 70% to 80% of the data and 20% of the system read I/Os.

Systems with three tiers of disks are not very common in modern designs. However, there are reasons that it makes sense to implement three tiers. One example may be when expanding an existing system that already has two types of drives, and the customer is looking to add high performance SSDs or high capacity 7K drives to their system. Another example is when the customer has one or two applications that they require the highest performance but do not have the budget to run their whole environment on SSDs. In this case, the customer would use a custom storage profile to allow certain volumes to span different tiers of disk. Typically, Dell EMC recommends spanning only two tiers for a volume. As an example, allow the volumes for the high-performance applications to span tiers 1 and 2 (or maybe 1 and 3), while using a storage profile for the rest of the volumes to span tiers 2 and 3. This would deliver two different I/O performance characteristics depending on the storage profile in use on the particular volume.

2.9 Different capacity drives within a tier

When a customer adds new drives to an existing system, they usually add the same class and capacity of drives. On occasion, drives of the same capacity are no longer available due to increasing capacity points by the manufacturer. In these cases, care should be given to prevent a performance issue. This type of issue can arise when adding only a few drives to larger disk folder.

Example: An existing system runs at an average of 2,000 IOPS and has 100 x 1 TB 7,200 rpm drives. To expand the storage capacity, 12 x 2 TB 7,200 rpm drives are added into the disk folder. A rebalance and restripe is automatically run and the system performs as expected. However, once the data has expanded beyond the original 1 TB of each drive in the disk folder, all additional space is now striped only across the 12 new 2 TB drives. This space will perform only as many IOPS that the 12 drives allow, which is below the original 2,000 IOPS required.

This performance issue could be avoided by adding enough larger drives into the disk folder that would satisfy the needed IOPS once the larger drives are servicing the used additional capacity. One way to address this issue is to use drives that are closest in capacity to the existing drives in the disk folder, rather than use the largest capacity drives available. This would provide more spindles for the same additional capacity.

3 SC Series features

This section describes the different SC Series features and how they can be configured to optimize SC Series storage performance.

3.1 Data Progression

Data Progression is the data movement engine in the SC Series array that converts data to a more efficient RAID type when a snapshot is taken. It also moves data to an appropriate tier of storage based upon the data's access frequency in conjunction with the storage profile of the volume.

3.2 Replication

SC Series arrays support two types of replication: synchronous and asynchronous.

Asynchronous replication: The system acknowledges a server write I/O as soon as the source system can guarantee the write will be committed to disk. Then, the write is replicated to the target system at a later point in time.

Synchronous replication: The server write I/O is not acknowledged until both the source and target systems can guarantee the write will be committed to disk. There are two synchronous replication modes available:

- **High Availability Mode:** This accepts I/O requests to the primary volume when the secondary volume is unavailable (or when latency is too high) to avoid interrupting service. However, if writes are accepted to the primary volume, the secondary volume data becomes stale.
- **High Consistency Mode:** This prevents I/O requests to the primary volume when the secondary volume is unavailable to make sure that the volumes remain identical. The primary volume cannot be modified during this time, which can interrupt operations.

The SCv20x0 platform only supports asynchronous replication.

3.2.1 Replication topology

Mixed topology allows the source volume to be replicated to two destination volumes where one is synchronous and the additional replications are asynchronous (1 to N asynchronous destinations where N is defined by SCOS, currently 2).

Cascade topology allows asynchronous replications to be chained to synchronous or asynchronous replication destination volumes.

Hybrid topology is created by combining mixed and cascade topology types.

3.3 Live Migrate and Live Volume

Live Migrate and Live Volume expand on the replication capabilities of the SC series system.

Live Migrate is included with the purchase of an SC Series array and allows for volume movement with no down time to another SC Series array. No additional license is needed for Live Migrate, which is available with SCOS 7.1 or higher.

Live Migrate leverages replication. It maps the volume to both instances of the volume, although the server sees both volumes as one volume, but with multiple paths, and will send reads and writes down any path. In

In addition, the included Volume Advisor helps automate the process. The combination of this feature and the Volume Advisor allows data to move freely and transparently across local, campus, metro, or geographically dispersed distances, without interrupting workloads or reconfiguring hosts.

Live Volume lets you perform maintenance on an active system without affecting production workloads, replicating new writes automatically as you go, and swapping primary paths manually as needed, or automatically. You can also take impromptu disaster-avoidance precautions to quickly prepare for potential outages, such as hurricanes, without predefined disaster recovery mappings.

3.4 Snapshots

Snapshots allow the creation of local recovery points for any volume or group of volumes. These snapshots do not take any additional space, but freeze the pages of a volume as read only. Any new writes or updates to frozen pages will take up additional space because the frozen pages cannot be modified. These snapshots can be automatically expired after a user-selected time period has passed.

Snapshots may increase the overall bandwidth needed on the SAS backend, which is due to the backfilling of the frozen data when updates are made to a frozen page. Consider the overhead of snapshots when planning both performance and capacity requirements for the storage system.

Dell EMC recommends the practice of staggering volume snapshot operations (such as creation or deletion), to minimize the bandwidth load on the SAS backend. This can be accomplished by using different snapshot schedules for different sets of volumes.

Snapshots are deleted by the system asynchronously; when a snapshot is in the process of being deleted, it is marked as expired. The system then coalesces the pages into the unexpired snapshot to ensure that any unexpired snapshot can be recovered.

3.5 Deduplication and Compression

SC Series compression is available with SCOS 7.0 or later.

Supported SC Series volumes have the option to be compressed, or deduplicated and compressed. Data reduction is not supported on the SCv20x0 models, and SCv30x0 models only have compression available. The data reduction process is executed during On-Demand Data Progression (initiated by a snapshot) or during the daily Data Progression cycle. The data reduction process is run by the system and may have an overall effect on available resources.

An individual volume's data not contained in a snapshot will not be compressed; the compression works on the snapshot data only and may impact the snapshot access (accessible frozen pages).

Data reduction occurs on the lowest tier of the media type (lowest tier of flash or spinning disk).

Rules for deduplication and compression (for systems with flash) include the following:

- Tier 1 only (flash): Deduplication and compression
- Tier 1 flash, tier 3 HDD: Deduplication and compression on both tiers
- Tier 1 write-intensive flash, tier 2 read-intensive flash: Deduplication and compression on tier 2 only
- Tier 1 write-intensive flash, tier 2 read-intensive flash, tier 3 HDD: Deduplication and compression on tier 2 and tier 3 only
- Tier 1 flash, tier 2 HDD, tier 3 HDD: Deduplication and compression on tier 1 and tier 3 only

Rules for compression (HDD-only systems) include the following:

- Tier 1, tier 2, tier 3: Compression on tier 3 only
- Tier 1, tier 3: Compression on tier 3 only

In a three-tier system, only tier 1 and tier 3 will be deduplicated and compressed because the amount of data-reduction domains may exceed the memory requirements needed for a failover.

For volumes that have data reduction turned on, the snapshot schedules should be staggered to avoid putting too much pressure on the system resources, such as the SAS chain back end. In addition, schedules of groups of volumes should allow for each set to complete the snapshot processing (for instance, spanning a couple of hours between sets for high-change-rate volumes).

3.6 Volume Quality of Service

Dell EMC recommends using the volume Quality of Service feature on workloads which might monopolize SC Series resources and starve other applications of their required performance.

Example: Limit the bandwidth available to large-block applications which may be increasing the latency on other small-block workloads.

For more information, see [Volume Quality of Service Best Practices with Dell SC Series Storage](#).

4 Application considerations

This section presents best practices for optimizing SC Series storage with certain applications.

4.1 VMware

Dell EMC has a robust set of best practices for SC Series systems when implementing VMware® applications. Find the latest resources at [SC Series technical documents and videos](#).

4.1.1 Virtual Volumes

SC Series storage supports VMware® vSphere® Virtual Volumes™ (vVols) through Dell Storage Manager. For vVol best practices, refer to the document *Dell EMC SC Series and VMware vSphere Virtual Volumes Best Practices* on [SC Series technical documents and videos](#)

4.2 Transactional applications

SC Series systems require high concurrency to deliver the maximum performance (IOPS). Therefore, it is a best practice to set HBAs to a large queue depth initially. For systems that are configured with many hosts or volumes, host HBA settings may need to be adjusted to throttle a server that is generating a significant amount of I/O that affects the performance of other servers. Consult the documentation for your OS or HBA to adjust queue depth settings.

4.3 Sequential applications

For workloads which require high bandwidth for sequential streaming data, use all front-end ports available. A system fully populated with front-end and back-end ports allows for maximum bandwidth to be achieved. It may also be beneficial to fully provision the volumes before using them. This fully allocates the capacity before application use, in a consistent manner, which can improve subsequent sequential access.

It is not recommended to fully allocate volumes that will be used with snapshots or replication, which can lead to inefficient use of capacity within the system.

Thick provisioning can be achieved by preallocating the volume in Dell Storage Manager.

5 Conclusion

This best practices guide provides configuration and usage recommendations for SC Series storage systems in general use cases.

For a detailed discussion of the reasoning or methodology behind these recommendations, or for additional guidance around specific use cases, see the additional resources in appendix A, or contact your Dell EMC representative.

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

For additional information regarding any of the topics covered in this white paper, or the areas listed below, refer to the following resources available on [SC Series technical documents and videos](#).

- Dell EMC SC Series: Data Protection
- Dell EMC SC Series: Databases
- Dell EMC SC Series: Dell Storage Manager
- Dell EMC SC Series: Exchange Server
- Dell EMC SC Series: FluidFS/NAS
- Dell EMC SC Series: Linux/Unix
- Dell EMC SC Series: Oracle
- Dell EMC SC Series: SQL Server
- Dell EMC SC Series: SharePoint
- Dell EMC SC Series: Storage Center
- Dell EMC SC Series: Virtualization
- Dell EMC SC Series: Windows Server