

Dell PS Series Storage: Choosing a Member RAID Policy

This paper describes the supported RAID policies for Dell PS Series arrays and provides information on the RAID policies under normal and failure conditions to help administrators choose the appropriate RAID policy.

Dell Storage Engineering
February 2017

Revisions

Version	Date	Description
1.0	2007	Initial Release
4.1	2009	Added RAID 6
4.2	June 2012	Added large disk recommendations and RAID information for 24 disk chassis
4.3	August 2012	Added RAID configuration recommendations and changes reflected in PS firmware v6.0
4.4	April 2013	Added info regarding improved RAID reliability score in SANHQ
4.5	December 2014	Added 1 TB Raid Policy recommendation
4.6	December 2015	Added PS6610, PS4210, large drive RAID 6 enforcement
4.7	February 2017	Added 6 TB RAID 6 recommendation

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 © 2007 - 2016 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [2/1/2017] [Best Practices] [TR1020]

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

Table of Contents

1	Introduction to member RAID policies	4
1.1	RAID policy descriptions.....	5
2	RAID policy availability and performance comparisons	6
2.1	RAID policy disk failure protection and reliability.....	6
2.2	RAID policy impact on performance	9
2.3	Spare disks	12
2.4	Examples of data protection with different RAID policies.....	13
2.5	Relative capacity utilization by RAID policy.....	16
3	Setting the member RAID policy	17
4	Displaying the RAID level space distribution.....	18
5	PS Series array disk layout	19
6	Converting or migrating from a member RAID policy.....	20
6.1	Large disk policy for RAID 50 and RAID 10	20
6.2	Online conversion	20
6.3	Array evacuation.....	22
6.4	Online volume migration.....	22
6.5	Backup and restore	24
7	Summary	25
A	12-Disk PS Series array RAID configurations	26
B	14-Disk PS Series array RAID configurations	27
C	16-Disk PS Series array RAID configurations	28
C.1	RAID 50	28
C.2	RAID 5	29
C.3	RAID 6	29
D	48-Disk PS Series array RAID configurations	31
E	42 or 84-disk PS Series RAID configurations	32
F	Technical support and resources	33

Executive summary

One of the most important decisions to make when deploying a new storage solution is which RAID type(s) to use. Each RAID type has its advantages and disadvantages in terms of protection, efficiency (usable storage), and performance—both during normal operations and in terms of rebuild time. Regardless of the storage solution, RAID decisions will be based on the same set of criteria. First and foremost, the application's performance needs must be met, and this should be the primary decision factor when selecting a RAID policy. Once the application(s)'s needs are addressed, choosing a RAID policy comes down to two factors: cost—in terms capacity utilization efficiency—and protection level.

This document discusses RAID types available with Dell™ PS Series Arrays and how to select the right RAID policy in terms of performance, protection, and efficiency.

Preface

PS Series arrays optimize resources by automating performance and network load balancing. Additionally, PS Series arrays offer all-inclusive array management software, host software, and free firmware updates.

Table 1 Software and firmware used for the preparation of this paper

Vendor	Model	Software Revision
Dell™	PS Series Firmware	v5.2.x, v6.x ,v7.x, v8.x, v9.x and later
Dell™	PS Series Array Hardware	PS6000, PS6100, PS6110, PS6210, PS4000, PS4100, PS4110, PS4210, PS6500, PS6510 and PS6610 Series

The information in this guide is intended for storage administrators of Dell PS Series arrays.

1 Introduction to member RAID policies

PS 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, PS Series arrays support several different RAID types with each type optimized to maximize performance of the PS SAN architecture. When an array is added as a member to a PS Series group, the RAID policy for that member is chosen. Depending on the PS Series array model and number of disks in the array, RAID characteristics are implemented with various numbers of disks dedicated to data, parity, mirrors, and hot spares.

The RAID policy choices for a member consist of RAID 6, RAID 10, RAID 50 and RAID 6 Accelerated (Hybrid arrays only).

The table below lists the different RAID configurations in PS Series firmware v5.2.x and higher, options supported with PS Series arrays, and when a RAID policy is a best practice.

Table 2 RAID configurations and best practices

RAID Policy	Configurable in GUI	Configurable in CLI	Best practice for business critical data?
RAID 6	Yes	Yes	Yes
RAID 10	Yes	Yes	No for HDD or SSD larger than 6 TB.
RAID 50	Yes ¹	Yes	No for 7.2K disks 1 TB and larger
RAID 50	Yes ¹	Yes	Yes 10K, 15K, SATA or NL-SAS < 1 TB
RAID 5	No	Yes ²	No

¹ RAID 50 is not available using the GUI on PS6610 arrays

² RAID 5 is only available via the PS command line

Since the release of PS Series Firmware version 6.0, RAID 5 is no longer an option in the Group Manager GUI. In the PS Series, RAID 50 actually is preferred over RAID 5 because of the underlying disk set configuration as described in the Appendix C.1 and C.2.

In addition, RAID 50 is not preferred for arrays with 1 TB or larger HDD disks. This change is the result of the growing Class-2 HDD capacities and the degree of data protection that different RAID policies offer. RAID 6 and RAID 10 offer significantly higher levels of resiliency with very high capacity disks. This document shows examples of these differences and discusses options for choosing an appropriate RAID policy.

Note: RAID 6 is the Dell recommendation for PS Series arrays populated with SATA or NL-SAS disks 1 TB or larger and SSD or HDD 6 TB or larger. RAID 6 minimizes the risk of data loss during a RAID reconstruction thereby maximizing protection of the RAID set. Therefore, this recommendation is strongly encouraged for ALL customers when choosing a RAID policy for their arrays.

Existing configurations that use high capacity disks are encouraged when starting a migration plan that follows the best practices described in this document. Refer to “Converting or migrating from a member RAID policy” [Converting or Migrating](#) for details.

1.1 RAID policy descriptions

The following RAID policies are available using the PS Series Group Manager GUI or the CLI:

- **RAID 6** – One or more RAID 6 sets, with one or two spare disks. RAID 6 delivers multiple parity protection and provides the highest availability of any of the supported RAID policies because it can accommodate two drive failures at any one time. RAID 6 offers reliable performance for workloads consisting mainly of sequential reads and writes, or random reads. For workloads consisting mainly of random write operations, RAID 10 would be a better RAID policy choice due to the performance cost of extra writes to parity within the RAID 6 set during write operations. Unlike with the other RAID policies, a RAID 6 set can survive up to two drive failures at one time. The RAID set remains degraded until both failed disks are replaced and data reconstruction is complete.
- **RAID 10** – Striping on top of multiple RAID 1 (mirrored) sets, with one or two spare disks (depending on the total number of disks in the array). RAID 10 provides excellent performance for random writes, in addition to high availability. RAID 10 also offers the best performance when the RAID set is degraded. However, since the disks are mirrored, RAID 10 provides the least capacity and highest cost.
- **RAID 50** – Striping on top of multiple RAID 5 (distributed-parity) sets, with one or two spare disks (depending on the total number of disks in the array). RAID 50 provides a balance of performance, availability, and capacity. However, RAID 50 offers lower performance than RAID 10 when the RAID set is degraded and is not recommended for arrays with disks 1TB or greater.
- **RAID 5** - One or more RAID 5 sets, with at least one spare disk. RAID 5 is similar to RAID 50, with more capacity but significantly lower reliability and availability when compared with the other RAID types. RAID 5 also suffers more performance impact than RAID 10 and RAID 50 when the RAID set is degraded.

Note: RAID 5 should only be used for non-critical data. RAID 5 is only mentioned in this document to accommodate older systems.

For the purposes of this document, RAID 50 was used as a baseline to compare various aspects of the PS Series RAID options. RAID 50 was chosen because it has historically been popular with users due to its good balance of capacity and performance.

2 RAID policy availability and performance comparisons

Although all RAID levels provide good performance and data protection, there are some differences. When choosing a member RAID policy, the performance and availability needs of the workload should be identified to select a RAID policy that meets those needs. If the workload has mixed requirements in terms of performance and availability, consider mixing RAID types in a multi-member group.

There are many factors to consider when choosing a RAID level for a PS Series member, as shown in Table 3. The first column lists workload requirements, with the other columns listing the relative position of each requirement for each RAID policy.

Table 3 RAID policy comparison.

RAID	Reads (Random/Sequential)	Writes (Random/Sequential)	Relative Cost	Relative Protection	Rebuild Performance
10	Excellent/Excellent	Excellent/Good	Highest	Excellent	Best ¹
50	Excellent/Good	Good/Good	Moderate	Adequate	Good
6	Excellent/Good	OK/Good	Moderate	Best	OK
5	Excellent/Good	OK/Good	Lowest	Poor	OK

¹RAID 10 is not recommended for any SSD or HDD larger than 6 TB.

Note: RAID performance and availability is 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.

2.1 RAID policy disk failure protection and reliability

In principle, RAID policies 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.

A comparison of 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 – MTBF
- Controller RAID policy – member RAID setting
- RAID geometry – how RAID is constructed on the member

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. The graphic below depicts the relative reliability of RAID policies for the different disk options for PS Series arrays.




Figure 1 RAID reliability scores for different PS Series disk types

NOTE: The **y-axis** is the relative average reliability score for several PS Series disks. RAID 5 is deliberately left off this graph to reflect current PS Series best practices.

Key observation: RAID 6 demonstrates a much greater reliability than other RAID types. The extreme reliability of RAID 6 is due to its ability to protect against a double disk failure, where a read error occurs or a second disk fails during rebuild of failed disk. As disk counts and capacities have increased, RAID 6 has become a more important practice to ensure overall data protection in storage arrays.

The reliability scores in the graphic above represent relative reliability for a given RAID choice based on type of disks, the number of disks in the RAID set, rebuild times for the RAID set, and the possibility of a second disk (or 3rd if RAID 6) failing during a RAID rebuild operation.

SAN Headquarters v2.6 and later includes this intelligence to enable a higher level of awareness for customers that are concerned about RAID reliability and how they have their PS Series groups configured. SANHQ will measure the reliability score of an array from -5 to 5 and give it a rank of Excellent, Fair or Poor based on the score.

- Excellent = 1.19 or greater
- Fair = 0.x to 1.18
- Poor = 0 and below



Figure 2 SANHQ representation of array reliability

SAN Headquarters RAID Policy Reliability Score

The RAID policy reliability score in SAN Headquarters is calculated at the member level from the drive speed, size, disk mean time between failure (MTBF) hours, controller type, RAID policy, RAID geometry (number of disks in each RAID set), workload, and rebuild times. These factors result in a calculated score ranging from approximately -5.00 to $+5.00$. A score of 1.19 is the threshold that Dell recommends as a reliable RAID policy for the group. For example, the threshold score validates the maximum disk configuration in a RAID 50 policy for a single member. A RAID policy with a score higher than 1.19 increases that effectiveness.

Note: Remember to consider other aspects of the storage system, like capacity and performance, when choosing a reliable RAID policy.

For more information on using SANHQ and FAQs on RAID reliability see the "RAID Policy Evaluation and Reliability" section in the latest version of the *SAN Headquarters Installation and User's Guide* at <http://eqsupport.dell.com> then "Downloads->SAN Headquarters".

2.2 RAID policy impact on performance

A PS Series group provides storage to a broad range of environments and delivers good performance for a variety of workloads and applications, regardless of the RAID level on the member. However, for some applications, choosing the correct RAID level can make a difference in performance, under both normal operating conditions and failure conditions.

Determining which RAID level is best for an application depends on the following factors:

- Workload (heavy or light)
- How the application performs I/O (small random transfers or large sequential transfers)
- Read/write ratio
- Impact of degraded and reconstructing state

For example, video editing, streaming media, and disk-to-disk backup applications mainly perform large sequential transfers. Database and mail applications perform small random transfers. However, if multiple applications are accessing the storage, the total load may become random in nature, even if individual transfers are sequential.

An important tool in assessing performance of different RAID policies is a suite of benchmarks called “Four Corner” testing. The graphic below illustrates the four combinations of I/O characteristics highlighted above:

- 100% sequential writes
- 100% sequential reads
- 100% random writes
- 100% random reads

Note: The graphic below illustrates four corner testing results as compared to (or normalized to) RAID 50 performance for each corner. This graphic does not illustrate the differences between corners (such as random write versus sequential read), but rather just compares RAID performance within a corner (RAID 6 versus RAID 10 for random read). Absolute performance comparisons are discussed later.

While there are modest differences between RAID policies within a corner (aside from random writes), there can be significant differences in the absolute IOPS performance when comparing between corners, that is, random write to sequential read.

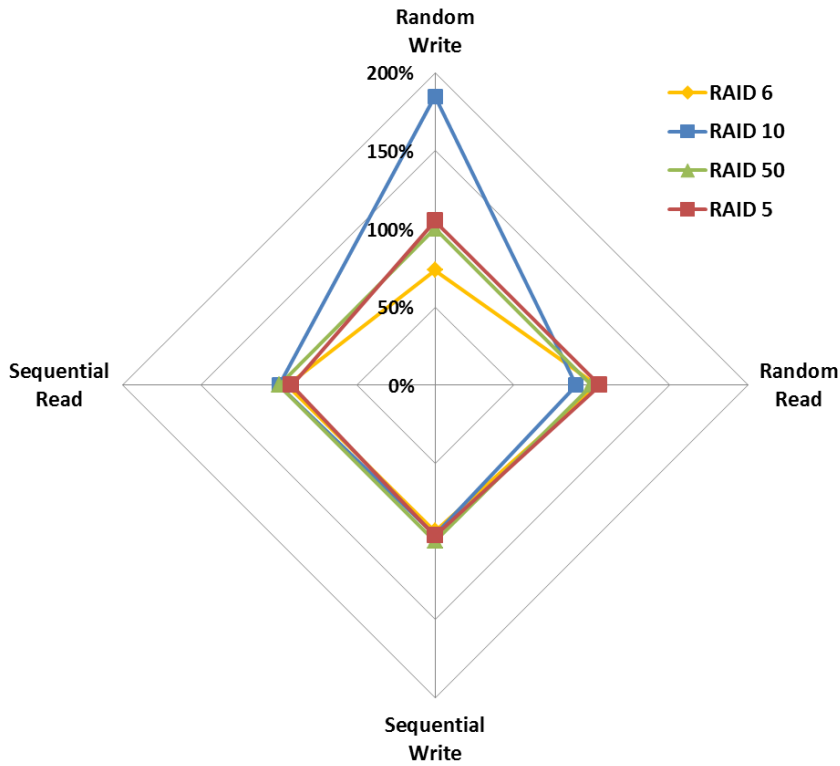


Figure 3 Corner testing results

Key Observation: As depicted by Figure 3, PS Series had considerable variation in performance amongst the RAID levels for random write I/O, as compared to the other three corners. Generally speaking, RAID 10 is very effective with applications that display small, random write intensive I/O, whereas RAID 6 pays a slight performance penalty for writing out the two parity data to two different disks.

The real world is more complicated, so a variety of workloads are benchmarked that reflect a variety of conditions. The graphic below depicts results for each RAID policy using a 50/50 mix of reads and writes for random and sequential I/O.

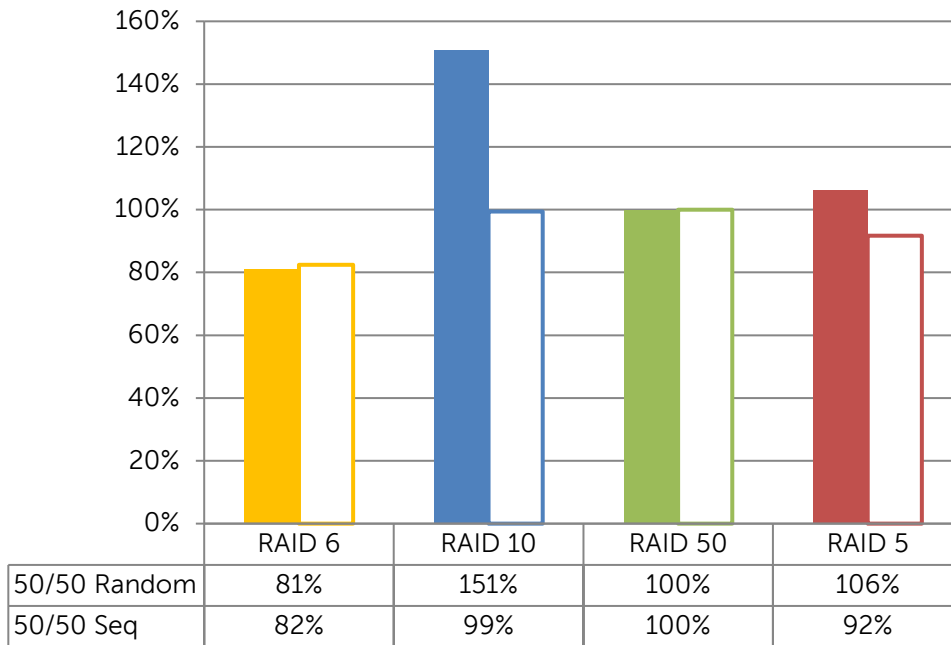


Figure 4 Relative performance by RAID policy

Key Observation: RAID 10 still offers the best random performance.

Another factor to consider, when comparing the performance of different RAID levels is the performance impact of an individual RAID set operating in a degraded mode.

Degraded: When a disk in a RAID set fails and there is no spare disk to replace it, the set is still functional, but degraded. Another disk failure (except in RAID 6) could result in a complete loss of data in the set.

- With RAID 6, there is moderate impact on read and write performance on a heavily loaded system. The impact difference is negligible compared to RAID 5 or 50, even with two parity calculations to reconstruct the data. However, RAID 6 can survive the simultaneous failure of any two disks in a set.
- With RAID 10, there is only minimal impact on read performance in the RAID 1 mirror set that has experienced the failure. There is negligible impact on the overall read performance.
- With RAID 5/RAID 50, there is moderate impact on read performance and random write performance, due to the parity reconstruction overhead.

Degraded Reconstructing: When a disk in a RAID set fails and a spare disk is available, the amount of time needed to reconstruct the data from parity information must be considered in the context of both heavy and light workloads.

- With RAID 6, data reconstruction is equal to RAID 5 or RAID 50 with a single disk failure. With a double disk failure, RAID 6 can involve slightly more overhead than a RAID 5 or RAID 50 because two parity calculations must be performed. This slight difference in overhead is outweighed by the additional level of protection RAID 6 offers. RAID 6 will continue to protect the RAID set should one occur whereas other RAID types will not.

- With RAID 10, data reconstruction involves only two disks (one active and one spare) out of the total number of disks in the array, so the overall performance impact is minimal.
- With RAID 5 and RAID 50, data reconstruction involves multiple disks, so the overall performance impact is significantly higher than other RAID types. For example, reconstruction in a six-disk RAID set involves six disks (five active and one spare). It is important to note that reconstruction time is directly impacted by the size of the disks.

2.3 Spare disks

Depending on the RAID policy and the total number of disks in each PS 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 another 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, 50, or 5 set that has a multi-disk failure event or a disk failure during a RAID rebuild operation.

2.3.1 If a spare disk is not available

- 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 loss of data.
- RAID 10, RAID 5, and RAID 50 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 loss of data.

Note: RAID 5 has the highest risk of data loss because it has the least protection of all the RAID choices during a rebuild operation.

When a disk in a RAID set fails, an array behaves in the following way:

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, RAID 5, or RAID 50 set: The member might be set offline, and any volumes and snapshots that have data located on the member are set offline.

If another disk fails in a degraded RAID 6 set: The member is not set offline but 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.

2.3.2 When a failed disk is replaced

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, performance returns to normal.

2.4 Examples of data protection with different RAID policies

The following table and figures show the results of a 24-disk, PS6210 series array responding to failures in the same RAID set based on the member RAID policy.

Note: The figures show a logical disk layout when an array is initialized for the first time. The actual physical layout of disks can change and evolve due to maintenance and administrative actions. Spare disks can move as they are utilized to replace failed disks and newly added disks become the spares. It is not possible to determine which physical disks are associated with each RAID set. This information is dynamic and maintained by the PS Series firmware.

Table 4 lists the RAID set relationship for each RAID policy in a 24-disk PS6210 series array. See Appendix B for PS65xx series arrays and Appendix C for PS6610 series arrays.

Table 4 PS Series PS6210 Series array RAID policy and disk set relationships

RAID policy	Spare disks	Disk RAID set relationship	Best practice
RAID 6	1 Spare Disk	(10+2) (9+2)	Yes
RAID 10	2 Spare Disks	(6+6) (5+5)	No for 6 TB and larger disks
RAID 50	2 Spare Disks	(5+1 5+1) (4+1 4+1)	For selected configurations ¹

¹ RAID 50 is not recommended for arrays with 1 TB or larger disks.

Note: The disk RAID set relationship indicates the number of data disks to parity or mirror disks ([data disk] + [parity|mirror disk]).

2.4.1 RAID 6

The following figures show how RAID 6 provides capacity and protection from single or dual disk failures for a 24-disk, PS6210 series array. RAID 6 provides higher availability because each RAID 6 set can survive two disk failures. This level of RAID is best used with larger physical disk sizes and for data that requires the highest level of protection and reliability.

In the following figures, the member RAID policy is RAID 6 with spares. There is only one spare in a 24-disk array with RAID 6. RAID and the spare disk protect the data. The second figure shows the reliability that RAID 6 offers by allowing a second disk failure protection in each RAID set.




Figure 5 Surviving disk failures with RAID 6 as the member RAID level




Figure 6 Example of RAID 6 surviving a double disk failure in each RAID set

2.4.2 RAID 10

The following figures show how using RAID 10 can increase (but cannot guarantee) the chance that an array can survive multiple disk failures in a 24-disk, PS6210 Series array. RAID 10 can survive only one simultaneous disk failure per RAID 1 mirror but allows multiple disk failures in a single array as long as the failed disks fall in different mirrored sets. This RAID level is appropriate for mission critical data with high write intensive random I/O requirements.

Figure 7 represents a RAID 10 configuration protected by RAID and two spare disks. In the event of a disk failure, the spare disks automatically replace the failed disks. After reconstruction, data is again protected by RAID.

Note: RAID 10 is not recommended for HDD or SSD larger than 6 TB.




Figure 7 Surviving disk failures with RAID 10 as the member RAID policy

2.4.3 RAID 50

The following figures show how RAID 50 provides more capacity than RAID 10 while continuing to provide protection from disk failure and good performance and availability for a 24-disk, PS6210 Series array. RAID 50 has less availability than RAID 10 due to fewer RAID sets that make up the overall RAID 50 policy. RAID 50 can survive only one simultaneous disk failure per RAID 5 set but allows multiple disk failures in a single array as long as the failed disks fall in different RAID 5 sets.

In the following image, the member RAID policy is RAID 50 with spares. The first image shows no disk failures. The second image shows two disk failures in different RAID 5 sets, and failures handled in the spare disks. Spare disks protect the RAID sets up to two disk failures.




Figure 8 Surviving disk failures with RAID 50 as the member RAID policy

Note: RAID 50 is not recommended for HDD larger than 1 TB.

2.5 Relative capacity utilization by RAID policy

Selection of the RAID policy impacts the capacity utilization of the RAID set in the array. Some RAID levels are more efficient and provide greater usable capacity for a given amount of raw capacity. The graph below illustrates the relative capacity utilizations of the various RAID policies supported by PS Series arrays. These metrics are normalized to RAID 50 (which is depicted as 100%).

Key observation: As illustrated, parity RAID policies (RAID 50, RAID 6, and RAID 5) offer comparable utilization, whereas RAID 10 with mirroring offers less capacity in comparison.




Figure 9 Relative capacity by RAID policy

3 Setting the member RAID policy

When a new group is created, or a new member is added to an existing group, the RAID policy must be selected before the space on the member is available.

Once the RAID policy is selected the disks are automatically configured according to the designated RAID level, with the appropriate number of spare disks.

Note: Starting with PS Series firmware v6.0.0, RAID 5 is not an available option in the GUI and RAID 50 is no longer preferred for arrays with disks 1 TB or greater regardless of the firmware level. For larger than 6 TB disks RAID 10 is not a recommended practice.

To set the member RAID policy in the GUI:

1. Click **Members** > *member_name*
2. A dialog box opens, warning that RAID is not configured on the member. Click **Yes** to configure the member.
3. In the **General Settings** dialog box, click **Next**.
4. In the RAID Configuration dialog box, select the RAID policy.
5. By default, the space on the member is available immediately, before the initialization completes and the batteries are fully charged. Select the option to wait until initialization completes. Click **Next**.
6. Review the summary information and click **Finish** to initialize the member.

To set the member RAID policy in the CLI, enter the following command:

```
> member select member_name raid-policy policy
```

Specify RAID 6, RAID 10, RAID 50 (if presented) for the *policy* variable.

4 Displaying the RAID level space distribution

To display the RAID level distribution for PS Series group members, open the Group Manager GUI. In the **Group Disk Space** panel, select the RAID level space distribution button.

A pie chart displays the RAID level distribution.

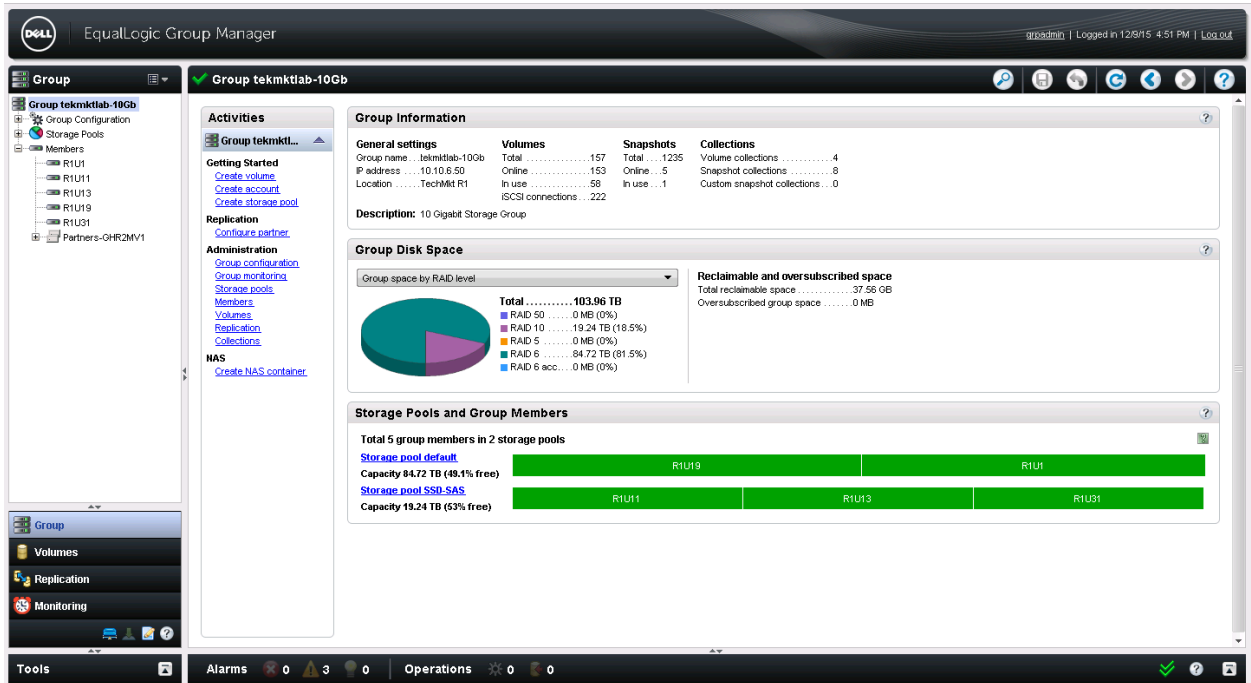


Figure 10 PS Series Group RAID level

5 PS Series array disk layout

To display the disks in a PS Series array, click:

Members > *member_name* > **Disks** tab

The **Member Disks** window displays a graphical representation of the disks in an array and also the status of each disk. The following figure shows an example of the member **Disks** window for a PS6610ES array with 14 x 400 GB SSD disks and 70 x 2 TB NLSAS disks, using RAID 6 accelerated, with one spare for each disk type.

The screenshot shows the Dell EqualLogic Group Manager interface. The main window is titled "Member R9U1 PS6610ES" and is in the "Disks" tab. It displays a "Disk Array Summary" grid and an "Installed Disks" table.

Disk Array Summary

0	1	2	3	4	5	6	7	8	9	10	11	12	13
14	15	16	17	18	19	20	21	22	23	24	25	26	27
28	29	30	31	32	33	34	35	36	37	38	39	40	41
42	43	44	45	46	47	48	49	50	51	52	53	54	55
56	57	58	59	60	61	62	63	64	65	66	67	68	69
70	71	72	73	74	75	76	77	78	79	80	81	82	83

Installed Disks

Total disks installed: 84

Slot	Type	Model	Revision	Size	Speed	Status	Errors	Health	Sector size
0	SAS SSD	LB406M	E327	372.61 GB	SSD	● spare	0	healthy	512 bytes
1	SAS SSD	LB406M	E332	372.61 GB	SSD	● online	0	healthy	512 bytes
2	SAS SSD	LB406M	E327	372.61 GB	SSD	● online	0	healthy	512 bytes
3	SAS SSD	LB406M	E332	372.61 GB	SSD	● online	0	healthy	512 bytes
4	SAS SSD	LB406M	E332	372.61 GB	SSD	● online	0	healthy	512 bytes
5	SAS SSD	LB406M	E332	372.61 GB	SSD	● online	0	healthy	512 bytes
6	SAS SSD	LB406M	E332	372.61 GB	SSD	● online	0	healthy	512 bytes
7	SAS SSD	LB406M	E332	372.61 GB	SSD	● online	0	healthy	512 bytes
8	SAS SSD	LB406M	E327	372.61 GB	SSD	● online	0	healthy	512 bytes
9	SAS SSD	LB406M	E327	372.61 GB	SSD	● online	0	healthy	512 bytes
10	SAS SSD	LB406M	E332	372.61 GB	SSD	● online	0	healthy	512 bytes
11	SAS SSD	LB406M	E327	372.61 GB	SSD	● online	0	healthy	512 bytes
12	SAS SSD	LB406M	E332	372.61 GB	SSD	● online	0	healthy	512 bytes
13	SAS SSD	LB406M	E332	372.61 GB	SSD	● online	0	healthy	512 bytes
14	SAS HDD	ST2000NM0023	GE0A	1.82 TB	7200 RPM	● online	0	healthy	512 bytes
15	SAS HDD	ST2000NM0023	GE0A	1.82 TB	7200 RPM	● online	0	healthy	512 bytes

Figure 11 Member Disks window

6 Converting or migrating from a member RAID policy

There are different methods available to convert or migrate from one RAID policy to another with PS Series arrays. These options can include:

- Online conversion
- Array evacuation
- Online volume migration
- Backup and restore

6.1 Large disk policy for RAID 50 and RAID 10

Beginning with version 7.0.0 firmware the Group Manager GUI and the CLI allow users to configure RAID 50 for members using disks with a capacity of less than 3 TB only. The CLI allows users to configure RAID 50 for members using disks with a capacity of 3 TB or greater, although this configuration is not recommended due to decreased RAID protection and reliability for large disk sizes.

For the PS models that support greater than 6 TB disks, RAID 10 is not recommended. RAID 6 should be used instead.

In addition, Dell recommends against using RAID configurations that do not use spare disks. You should convert all group members that are using a no-spare RAID policy to a policy that uses spare disks.

Table 5 Group Manager GUI RAID Policy rules

RAID Policy	Configurable in GUI	Best Practice for Business-Critical
RAID 6	Yes	Yes
RAID 10	Yes	No for HDD or SSD capacity greater than 6 TB.
RAID 50	Yes for disk capacity smaller than 3 TB	No for HDD disk capacity of 1 TB and larger.
RAID 5	No	No

6.2 Online conversion

With an online conversion, the migration occurs while the member remains online (in the group) if the new RAID policy provides the same or more space as the current policy. Conversions can take a long time and should be scheduled during off-peak times or when the workload is light.

The following online conversions are available in the Group Manager GUI and the CLI:

- RAID 10 to RAID 50, or RAID 6
- RAID 50 to RAID 6

Note: RAID 5 and RAID 6 have no online conversion options. The only way to convert RAID 5 or RAID 6 to a different RAID type is to reset the array and initialize with a different RAID policy.

Migration using Group Manager

To convert the member RAID policy, open the Group Manager GUI and click:

Members > *member_name* > **Modify RAID configuration**

The **Modify RAID Configuration** wizard displays the current RAID policy radio button selected.

Note: If a certain conversion option is not allowed, it will appear grayed out.

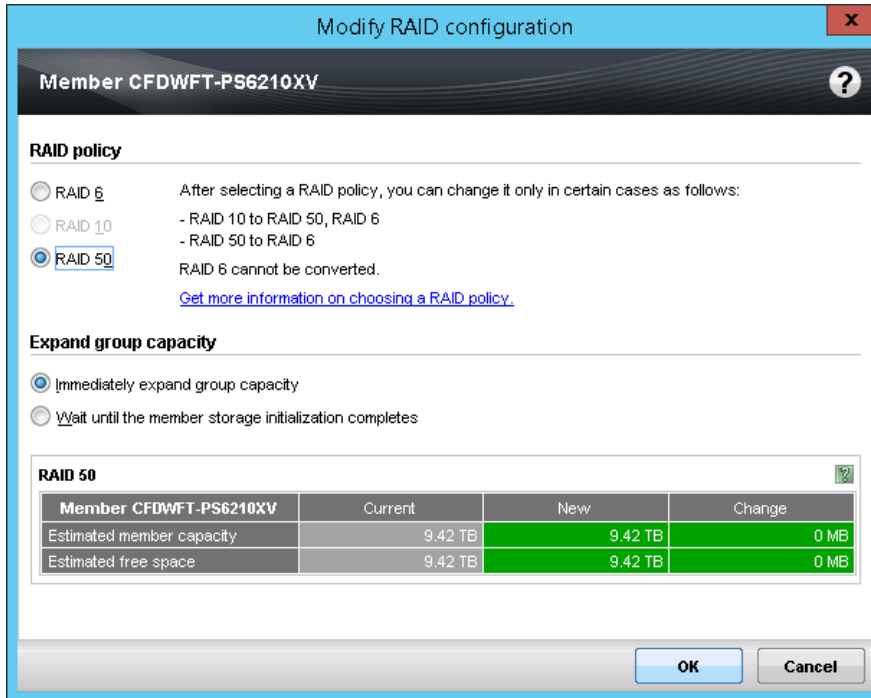


Figure 12 Modify RAID configuration wizard

When a new RAID policy is selected, the title of the table changes to the new RAID policy, also the capacity will reflect the new capacity as shown in the following figure.

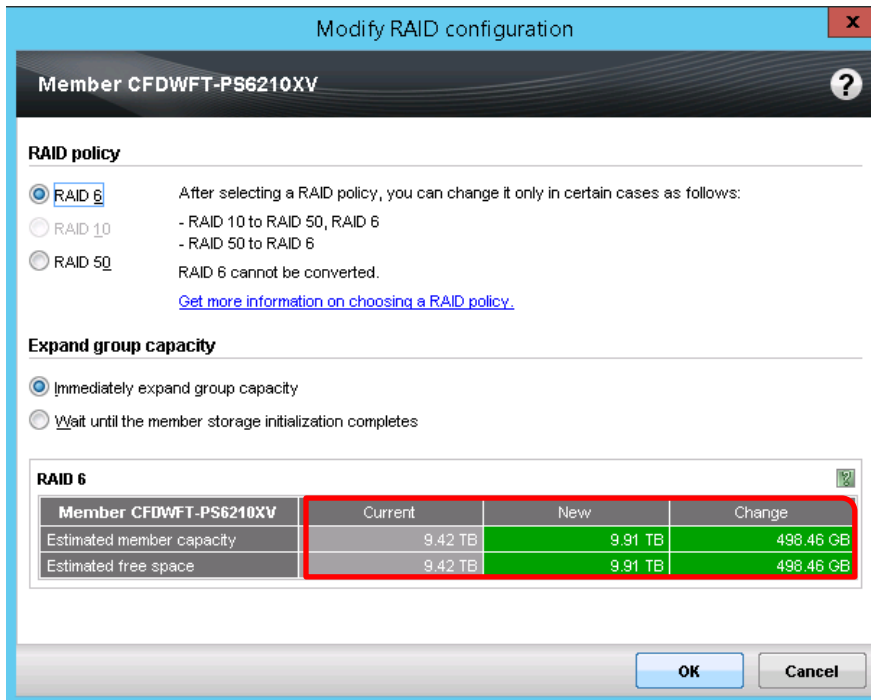


Figure 13 Selecting RAID policy

6.3 Array evacuation

If the current RAID policy does not support an online conversion and the group has other members with sufficient free space, a member can be converted to any other RAID policy by removing the member from the group (which resets the array) and adding it back to the group, then selecting a different RAID policy.

The following RAID policies must use an offline migration method:

- RAID 5 to RAID 6 or RAID 10
- RAID 6 to RAID 10
- RAID 50 to RAID 10

6.4 Online volume migration

Another option is migrating volumes to another RAID type. There are three ways to move individual volumes to specific RAID types.

- Move volume
- RAID preference designation
- PS Series replication

6.4.1 Move volume

The Move volume method requires an array that is already configured with the RAID policy and enough free space to host the volume data.

Once those criteria have been met, the volume can be migrated to another storage pool by choosing Move volume option in the Group Manager GUI and selecting the desired storage pool.

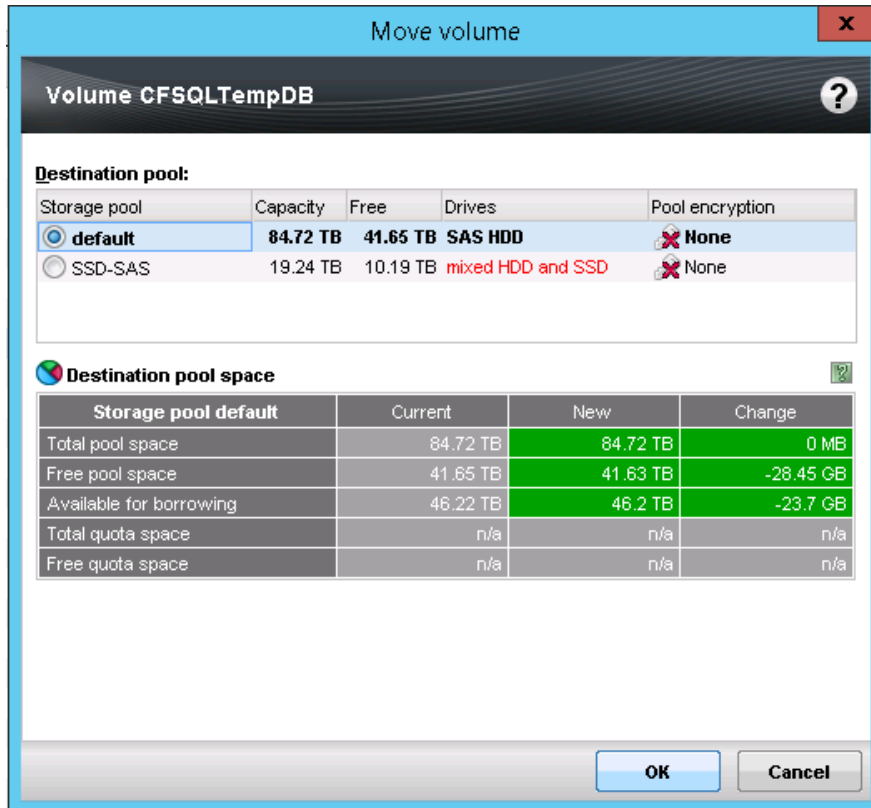


Figure 14 Move volume window

6.4.2 RAID preference

The other migration option is to set the RAID preference for the volume. This can be done using the **Modify Settings** option for the volume. Once in **Modify Settings**, choose the **Advanced** tab and set the RAID preference for the volume. The volume will be moved to the appropriate RAID policy if available.

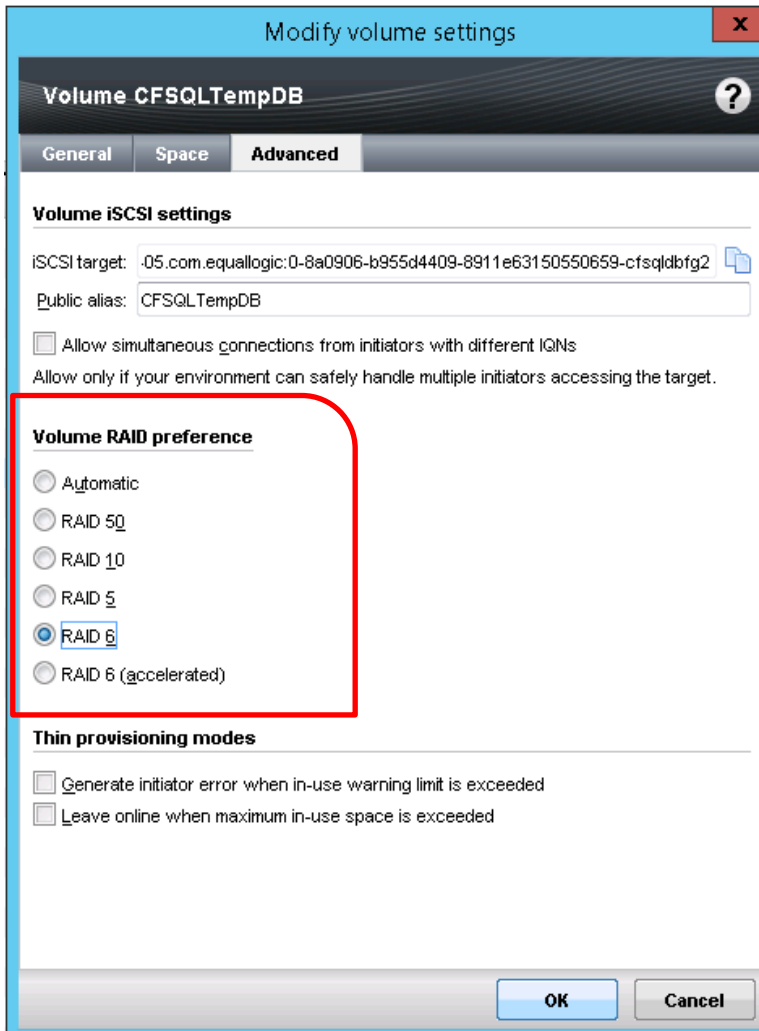


Figure 15 Modify volume settings window

6.4.3 PS Series replication

PS Series replication can be used to migrate volume data. Volumes can be migrated using either standard asynchronous replication or in firmware version 6.x and higher synchronous replication.

Asynchronous replication moves data between PS Series groups while synchronous replication moved data between storage pools within the same PS Series group. For more information regarding replication options and setup, refer to the *PS Series Group Administration guide* <http://eqsupport.dell.com> then “Downloads -> PS Firmware”.

6.5 Backup and restore

Lastly, traditional backup and restore options can also be used to migrate data from one array to another with a different RAID policy. These methods usually require downtime and can be lengthy depending on the size of the data being moved.

7 Summary

Understanding and using proper RAID levels can be beneficial from a performance perspective and crucial for protection against failures in business critical environments. Dell recommends choosing RAID policies based on evolving storage technologies and protecting storage needs of customers. PS Series arrays allow administrators the flexibility to choose and modify RAID configurations that best meet the needs of their data center applications.

A 12-Disk PS Series array RAID configurations

Table 6 PS4210, PS4100 and PS4110 12-disk RAID Set Relationship

RAID Policy	Spare Disks	RAID Set Relationship	Best Practice
RAID 6	1 Spare Disk	(9+2)	Yes
RAID 10	2 Spare Disks	(5+5)	No for HDD or SSD greater than 6 TB.
RAID 50	2 Spare Disks	(4+1)(4+1)	For selected configurations ¹
RAID 5	1 Spare Disk	(10+1)	Not for business critical data

¹ RAID 50 is not recommended for arrays with 1 TB or larger disks. The Dell Group Management GUI will not allow for any members with 3 TB or higher to be configured with RAID 50 (firmware 7.0.0 and higher)

B 14-Disk PS Series array RAID configurations

The following table shows the blade version of the PS Series arrays.

Table 7 PS-M4100 and PS-M4110 14-disk RAID Set Relationship

RAID Policy	Spare Disks	RAID Set Relationship	Best Practice
RAID 6	1 Spare Disk	(11+2)	Yes
RAID 6 Accelerated	1 Spare HDD	(6+2 HDD) (6+2 SSD)	N/A
RAID 10	2 Spare Disks	(6+6)	No for HDD or SSD greater than 6 TB.
RAID 50	2 Spare Disks	(5+1)(5+1)	For selected configurations ¹
RAID 5	1 Spare Disk	(12+1)	Not for business critical data

¹ RAID 50 is not recommended for arrays with 1 TB or larger disks. The Dell Group Management GUI will not allow for any members with 3 TB or higher to be configured with RAID 50 (firmware 7.0.0 and higher)

C 16-Disk PS Series array RAID configurations

The following figures shows the response of a 16-disk, PS6000 Series array to failures **in the same RAID set**, based on the member RAID policy.

The figures show a logical disk layout. The actual physical layout of disks can change due to administrative actions. It is not possible to determine which disks are associated with each RAID set.

Note: Adhere to the recommendations in the main section of this document when choosing a member RAID policy

RAID 10

The following figures shows that using RAID 10 can increase (but cannot guarantee) the chance of an array surviving multiple disk failures in a 16-disk, PS6000 array. This RAID level is appropriate for mission critical data with high write intensive random I/O requirements.

In Figure 16, the member RAID policy is RAID 10 with spares. Spares and RAID protect the data.

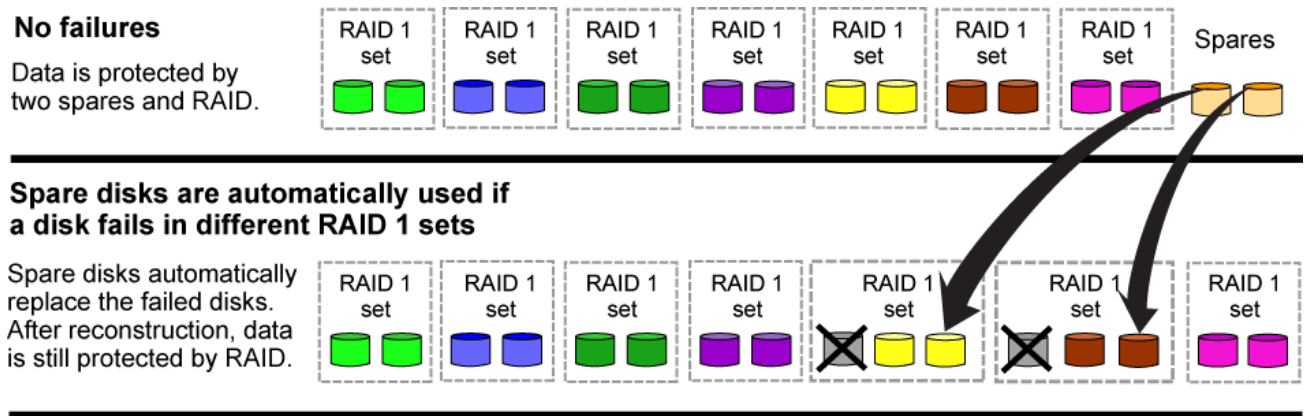


Figure 16 Surviving disk failures with RAID 10 as the member RAID policy

C.1 RAID 50


The following figures shows that RAID 50 provides more capacity than RAID 10 while continuing to provide protection from disk failure and good performance and availability for a 16-disk, PS6000 array.

RAID 50 is not recommended for arrays with 1 TB or larger disks. The Dell Group Management GUI will not allow for any members with 3 TB or higher to be configured with RAID 50 (firmware 7.0.0 and higher)

In Figure 17, the member RAID policy is RAID 50 with spares. Spares and RAID protect the data.

No failures

Disks are protected by two spares and RAID.



Spare disks are automatically used if a disk fails in each RAID 5 set

Spare disks automatically replace the failed disks. After reconstruction, data is still protected by RAID.




Figure 17 Surviving disk failures with RAID 50 as the member RAID policy


C.2 RAID 5

Figure 18 shows that RAID 5 provides capacity and protection from disk failure but less availability for a 16-disk, PS6000 array. RAID 5 provides less availability because there is only one set and one spare. This level of RAID is not recommended for any production PS Series arrays.

In the following figure, the member RAID policy is RAID 5 with spares. Spares and RAID protect the data.

No failures

Disks are protected by one spare and RAID.



Spare disk is automatically used if a disk fails in the RAID 5 set

The spare disk automatically replaces the failed disk. After reconstruction, data is still protected by RAID.




Figure 18 Surviving disk failures with RAID 5 as the member RAID level

C.3 RAID 6

Figure 19 shows that RAID 6 provides capacity and protection from single or dual disk failures for a 16-disk, PS6000 array. RAID 6 provides less availability because there is only one set and two spares. This level of RAID is best used with larger physical disk sizes and for data that requires the highest level of protection and reliability.

In the following figure, the member RAID policy is RAID 6 with spares. There is only one spare in a 16-disk array with RAID 6. RAID and the spare disk protect the data. Even after the spare is used, the RAID set can survive the failure of any two disks at the same time.


No failures

Disks are protected by one spare and RAID.



Spare disk is automatically used if a disk fails in the RAID 6 set

Spare disk automatically replaces the failed disk. After reconstruction, data is still protected by RAID.



RAID protects data if two disks fail in the RAID 6 set

The RAID set is degraded, but the array is still functional.



Figure 19 Surviving disk failures with RAID 6 as the member RAID level

D 48-Disk PS Series array RAID configurations

Table 8 RAID set relationships for a 48 disk PS Series 65xx array

RAID Policy	Spare Disks	RAID Set Relationship	Best Practice
RAID 6	1 Spare Disk	(12+2) (12+2) (12+2) (3+2)	Yes
RAID 6 Accelerated	1 Spare HDD ²	(12+2 HDD) (12+2 HDD) (10+2 HDD) (5+2 SSD)	N/A ¹
RAID 10	2 Spare Disks	(7+7) (7+7) (7+7) (2+2)	No for HDD or SSD greater than 6 TB.
RAID 50	2 Spare Disks	(6+1, 6+1) (6+1, 6+1) (6+1, 6+1) (3+1)	For selected configurations ²
RAID 5	2 Spare Disk	(12+1) (12+1) (12+1) (6+1)	Not for business critical data

¹ RAID 6 Accelerated is the only available option for hybrid (mix of HDD and SSD) arrays.

² Firmware v8.0 and higher provides an optional PS Series CLI enabled feature which will allow for replacement of the failed SSD drive prior to automatically rebuilding the RAID. For more information see “RAID Rebuild Delay” in the PS Series Firmware release notes at <https://eqsupport.dell.com> then “Downloads -> PS Firmware”.

³ RAID 50 is not recommended for arrays with 1 TB or larger disks. The Dell Group Management GUI will not allow for any members with 3 TB or higher to be configured with RAID 50 (firmware 7.0.0 and higher)

E 42 or 84-disk PS Series RAID configurations

Dell Storage PS6610 arrays are offered in 84 or 42 drive configurations. The hybrid version of the PS6610 is only available in the 84 drive configuration. The following tables show the RAID relationships for these configurations.

Table 9 RAID set relationship for an 84 drive PS Series 6610 array.

RAID Policy	Spare Disks	RAID Set Relationship	Best Practice
RAID 10	2 Spare Disk	(7+7) (7+7) (7+7) (7+7) (7+7) (6+6)	R10- No for HDD or SSD greater than 6 TB.
RAID 6	2 Spare Disk	(12+2) (12+2) (12+2) (12+2) (12+2) (10+2)	Yes
RAID 6 Accelerated	1 Spare HDD 1 Spare SSD	(12+2 HDD) (12+2 HDD) (12+2 HDD) (12+2 HDD) (11+2 HDD) (11+2 SSD)	N/A ¹
RAID 50 RAID 5	2 Spare Disks	(12+2) (12+2) (12+2) (12+2) (12+2) (10+2)	Not for business critical data ²

¹ RAID 6 Accelerated is the only available option for hybrid (mix of HDD and SSD) arrays.

² RAID 5 and RAID 50 are not available via the Group Manager UI on PS6610 arrays.

Table 10 RAID set relationship for a 42 drive PS Series 6610 array.

RAID Policy	Spare Disks	RAID Set Relationship	Best Practice
RAID 10	2 Spare Disk	(7+7) (7+7) (6+6)	R10 – No for HDD or SSD greater than 6 TB.
RAID 6	2 Spare Disk	(12+2) (12+2) (10+2)	Yes
RAID 50 RAID 5	2 Spare Disks	(12+2) (12+2) (10+2)	Not for business critical data ¹

¹ RAID 5 and RAID 50 are not available via the Group Manager UI on PS6610 arrays.

F Technical support and resources

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

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

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell Storage platforms.

For an updated *Dell Storage Compatibility Matrix* visit the following URL:

<http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/20438558>.