

Understanding RAID with Dell SC Series Storage

Dell Engineering
September 2016

Revisions

Date	Description
February 2016	Initial release
September 2016	Added changes for redundancy for SCOS 7, 7.1

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2016 Dell Inc. All rights reserved. Dell and the Dell EMC logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Table of contents

Revisions.....	2
Executive summary.....	5
Acknowledgements.....	5
1 SC Series storage basics.....	6
1.1 Tiers.....	6
1.2 RAID redundancy.....	6
1.3 RAID level and tiers.....	7
1.4 RAID levels and redundancy.....	8
1.5 Dual redundancy requirements.....	8
2 RAID rebuild, availability, and efficiency comparison.....	10
2.1 SC Series RAID-level disk failure protection and reliability.....	10
2.2 Space efficiency by RAID level.....	11
2.3 SC Series RAID comparisons.....	11
3 Spare disks.....	13
3.1 No spare available.....	13
3.2 Disk failure.....	13
3.3 Replaced disk.....	13
4 SC Series RAID level use.....	15
4.1 SC Series disk and RAID layout.....	15
4.2 Application data writes.....	16
4.3 Highest performing RAID write advantage.....	17
5 Snapshots and Data Progression.....	19
5.1 Data Progression.....	19
5.2 On-Demand Data Progression.....	19
5.3 Data Progression and storage profiles.....	19
5.4 Capacity within tiers.....	20
6 RAID management.....	23
6.1 RAID rebalancing considerations.....	23
6.2 Modifying tier redundancy.....	23
6.3 Creating a storage type.....	26
6.4 Creating a storage profile.....	30
7 Design considerations.....	32

7.1	SSD considerations	32
7.2	SSDs in Tier 1 dual redundancy.....	32
7.3	Techniques for specific workloads.....	32
8	Summary	34
A	Additional resources.....	35

Executive summary

Dell™ SC Series storage is designed to take advantage of redundant components, including RAID, in order to achieve the greatest level of performance and availability. Core to the SC Series Fluid Data™ Architecture is the automatic placement of data across the available disks. The disks are logically grouped by type based on speed and capacity. RAID protection is built on a page architecture that provides data optimization for availability, capacity, and performance. SC Series arrays automate the most efficient data placement. From a performance and availability perspective, all applications write to the highest performing RAID level and disks. As time passes, data with lower activity moves to capacity efficient RAID levels and disk types.

This document provides a general understanding of how RAID is used in SC Series arrays, RAID levels and disk tier usage, as well as large SSD considerations and appropriate selection of data protection based on performance, availability, and efficiency.

Acknowledgements

Authored by: Chuck Farah

Illustrated and edited by: Camille Reaves

1 SC Series storage basics

SC Series arrays include redundant, hot-swappable components (for example, physical disks, control modules, fans, and power supplies) for a no-single-point-of-failure configuration. Along with redundant hardware, several different RAID levels are supported with each configuration optimized to maximize performance, availability, and capacity of the SC Series architecture.

1.1 Tiers

When configured, all SC Series disks act as a single pool of storage virtualizing the RAID configurations, disk types, and speed. The disks are grouped by speed into as many as three tiers. Each tier is determined by the speed of the disk from the fastest (Tier 1) to slowest disk (Tier 3).

Table 1 shows each disk type and the typical tiers assigned to those disk based on speed. However, many variations of tiering are possible and not all combinations are included.

Table 1 Example disk types and typical associated tiers

Disk type	Typical tier	Notes
SSD–write intensive (SSD-WI)	Tier 1	Provides best performance
SSD–read intensive (SSD-RI)	Tier 2	Provides good read and write performance. Use in Tier 1 in absence of SSD-WI.
SSD (other types)	Tier 1	Use when SSD-WI are absent. Otherwise, use in any tier.
15K HDD	Tier 1	Use when SSDs are absent. Otherwise, use in Tier 2 or Tier 3.
10K HDD	Tier 2	May also be used in Tier 1 if faster HDDs or SSDs are not available. Tier 3 is also appropriate in some cases.
7.2K HDD	Tier 3	Only used in Tier 1 when these are the only disk types available.

In cases where more than three speeds of disks are present, the third tier contains the lowest-speed disks. These disk types are just an example of the disks available today and may evolve over time.

Note: RAID tiering for the Dell Storage SCv2000 Series controllers moves data between RAID levels within a tier. No movement of data occurs between tiers on different disks.

1.2 RAID redundancy

In most configurations, all disks form a single pool of storage in a managed disk folder. Storage Center OS software defines the RAID protection for the disk pool. RAID redundancy levels provide fault tolerance for a disk failure. Redundancy options may be restricted depending on the disk size.

The RAID redundancy options are:

- Non-redundant: Uses RAID 0 in all classes, in all tiers. Data is striped but provides no redundancy. If one disk fails, all data is lost. Do not use non-redundant storage for a volume unless the data has been backed up elsewhere.

Note: The non-redundant option is not available for the SCv2000 series.

- Single-redundant: Protects against the loss of any one drive. Single-redundant tiers can contain any of the following types of RAID storage:
 - RAID 10 (each disk is mirrored)
 - RAID 5-5 (striped across 5 disks)
 - RAID 5-9 (striped across 9 disks)
- Dual-redundant: Protects against the loss of any two disks. HDDs larger than 1.9TB should use dual redundancy, and in some cases it is mandated. Dual-redundant tiers can contain any of the following types of RAID storage:
 - RAID 10 dual mirror (data is written simultaneously to three separate disks)
 - RAID 6-6 (4 data segments, 2 parity segments for each stripe)
 - RAID 6-10 (8 data segments, 2 parity segments for each stripe)

1.3 RAID level and tiers

For SC Series arrays, RAID 10, RAID 10-DM, RAID 6-6, RAID 6-10, RAID 5-5, and RAID 5-9 are available. Table 2 shows the available RAID levels, description, and the typical tiers used. The random write performance rating indicates the best performing RAID with 1 being the highest.

Note: SC Series Fast Track is a licensed feature that improves performance by consolidating frequently accessed blocks of data to the outer tracks.

Table 2 SC Series RAID levels

RAID levels	Description best practice	Typical tier	Random write performance rating
RAID 10 Standard ¹ Fast ²	RAID 10 provides the highest level of random write performance.	Tier 1	1
RAID 10 – DM (dual mirror) Standard ¹ Fast ²	RAID 10 with a third mirror for excellent efficiency and excellent random write performance.	Tier 3	2
RAID 5-5 RAID 5-9 Standard ¹ Fast ²	Not recommended for HDDs larger than 967 GB which should use RAID 6 or RAID 10 – DM. Good overall performance and efficiency.	Tier 1, Tier 2, Tier 3	3
RAID 6-6 RAID 6-10 Standard ¹ Fast ²	Offers best protection; Tier 3 will be written to by the system in most configurations.	Tier 3	4

¹ Standard track RAID level are tracks not allocated to Fast Track.

² Fast Track RAID level applies to HDDs with the Fast Track license; does not apply to SSDs.

RAID 5 is not preferred for arrays with 967 GB or larger HDDs. RAID 6 and RAID 10-DM offer significantly higher levels of resiliency with very high capacity disks.

1.4 RAID levels and redundancy

A storage type is a pool of storage with a single data page size and a specified redundancy level. SC Series arrays use storage types to logically group disks into tiers. In typical SC Series environments, a single storage type has all tiers of disks assigned. By default, data is migrated between tiers and RAID levels in 2 MB blocks (data pages). Data can be moved in smaller or larger blocks to meet specific application requirements.

Because Dell Enterprise Manager (now Dell Storage Manager), enables redundancy selection for the storage type at the time of installation or addition of new disks, the RAID levels are automatically allocated according to either single or dual redundancy settings.

Figure 1 shows a typical two-tier SC Series array with single redundancy configured in Tier 1 and dual redundancy configured in Tier 3.

Name	Track	Used	Free	Allocated	% Full	Growth	Est. Full
Tier 1 Storage [Non Allocated: 1.4 TB] [Disks: Write-Intensive SSD] [Redundant]							
RAID 10	Standard	158.12 GB	233.85 GB	391.97 GB	40%	-630 KB	
RAID 5-5	Standard	13.87 GB	24.54 GB	38.42 GB	36%	0 MB	
		172 GB	258.39 GB	430.39 GB	39%	-630 KB	
Tier 3 Storage [Non Allocated: 7.41 TB] [Disks: 7K] [Dual Redundant]							
RAID 10-DM	Fast	0 MB	582.52 GB	582.52 GB	0%	945 KB	Dec 13, 3008
RAID 10-DM	Standard	0 MB	1.44 TB	1.44 TB	0%	118.62 MB	Oct 19, 2020
RAID 6-10	Fast	209.09 GB	2.07 TB	2.27 TB	8%	338.31 MB	Nov 15, 2027
RAID 6-10	Standard	3.03 TB	759.41 GB	3.78 TB	80%	-1.69 MB	
		3.24 TB	4.82 TB	8.06 TB	40%	456.15 MB	
Free Space	5.07 TB						
Compression Savings	0%						
Compression Ratio	N/A						
Disk space saved through compression	0 MB						

Figure 1 Enterprise Manager storage type displaying Tier 1 as single redundancy and Tier 3 as dual redundancy.

Note: Fast indicates the Fast Track RAID levels while Standard indicates non-Fast Track RAID levels. Fast Track is a licensed feature that improves performance by consolidating frequently accessed blocks of data to the outer tracks.

1.5 Dual redundancy requirements

Redundancy requirements for each disk tier are based on the size of the disks in the tier. Dual redundancy may be required by the array during an installation or change to a storage type redundancy. The rules

followed for RAID level redundancy are based on capacity and on whether the disks are added to an existing or new storage type. Currently, disks up to 1.9 TB will default to dual redundancy for HDDs and up to 3.9 TB for SSDs. Beyond those capacities, dual redundancy may be required. The requirements and recommendations (defaults) for dual redundancy are shown in Table 3.

Table 3 Dual redundancy requirements based on capacity

Disk type	Recommended	New storage type (required)	Existing storage type (required)
HDD	967 GB to 1.9 TB	2.0 TB or larger	2.79 TB or larger
SSD	1.8 TB to 3.9 TB	4.0 TB or larger	4.0 TB or larger

For new SC Series arrays, HDDs 2 TB and above are required to use dual redundant RAID level (RAID 10-DM or RAID 6). For existing arrays with 2.79 TB or larger HDDs or 4.0 TB SSDs, dual redundancy is also required.

Note: For SSDs, these rules may change over time as technology advances. These rules apply to SCOS 7.0.1–7.1.1. Refer to the release notes for the appropriate SCOS version.

2 RAID rebuild, availability, and efficiency comparison

In general terms, rebuild rates, availability, and capacity efficiency are dependent on the RAID level.

Figure 2 depicts the rebuild rate based on RAID levels for the different SC Series disk options. The results are from rebuilds with a light random, 64KB, 70-percent read workload simulation.

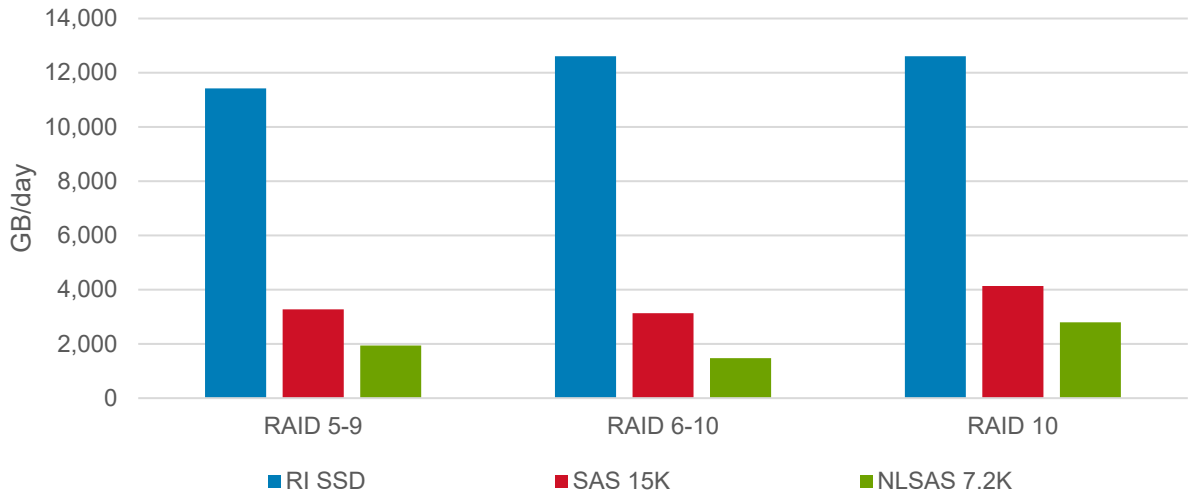


Figure 2 Rebuild rates for different RAID levels

Observations from Figure 2 show that SSDs have a great advantage over spinning disks for rebuilds with I/O activity. Spinning disks clearly show the progression of rebuild efficiency between RAID levels. RAID 10 experiences the greatest rebuild rate while RAID 6 is slightly lower than RAID 5.

2.1 SC Series RAID-level disk failure protection and reliability

In principle, RAID levels are primarily designed to protect against data loss with various levels of performance and capacity utilization. That data protection capability is provided by the ability of a RAID set to rebuild a failed disk while still servicing I/O requests.

Comparing the resilience of different RAID policies to protect against data loss in the event of a disk failure relies on a statistical analysis involving the following factors:

- Disk protocol, size, and RPM – physical characteristics of the disk
- Disk failure rates reflecting mechanical reliability – mean time before failure (MTBF) and annual failure rate (AFR)
- RAID level – RAID 10, RAID 10-DM, RAID 5-9, RAID 5-5, RAID 6-10, RAID 6-6
- RAID geometry – the RAID set construction or number of disks in the RAID
- RAID rebuild rate – the amount of time to rebuild a disk.

While there are modest variations in relative capacity utilization and performance, the levels of data protection provided by RAID policies vary dramatically, even when considering the same type of disk.

RAID 6 and RAID 10-DM provide the greatest protection, while SSDs have the fastest rebuild rates. 15K HDDs rebuild faster than 10K or NLSAS. Larger disks of any type are subject to the rate of rebuild to avoid a dual disk failure. As the number of disks and capacities have increased, RAID 6 or RAID 10-DM have become increasingly important to ensure overall data protection in storage arrays.

2.2 Space efficiency by RAID level

Each RAID level has capacity efficiencies when compared to the raw capacity available. The general efficiency of a specified RAID level is easy to compute by taking the disks that are not parity or mirror and dividing them by the total disks in the set.

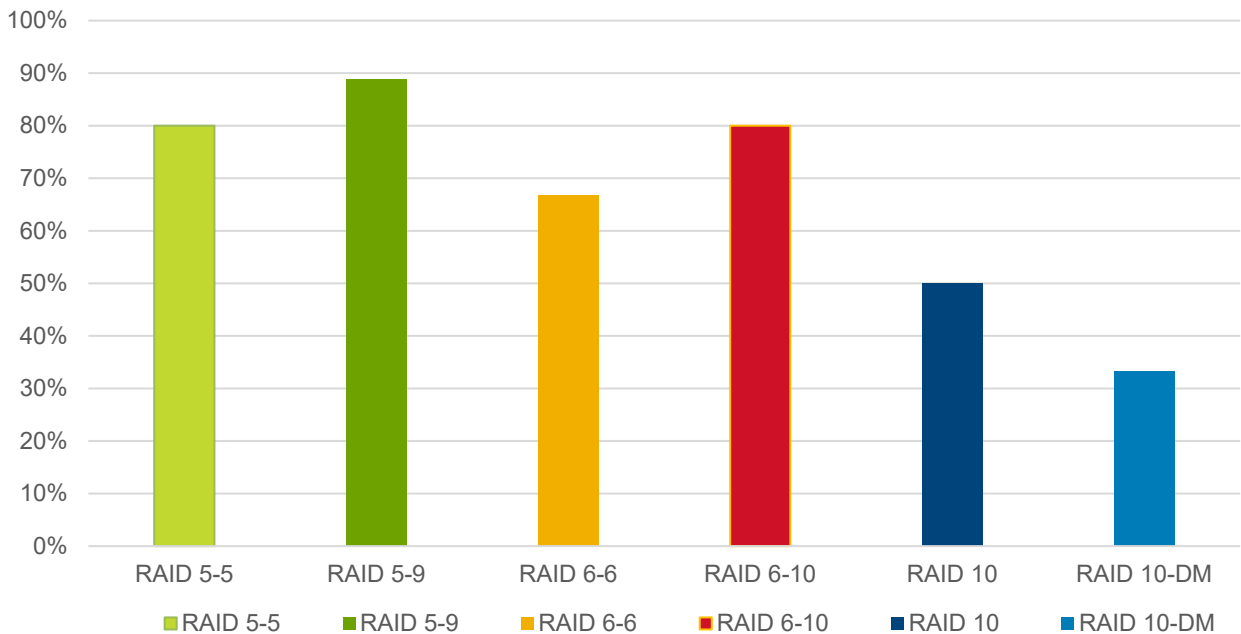


Figure 3 RAID level capacity efficiencies in SC Series arrays

2.3 SC Series RAID comparisons

Although all RAID levels provide good performance and data protection, there are some differences. RAID levels are chosen by the array to accommodate the performance and availability needs of that tier of disks. Deviations from the standard SC Series allocation of RAID levels should carefully consider the performance and availability impact on the workload.

RAID levels have different attributes for performance, reliability and rebuilds and are described in Table 4. The first column lists the RAID level while the other columns indicate the best suited workload along with rebuild performance and relative protection.

Table 4 RAID level and redundancy comparison.

RAID	Reads		Writes		Redundancy	Relative protection	Rebuild performance
	Random	Sequential	Random	Sequential			
10	Excellent	Excellent	Excellent	Excellent	Single	Excellent	Best
10-DM	Excellent	Excellent	Good	Good	Dual	Best	Best
5-5	Excellent	Excellent	Good	Excellent	Single	Good	Good
5-9	Excellent	Excellent	Good	Excellent	Single	Good	Good
6-6	Excellent	Excellent	OK	Good	Dual	Best	Good
6-10	Excellent	Excellent	OK	Good	Dual	Best	OK

Note: RAID performance and relative protection are directly affected during rebuild and reconstruction. RAID reconstruction times increase substantially as physical disk sizes increase. The time for reconstruction increases the potential for a second disk failure which exposes the vulnerability of the data on the array.

3 Spare disks

Depending on the RAID level and the total number of disks in each SC Series storage array, one or more spare disks are automatically configured and used in the event of a disk failure. The use of spare disks is highly recommended as an additional level of protection should a disk failure happen. Spare disks will replace the failed disk and allow the RAID set to rebuild. It is important to understand that two spare disks cannot guarantee the survival of a RAID 10 or RAID 5 set that has a multi-disk failure event or a disk failure during a RAID rebuild operation.

SC Series arrays automatically assign a single disk as a hot spare using these conventions:

- 2U enclosures: One spare disk for every disk class (SSD, 15K, 10K, 7.2K, and so on)
- 5U enclosures: One spare disk for every 21 disks. However, no single row contains more than one spare disk.

3.1 No spare available

When a spare disk is not available, the following RAID conventions are used.

- RAID 6 is guaranteed to survive the simultaneous failure of any two disks. Data continues to be available, but the set is degraded. A third disk failure in a degraded set can result in data loss.
- RAID 10 and RAID 5 are guaranteed to survive one disk failure per RAID set. Data continues to be available, but the set is degraded. A second disk failure in a degraded set can result in data loss.

3.2 Disk failure

When a disk in a RAID set fails, the SC Series array takes the following actions.

If a spare disk is available: The spare automatically replaces the failed disk. Data from the failed disk is reconstructed on the spare disk and continues to be available. During reconstruction, the set that contains the failed disk is temporarily degraded. After reconstruction, performance returns to normal.

If a spare disk is not available: Data continues to be available, but the set is degraded.

If another disk fails in a degraded RAID 10 or RAID 5 set: With more disks present in the SC Series system, chances of surviving a dual disk failure increase. However, there is no guarantee that the data will stay online. If a second disk fails and that disk is required to rebuild the initial failing disk, then all I/O is halted because the data is inaccessible. Contact Dell Support to attempt to data recovery in these situations.

If another disk fails in a degraded RAID 6 or RAID 10-DM set: The SC Series array continues to be degraded until the failed disks are replaced. Data is reconstructed on the replacement disks. If the first reconstruction is still underway, performance can be further reduced. After both disks are reconstructed, performance returns to normal.

3.3 Replaced disk

When a failed disk is replaced, the SC Series array responds as described below.

If a spare disk was used: Data has already been reconstructed on the spare disk, so the new disk becomes a spare.

If a set is degraded: Data is reconstructed on the new disk and after reconstruction, and performance will return to normal.

4 SC Series RAID level use

By default, an SC Series array groups all disks in a single disk folder. This disk folder contains all the various speeds and capacities of disks. Redundancy levels can be configured for each tier of disk. Tier 1 disks are the fastest disks in SC Series storage, while other disks are grouped according to their speed in the other two available tiers (Tier 2 and Tier 3). RAID is allocated within a tier across all of the disks. However, unlike other storage architectures the SC Series, tiers of disks may have multiple RAID configurations simultaneously. When a volume is created, it is allocated across all of the tiers and RAID levels.

Tier 1 is typically configured with RAID 10 and RAID 5, and Tier 3 typically includes large disks that are automatically configured for dual redundancy or RAID 10-DM and RAID 6.

Figure 4 shows a typical disk folder with the different RAID levels and their tier associations.

Storage Types	Storage Chart	Name	Disk Folder	Allocated Space	Used Space	Free Space
		Assigned - Redundant - 2 MB	Assigned	8.48 TB	3.4 TB	5.07 TB
		Tier 1 RAID 10 - Standard		391.97 GB	155.82 GB	236.15 GB
		Tier 1 RAID 5-5 - Standard		38.42 GB	15.37 GB	23.05 GB
		Tier 3 RAID 10-DM - Fast		582.52 GB	0 MB	582.52 GB
		Tier 3 RAID 10-DM - Standard		1.44 TB	0 MB	1.44 TB
		Tier 3 RAID 6-10 - Fast		2.27 TB	187.34 GB	2.09 TB
		Tier 3 RAID 6-10 - Standard		3.78 TB	3.05 TB	739.67 GB

Figure 4 Tier 1 and Tier 3 with various RAID levels

The disk folder is labeled **Assigned** by default and the storage type is **Assigned – Redundant – 2MB**, which refers to a redundant storage type allocated with 2MB data pages.

4.1 SC Series disk and RAID layout

Within a disk folder, the SC Series automatically distributes the data and parity or mirrors to avoid hot spots across the available disks. Each tier automatically configures the appropriate RAID level, if a particular RAID level is not used the SC Series system removes the unused extents that make up that RAID set so that space is maximized in that tier.

For illustrative purposes, Figure 5 shows a possible RAID distribution in a twelve-disk enclosure with one global hot spare. Tier 1, which includes the Fast Track feature, has both RAID 10 and RAID 5-5 (RAID 5-9 is also an option).




Figure 5 SC Series example of RAID 5-5 and RAID 10 in Tier 1 over 12 disks with one spare.

The example in Figure 5 shows RAID distributed across all of the disks, and parity is staggered as much as possible while hot spots are minimized. Reading the graphic from left to right, RAID 5-5 allocates a parity every fifth disk, starting in the outer tracks or Fast Track regions, then wraps to the standard tracks. RAID 10 follows the same pattern. Notice that in the Fast track region, one segment in D11 is unallocated since RAID 10 needs a two disk relationship. The three unallocated regions in the standard sections of D9-D11 represent where the next RAID 10 or RAID 5-5 would start if the system needed that area.

4.2 Application data writes

By default, the SC Series array writes all application data to the highest performing RAID level in Tier 1, which is typically RAID 10, then later moves the data according to the user-defined schedules. Writing to RAID 10 is an architectural advantage since RAID 10 is the highest performing RAID level for random write profiles. Since random writes typically have the most impact on any storage environment, RAID 10 minimizes the effect on performance.

Figure 6 represents access from the application to the various RAID levels in the various tiers. The application may read from any tier and any RAID level, but only writes to Tier 1 and RAID 10 as long as space is available.

The SC Series array allows application writes to available space in other tiers and RAID levels if RAID 10 on Tier 1 is full. If RAID 10 on Tier 1 is full, the following precedence will occur for application writes to the SC Series array.

- If RAID 10 is full, writes will occur on RAID 5.
- If Tier 1 disks are full, writes will occur on the next tier.
- If Tier 3 RAID 10-DM is full, writes will occur on Tier 3 RAID 6.

Proper design of Tier 1 is essential to maximize the utilization and performance of the SC Series array.

When the SC array or disk folder consists only of large disks such as 2 TB HDDs or 4 TB SSDs or larger, the tier is dual redundant. For these situations, RAID 10-DM is used for writes, while RAID 6 is updated by Data Progression. Consider the performance implications for large Tier 1 disk solutions.

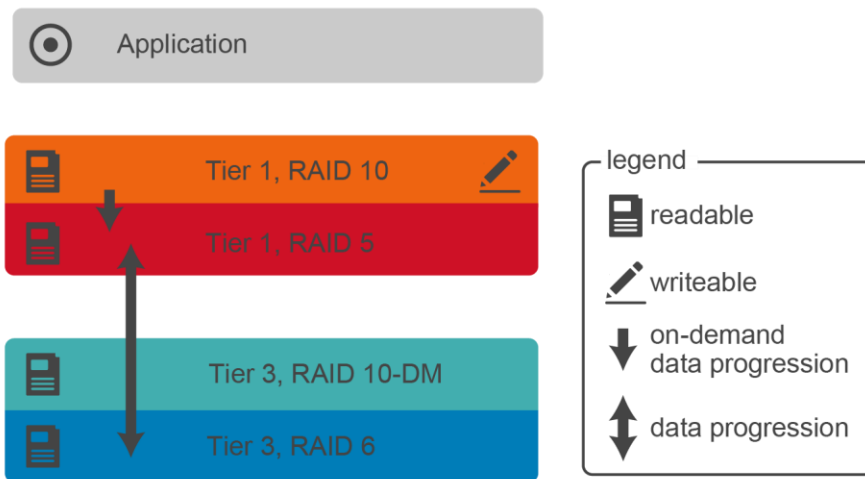


Figure 6 Application I/O writes only to RAID 10

Data Progression, the automated data movement by the SC Series array, frees up as much RAID 10 space as possible. On-demand data progression moves data in Tier 1 from RAID 10 to RAID 5.

4.3 Highest performing RAID write advantage

The SC Series architecture is designed to write to the highest performing RAID level. This section describes why RAID 10 is considered the highest performing RAID.

RAID technology applies the basic idea that multiple copies of data protect the larger set of data in the event of a disk failure. Different RAID configurations require additional I/O during updates to ensure this availability. For instance, RAID 5 has a single parity disk (or area on a disk) that is used to rebuild the failed disk. RAID 10 simply has a duplicate or mirror disk that is updated with the primary disk. In the case of RAID 5, when a write occurs, this parity needs to be modified by reading and recalculating and then writing the data back to the disks. From the theoretical standpoint, RAID 5 requires four I/Os; this is sometimes known as the RAID write penalty.

The RAID 5 write penalty consists of:

- Read from data
- Read from parity
- Write new parity
- Write new data

RAID 10, on the other hand, only writes the changed data to the primary and mirror disks, and therefore has about a 50-percent less burden on the disks during writes. The low overhead of operations during writes comes at the cost of usable capacity (50-percent efficient), however an SC Series array moves data to more efficient RAID levels or tiers according to their access patterns.

Figure 7 indicates the operational advantage during writes that RAID 10 has over RAID 5. In a single redundant tier, all application updates go to RAID 10 with only two I/Os, in contrast to the four I/Os needed for an application writing to RAID 5.

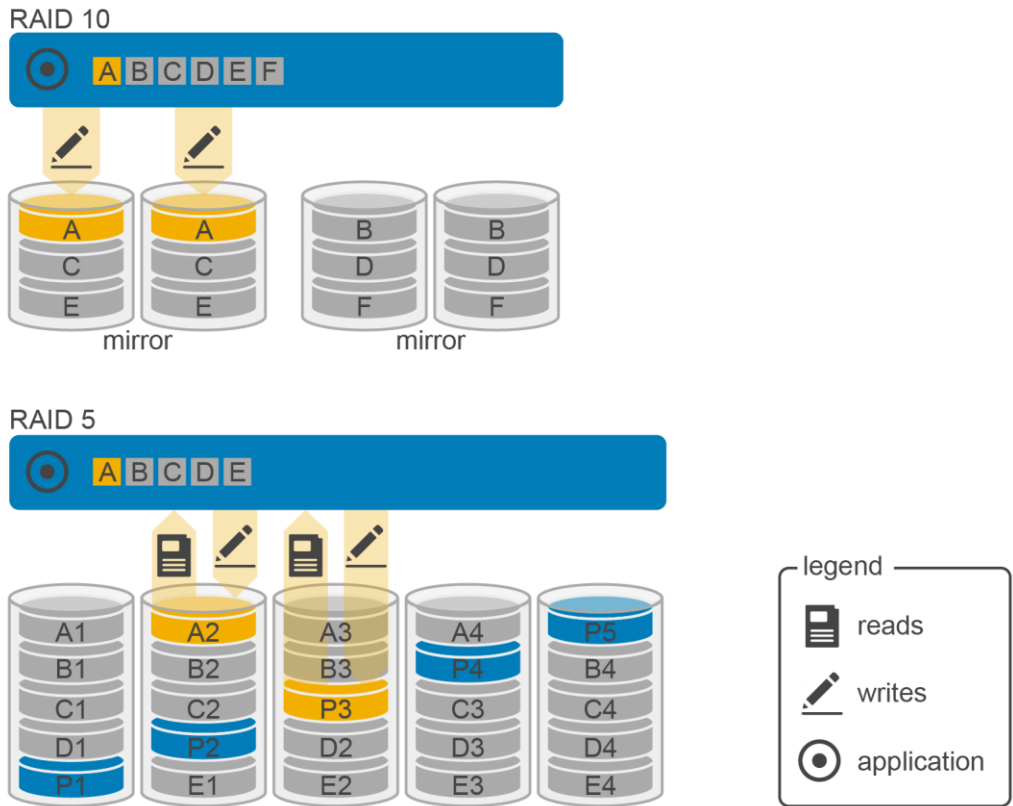


Figure 7 Difference in write penalty between RAID 10 and RAID 5

For most SC Series arrays, Tier 1 has RAID 10. However, if the tier requires dual redundancy writes, they occur on RAID 10-DM. Table 5 shows the Tier 1 and RAID-level write penalties and where application writes occur according to the redundancy setting.

Table 5 Tier 1 application RAID level write policy based on redundancy

Redundancy	RAID level	Write penalty	Application writes
Single	RAID 10	2	Yes
Single	RAID 5	4	No
Dual	RAID 10-DM	3	Yes
Dual	RAID 6	6	No

Note: SC Series arrays follow this policy as long as appropriate space is available for each RAID level and tier.

5 Snapshots and Data Progression

Over time, an SC Series array determines the appropriate movement of data based on the frequency of access. A snapshot is a point-in-time-copy (PITC) of a volume that provides fast recovery of data. Snapshots do not copy the data but simply freeze the data as read only. New writes to the volume allocate new space in the highest performing RAID level. Snapshots are integral to taking full advantage of the SC Series architecture that is intended to use the fastest disks from a performance perspective while efficiently using larger, slower disks to store less active data.

5.1 Data Progression

Data Progression moves data within a virtualized storage environment, between tiers and drive types, as well as among multiple RAID levels within the same tier for a constant balance of performance and cost.

Data Progression runs every 24 hours to assess disk use and moves data to disk space that is more efficient for the data usage. By default, Data Progression runs each day at 7 PM, but this schedule can be changed in the SCOS settings. The storage profile applied to each volume determines Data Progression behavior.

Note: Storage profiles are associated with the volume and are used to determine how to move data most effectively.

5.2 On-Demand Data Progression

Data Progression also moves snapshots. When a scheduled or manual snapshot is created, the data is frozen and moved to the tier specified by the volume storage profile. When Data Progression is run outside of the daily schedule, it is often referred to as On-Demand Data Progression (ODDP). Used by both HDDs and SSDs, ODDP may be triggered by the volumes several times a day.

5.3 Data Progression and storage profiles

A storage profile works with Data Progression to indicate how the data for the volume is moved between tiers. An SC Series array has several default storage profiles to choose from to accommodate most situations. For instance, using the recommended profile allows the volume to use all tiers. Active data will progress from Tier 1, RAID 10 to Tier 1, RAID 5 as a snapshot, and then move to RAID 5 or RAID 6 in lower tiers after it has not been accessed for a certain amount of time.

Snapshots prevent writes to data and then Data Progression moves the data to other RAID levels and tiers. Figure 8 is an example of Data Progression for volumes using the recommended storage profile in a typical two-tier environment.



Figure 8 Data Progression in a two-tier configuration

For volumes using the recommended storage profile in a three-tier configuration, Data Progression moves data to each tier as illustrated in Figure 9.



Figure 9 Data progression in a three-tier configuration

5.4 Capacity within tiers

Statistically, only about 20 to 30 percent of user data is actively serving the needs of the applications, while 70 to 80 percent is less active. For capacity-planning purposes, assume that 20 to 30 percent of the total capacity will be directed to Tier 1 disks and 70 to 80 percent can be accommodated by the other tiers of disks. Since the data movement is dynamic, visibility to the actual RAID-level allocations is best observed by the Dell Storage Manager Client or Enterprise Manager.

Figure 10 is an example of allocation of space by tier and RAID level.

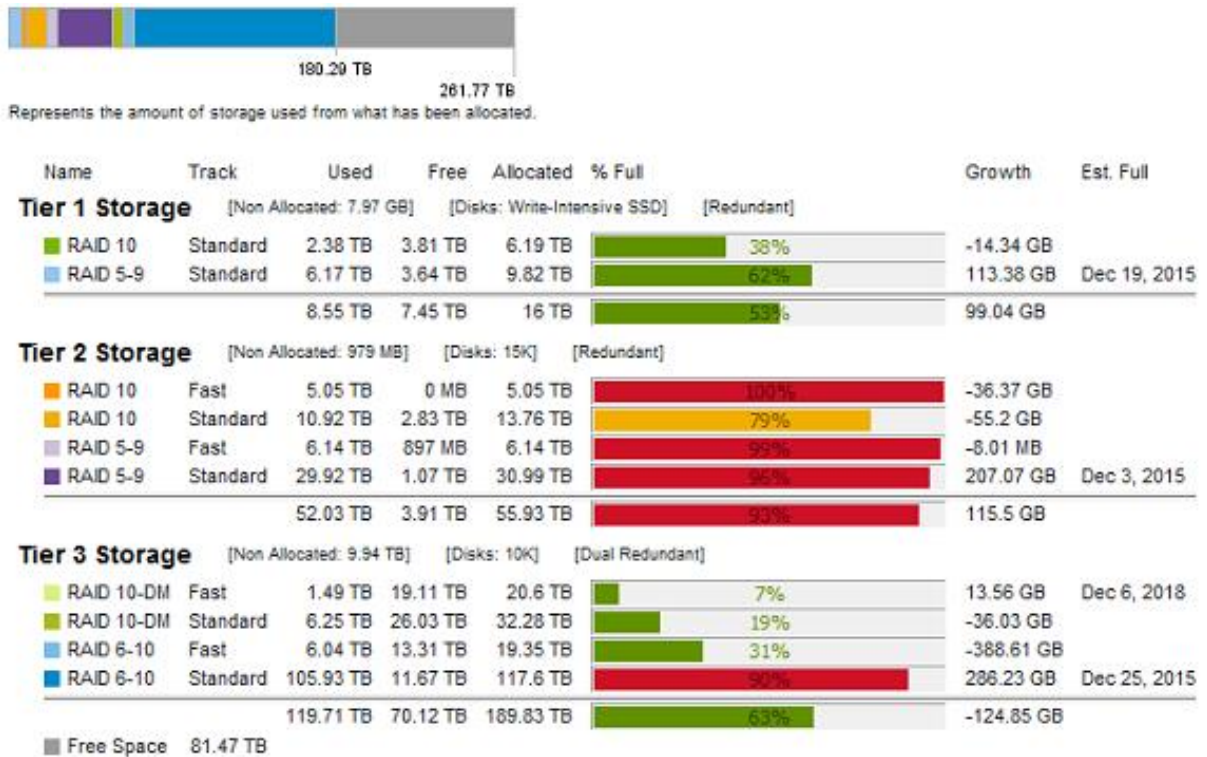


Figure 10 Storage types showing space utilization based on tier and RAID level

The SC Series arrays provide tools to show where the data will be placed during the next Data Progression cycles. As an example, the Enterprise Manager pressure report shows the direction that data is expected to travel with the next Data Progression (see Figure 11).

Assigned - Redundant - 2 MB

Index 3
 Disk Folder Assigned
 Redundancy Redundant
 Data Page Size 2 MB

Summary | Volumes | Historical Usage | **Pressure Report**

Tier 1 Storage

Disk Space Available in Tier: 8.01 TB Disk Space Allocated: 8.01 TB

RAID Level	Disk Track	Chart	Disk Allocated	Disk Used	Moving Up	Moving Down	Volume Allocated	Volume Used
RAID 10	Fast		0 MB	0 MB	0 MB	0 MB	0 MB	0 MB
RAID 5-9	Fast		0 MB	0 MB	0 MB	0 MB	0 MB	0 MB
RAID 10	Standard		4.84 TB	347.38 GB	0 MB	3.34 GB	2.42 TB	173.69 GB
RAID 5-9	Standard		3.16 TB	1.2 TB	0 MB	216 MB	2.81 TB	1.07 TB

Tier 2 Storage

Disk Space Available in Tier: 75.31 TB Disk Space Allocated: 75.3 TB

RAID Level	Disk Track	Chart	Disk Allocated	Disk Used	Moving Up	Moving Down	Volume Allocated	Volume Used
RAID 10	Fast		2.71 TB	2.71 TB	0 MB	792.91 GB	1.36 TB	1.36 TB
RAID 5-9	Fast		12.35 TB	12.32 TB	0 MB	23.53 GB	10.98 TB	10.95 TB
RAID 10	Standard		28.01 TB	7 TB	1.88 TB	1.23 TB	14.01 TB	3.5 TB
RAID 5-9	Standard		32.23 TB	17.41 TB	627.75 MB	651.47 GB	28.65 TB	15.47 TB

Tier 3 Storage

Disk Space Available in Tier: 145.55 TB Disk Space Allocated: 145.55 TB

RAID Level	Disk Track	Chart	Disk Allocated	Disk Used	Moving Up	Moving Down	Volume Allocated	Volume Used
RAID 10-DM	Fast		6.63 TB	3.03 TB	0 MB	1.53 TB	2.21 TB	1.01 TB
RAID 6-10	Fast		22.47 TB	6.62 TB	0 MB	1.05 TB	17.98 TB	5.3 TB
RAID 10-DM	Standard		9.98 TB	0 MB	0 MB	0 MB	3.33 TB	0 MB
RAID 6-10	Standard		106.46 TB	105.38 TB	0 MB	0 MB	85.17 TB	84.3 TB

Figure 11 Pressure report showing data movement as indicated by green up, orange down. Boxes added for illustration.

Note: Since Tier 1 is all SSD, Fast Track (Fast) RAID levels are not allocated.

This report shows the expected direction and quantity of data that will be moved during the next Data Progression cycle. In the table, the amount of data is indicated in the **Moving Up** or **Moving Down** columns. Data that moves down is first moved to the lower RAID, and then to the lower performing tier. Data that moves up is promoted to the next higher performing tier.

6 RAID management

SC Series arrays are designed to manage RAID levels automatically, however the following operations may be performed if the need arises;

- Rebalancing RAID: Redistributes data and typically occurs after changing the number of disks or changing the redundancy of the tier
- Modify tier redundancy: Increases or decreases the number of disks for parity or mirror
- Creating storage types: Associates disks to a redundancy level
- Create a storage profile: Allows for RAID level and tier binding at the volume level

6.1 RAID rebalancing considerations

When disks are added or removed or tier redundancy is modified, a RAID rebalance is needed to redistribute the data across all disks in a disk folder.

The RAID rebalance process follows these guidelines:

- RAID rebalancing begins with the smallest RAID sets (RAID 5-5).
- RAID rebalancing has the potential to fail or deadlock when moving data on some RAID sets. If this occurs, it may be necessary to add disk space to the array for RAID rebalancing to complete.

Depending on the total capacity used, a considerable amount of time may be needed to completely redistribute the data pages. In addition, ample space needs to be available before this operation can occur. Space available, amount of data to rebalance, the requirement for new RAID sets and load on the system all impact the time for the rebalance to complete.

6.2 Modifying tier redundancy

Each disk folder has a corresponding storage type. Tier redundancy may be changed from single parity or mirror to dual parity or mirror.

Before modifying tier redundancy, be sure there is sufficient space in the tier for a raid rebalance, and consider any impact a RAID rebalance may have on the system.

Figure 12 shows Tier 1 storage allocated as single redundancy.

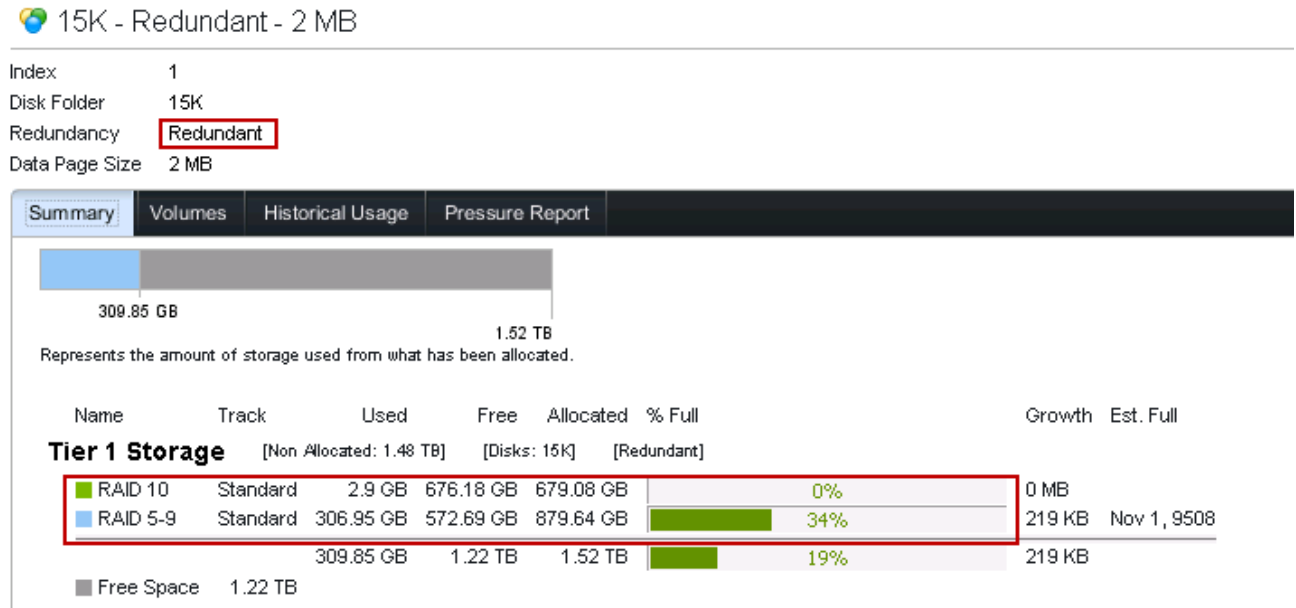


Figure 12 A single redundant tier

In the **Modify tier Redundancy** wizard, choices are available to change the redundancy to dual redundancy in each tier. The storage type is configured for all three tiers regardless of the presence of physical disks in those tiers. However, the redundancy may be selected based on tier as needed.

With dual redundancy, the RAID levels change from RAID 10 to RAID 10-DM, and from RAID 5 to RAID 6 as shown in Figure 13.




Figure 13 Dual redundancy tier

Note: A RAID rebalance occurs automatically after this change. **Modification in Progress** indicates the RAID rebalance is underway.

Some occasions require the tier redundancy to be changed from dual redundancy to single redundancy. This may only be achieved for SSDs smaller than 4 TB or HDDs smaller than 2 TB. Factors such as free space, system activity, time of rebalance, and impact on availability should be considered before making this change.

Figure 14 shows redundancy options available for each tier, and the change from a dual redundant to a single redundant storage type.

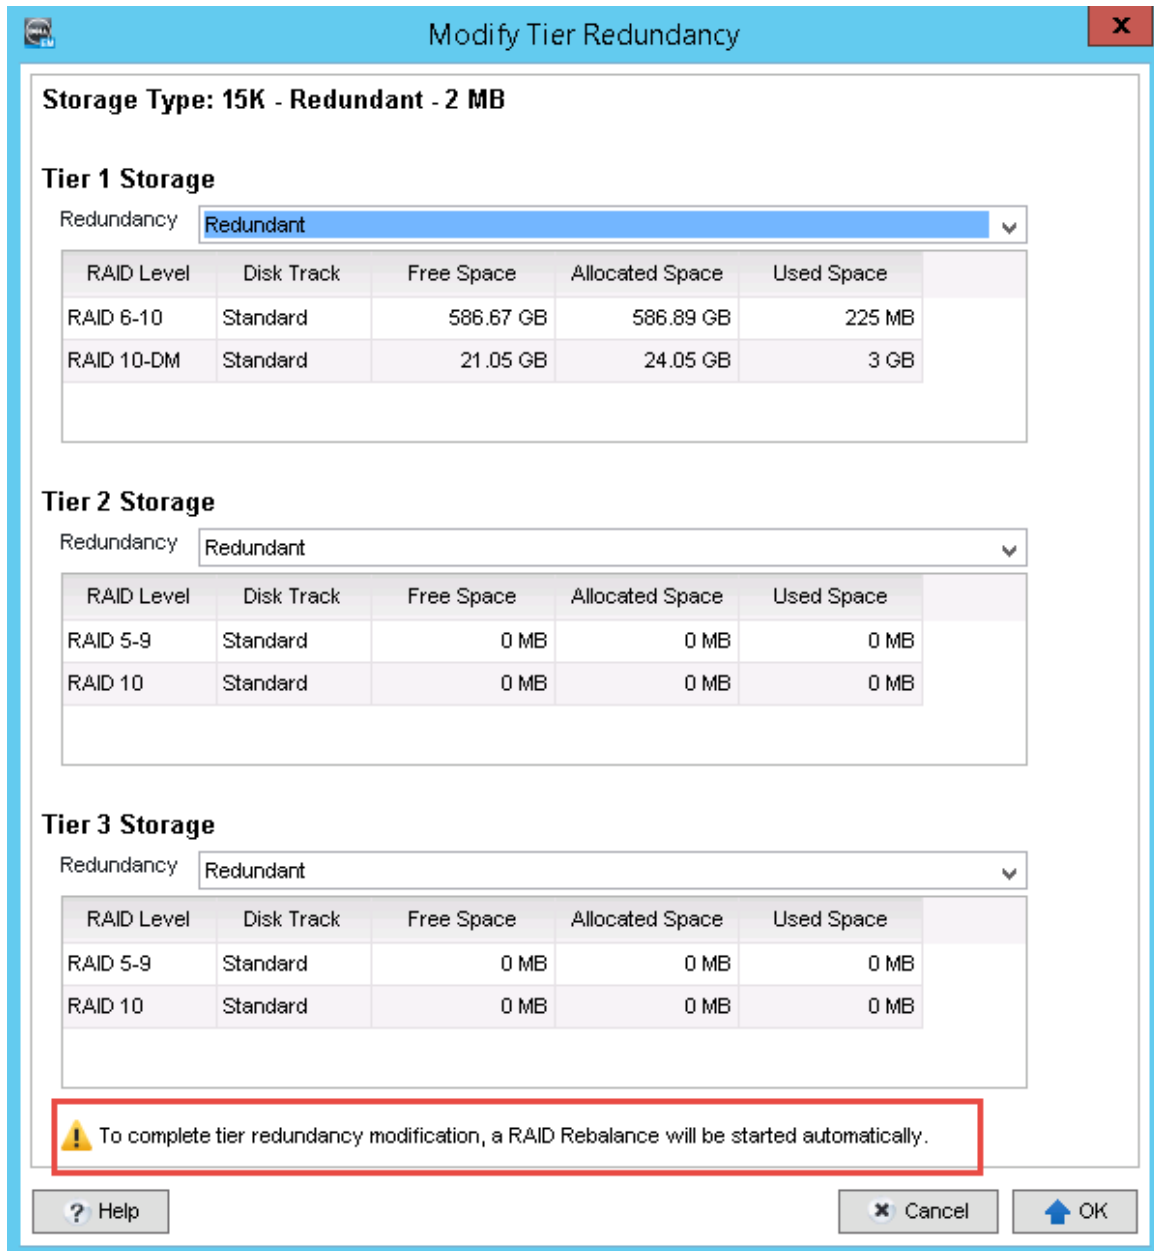


Figure 14 Changing from dual to single redundancy for all tiers within this storage type

Note: The RAID rebalance will be started automatically.

After the rebalance has completed, the storage type will change the RAID 10-DM and RAID 6 to RAID 10 and RAID 5, respectively.

Figure 15 shows the results of changing from dual to single redundancy.

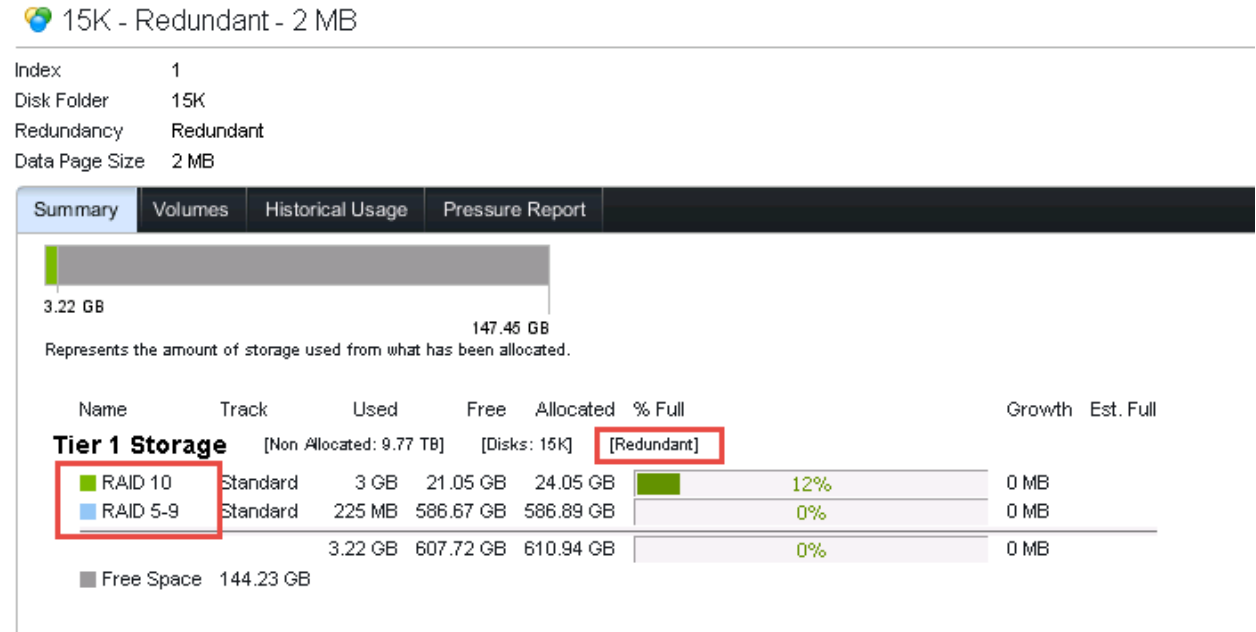


Figure 15 Dual to single redundancy

Note: Changing a dual redundant storage type to a single redundant storage type requires careful consideration and should be discussed with Dell Support prior to implementation.

6.3 Creating a storage type

A storage type is a pool of storage with specified redundancy levels. SC Series arrays assess the disks available in a disk folder and present the applicable storage type options. The default storage type and storage profiles provide the best performance. If site-specific requirements dictate that the new disks need to be isolated from other disks, create a disk folder for those disks and configure a storage type. A typical use case may be for an application that requires specific disks and redundancy to meet business, performance, or availability requirements.

Note: The SCv2000 Series controllers manage storage types automatically and do not allow storage type creation.

Figure 16 demonstrates a storage type created on a new disk folder with all Tier 1, 2, and 3 disks. For this storage type example, dual redundancy is chosen to meet a higher availability requirement.

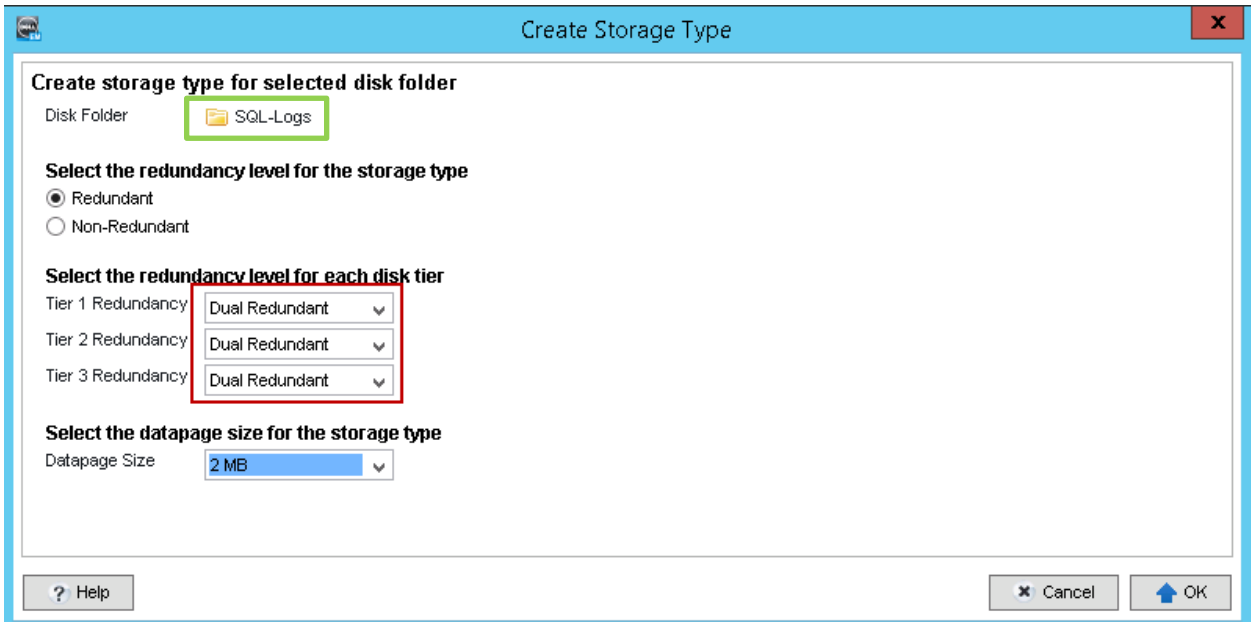


Figure 16 New storage type with dual redundant tiers

The storage type could be selected by volumes to take advantage of the disks in the **SQL-Logs** disk folder.

Note: Best practices recommend a single storage type per disk folder.

The results of creating a new storage type show that the new disks (in this example, the **SQL-Logs** disk folder) have both Tier 1 and Tier 3 disks with dual redundancy protection (RAID 10-DM and RAID 6).

Storage Types

Storage Types		Storage Chart	
Name	Disk Folder	Allocated Space	
Assigned - Redundant - 2 MB	Assigned	8.48 TB	
Tier 1 RAID 10 - Standard		391.97 GB	
Tier 1 RAID 5-5 - Standard		38.42 GB	
Tier 3 RAID 10-DM - Fast		582.52 GB	
Tier 3 RAID 10-DM - Standard		1.44 TB	
Tier 3 RAID 6-10 - Fast		2.27 TB	
Tier 3 RAID 6-10 - Standard		3.78 TB	
SQL-Logs - Redundant - 2 MB	SQL-Logs	24.04 GB	
Tier 1 RAID 10-DM - Fast		24.04 GB	
Tier 1 RAID 10-DM - Standard		96.84 GB	
Tier 1 RAID 6-6 - Standard		96.82 GB	
Tier 3 RAID 6-10 - Fast		163.04 GB	
Tier 3 RAID 6-10 - Standard		652.08 GB	

Figure 17 New disks are assigned to the **SQL-Logs** disk folder with a dual redundancy storage type

The volume will have the choice to use a specific storage type during creation. However, once created, this volume will not be able to be modified to use a different storage type.

Figure 18 shows the new storage type selected for a new volume.

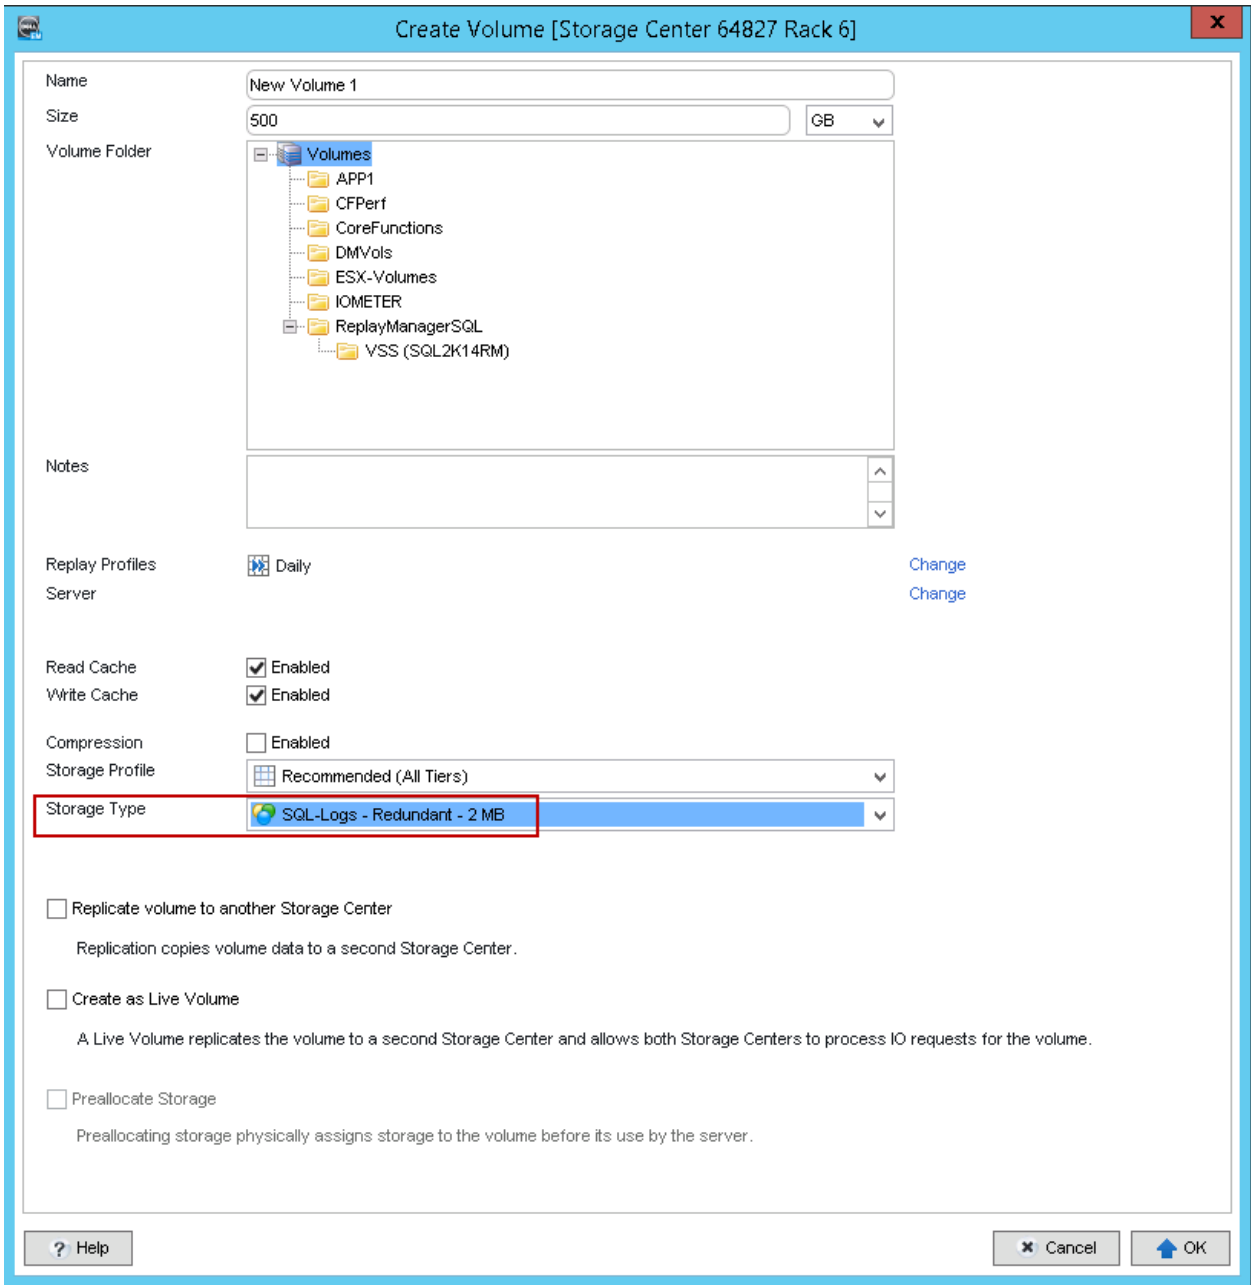


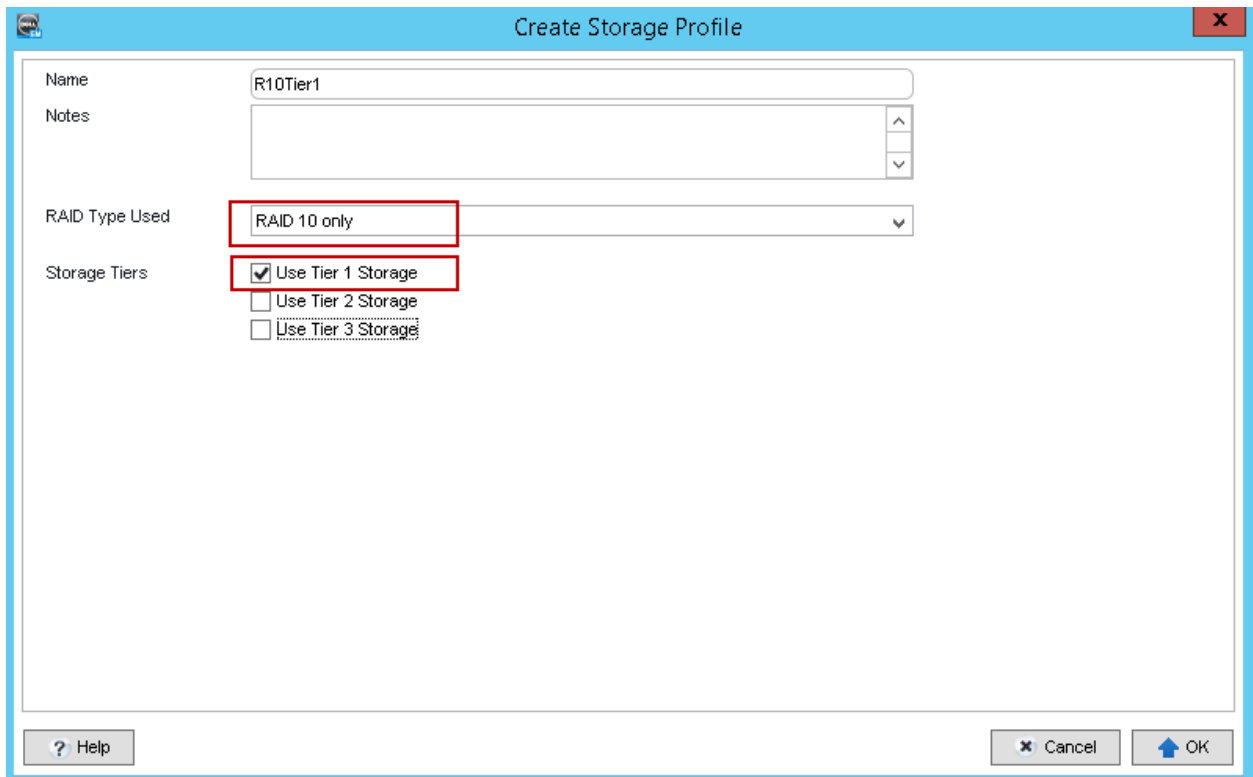
Figure 18 New storage type selected during new volume creation.

6.4 Creating a storage profile

A storage profile can further associate data to a particular tier as well as one or more RAID levels available in that tier. Occasionally, an application may need to override the automation built into the array and bind a volume to a particular RAID level and/or tier to accommodate performance, capacity or availability needs.

Note: Storage profiles cannot be created in the SCv2000 Series arrays.

A storage profile is easily created using the SCOS UI.



The screenshot shows a 'Create Storage Profile' dialog box with the following fields and options:

- Name:** R10Tier1
- Notes:** (Empty text area)
- RAID Type Used:** RAID 10 only
- Storage Tiers:**
 - Use Tier 1 Storage
 - Use Tier 2 Storage
 - Use Tier 3 Storage

Buttons at the bottom include: ? Help, Cancel, and OK.

Figure 19 Creating a custom storage profile with RAID 10 only in Tier 1

Each storage profile may be assigned to single or multiple volumes. Figure 20 shows the custom storage profile created in Figure 19 assigned to several volumes.

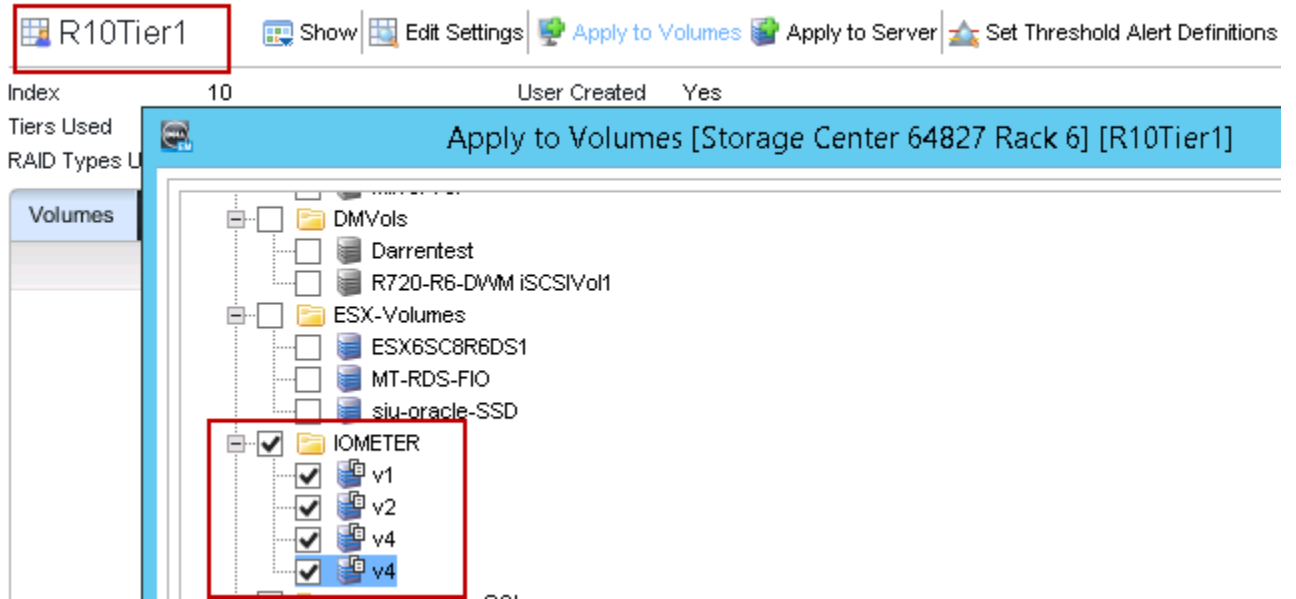


Figure 20 Several volumes assigned to the new RAID 10, Tier 1 only storage profile

7 Design considerations

Once the storage requirements have been determined, a solution may be designed with typically two or three tiers of storage. However, some environments may only use a single tier. Most often in multi-tiered solutions, the top tier is sized for 100 percent of the I/O and 20 to 30 percent of the capacity. However, specific customer needs may require deviation from this starting point. These tiers may consist of SSDs, HDDs, or a combination of the two. For example, in an all-SSD solution, Tier 1 consisting of write-intensive SSDs would be allocated as RAID 10, while Tier 2 containing read-intensive SSDs would be RAID 5. This solution could be further expanded with a third tier consisting of 7.2K HDDs for storage of cold data. Another possible solution consisting of only HDDs might have a 15K Tier 1, a 10K Tier 2, and a 7.2K Tier 3. The potential combinations are numerous and are best selected by consulting with a Dell Storage specialist or reseller.

7.1 SSD considerations

SSDs have increased in capacity while maintaining excellent performance. Dell has included 3.84 TB SSDs which may be used as Tier 1 disks. The default for 3.84 TB SSDs is dual redundancy, however that may be modified to single redundancy based on analysis of performance data or business needs. For disks 4 TBs and larger, dual redundancy is required, these SSDs would be configured in Tier 1 if no write-intensive SSDs exist.

7.2 SSDs in Tier 1 dual redundancy

Since dual redundancy affects the write penalty, consider a modest performance impact when designing with larger SSDs. The SC Series design criteria calls for Tier 1 to handle the IOPS load with additional disks for capacity. An SC Series array is designed to accommodate all IOPS in RAID 10. In the following definitions, **Total IOPS** references the back-end IOPS expected to the disks.

Single redundant sizing using RAID 10: Total IOPS = reads + writes * 2

When Tier 1 is dual redundant, the basic formula changes for the total IOPs supported. Since an additional mirror is needed for RAID 10-DM the write penalty is three.

Dual redundant sizing using RAID 10-DM: Total IOPs = reads + writes * 3

When designing solutions using large capacity SSDs, consider the total IOPS burden in high write environments.

When mixing large SSDs (greater than 3.9 TBs) with HDDs, the SSDs are in Tier 1 as RAID 10-DM and RAID 6. The HDDs are in Tier 2, Tier 3, or both; the redundancy level depends on the capacity of the disks.

Design SC Series arrays with the appropriate Dell Storage specialist to arrive at a solution that best fits the workloads and business requirements.

7.3 Techniques for specific workloads

SC Series storage is designed to handle applications with many different workloads. However, some workloads may need to deviate from the SC Series defaults and work with the available features of the array.

The following conditions are typical reasons that may alter a standard SC Series setup.

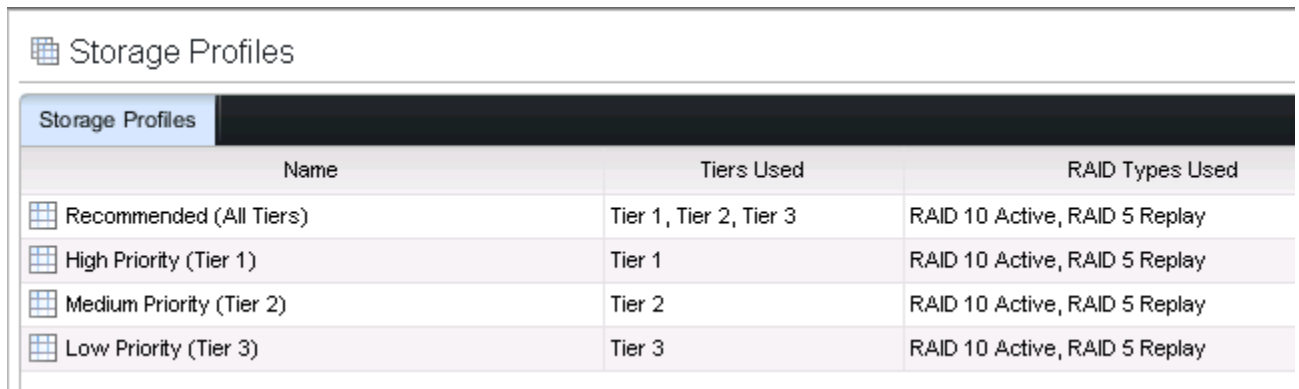
Site requirements for isolated disks: Multi-tenancy or other business requirements

High-write-bandwidth applications: Seismic recording, video recording, video editing, surveillance

For these types of applications, create a storage type with a set of disks to allow isolation of the applications, increase or decrease availability, and provide predictive performance.

The storage type will isolate disks and setup redundancy, but will not necessarily bind a volume to a particular disk type or RAID level. To further tune the applications, use a custom (or available) storage profile. For example, a custom RAID 10, Tier 1 only storage profile would configure the volumes to use the new storage type along with the storage profile that would isolate the disks to the volumes and force the RAID levels and tiers selected.

Several storage profiles are available depending on the type of SC Series arrays in the configuration. The following are the standard storage profiles for all SC Series arrays except the SCv2000.



Name	Tiers Used	RAID Types Used
Recommended (All Tiers)	Tier 1, Tier 2, Tier 3	RAID 10 Active, RAID 5 Replay
High Priority (Tier 1)	Tier 1	RAID 10 Active, RAID 5 Replay
Medium Priority (Tier 2)	Tier 2	RAID 10 Active, RAID 5 Replay
Low Priority (Tier 3)	Tier 3	RAID 10 Active, RAID 5 Replay

Figure 21 Standard storage profiles available

Applications that require a continuous amount of large block writes, typical of seismic recording, video editing, and surveillance, may follow a similar technique.

Design technique for high write applications:

- Use unique disks for the application and create a storage type to isolate these volumes to the appropriate disks.
- Use an existing storage profile, such as low priority, to allow the volumes to only write to Tier 3 disks.
- Use a custom storage profile to allow the volumes to only write to RAID 5/RAID6.

Using an existing storage profile: For this case, a dual redundant, low priority storage type would use only RAID 10-DM and RAID 6 in Tier 3.

Using a custom storage profile: Create a custom storage profile allowing volumes to use all tiers and only RAID 6. The volumes could then be created on the new storage type with the appropriate storage profile.

Dell recommends that the technical resources from the account team are involved in designing the SC Series solutions to meet the needs of the business.

8 Summary

SC Series arrays take advantage of disk tiering, RAID levels, and automation to fully optimize disks and data efficiently. RAID levels are allocated to meet the most demanding requirements of capacity, availability, and performance. Understanding the implementation of the SC Series virtualization and automation that allows data to move between types of disks and RAID levels helps in selecting proper protection and performance to match the business requirements.

A Additional resources

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

Dell TechCenter is an IT Community where you can connect with Dell customers and Dell employees for the purpose of sharing knowledge, best practices, and information about Dell products and installations.

Dell Storage technical documentation libraries are listed at:

<http://en.community.dell.com/techcenter/storage/w/wiki/2631.storage-solutions-technical-documents>.