



EQUALLOGIC

TECHNICAL SUPPORT INSTRUCTIONAL SERIES

Technical Report	Document ID: TR1020
PS Series Storage Arrays: Choosing a Member RAID Policy	Revision: A05
	Revision Date: August 28, 2009

TECHNICAL REPORT

PS SERIES STORAGE ARRAYS CHOOSING A MEMBER RAID POLICY

ABSTRACT

This Technical Report describes the supported RAID policies for PS Series group members. This report also provides information on the performance characteristics of the RAID policies under normal and failure conditions and will help you choose the RAID policy that is right for your applications and environment.



TR1020

V4.1

Copyright © 2009 Dell Inc. All Rights Reserved.

Dell EqualLogic™ is a trademark of Dell Inc.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

Possession, use, or copying of the documentation or the software described in this publication is authorized only under the license agreement specified herein.

Dell, Inc. will not be held liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change.

April 2009

PS Series Firmware Version 4.1

Dell, Inc.
300 Innovative Way
Suite 301
Nashua, NH 03062 USA
Tel: 603.579.9762
Fax: 603.579.6910

TABLE OF CONTENTS

Introduction to Member RAID Policies	3
RAID Policy Descriptions	3
RAID Policy Comparisons	4
RAID Level Performance Characteristics	5
RAID Policy Drive Failure Protection.....	6
Examples of Data Protection with Different RAID Policies	7
Setting the Member RAID Policy	11
Displaying the RAID Level Space Distribution.....	12
PS Series Array Drive Layout.....	13
Converting a Member RAID Policy	14
Documentation and Customer Support	15
Contacting Dell	15

INTRODUCTION TO MEMBER RAID POLICIES

PS Series storage arrays include redundant, hot-swappable components—for example physical drives, control modules, fans, and power supplies—for a no-single-point-of-failure configuration. Components fail over automatically without user intervention or disrupting data availability.

PS Series storage arrays protect data by using RAID technology and spare drives. When you add a member to a PS Series group, you choose the RAID policy for that member.

The RAID policy for a member consists of the following:

- RAID level (RAID 50, RAID 10, RAID 5, or RAID 6)
- Whether or not spare drives are automatically configured and used

After adding an array to a group, the storage in the member will be available once you set the RAID policy. The drives are automatically configured according to the designated RAID level, with the appropriate number of spare drives.

RAID Policy Descriptions

You can choose from the following RAID policies using the Group Manager GUI and the CLI:

- **RAID 10.** Striping on top of multiple RAID 1 (mirrored) sets, with one or two spare drives (depending on the total number of drives in the array). RAID 10 provides good 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 drives are mirrored, RAID 10 provides the least capacity.
- **RAID 50.** Striping on top of multiple RAID 5 (distributed-parity) sets, with one or two spare drives (depending on the total number of drives 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.
- **RAID 5.** One or more RAID 5 sets, with at least one spare drive. RAID 5 is similar to RAID 50, with more capacity (two additional drives) but lower availability and random write performance 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.
- **RAID 6.** One or more RAID 6 sets, with one or two spare drives. Delivers multiple parity protection and provides the highest availability of any of the supported RAID policies because it can survive the simultaneous failure of any two drives in a RAID set. Offers reliable performance for workloads consisting mainly of sequential reads and writes, or random reads. Not recommended for workloads consisting mainly of random writes. Unlike with the other RAID policies, when a RAID 6 set is degraded (two simultaneous drive failures), the member remains online, but performance is slower. The RAID set remains degraded until both failed drives are replaced and data reconstruction is complete.

Using the Group Manager CLI, you can specify these additional RAID policies:

- **RAID 50, no spares.** Striping on top of two RAID 5 sets, with no spares, if possible.
- **RAID 10, no spares.** Striping on top of multiple RAID 1 sets, with no spares, if possible.
- **RAID 6, no spares.** One or more RAID 6 sets with no spare drives, if possible.

Note: You cannot specify (or convert to) a no-spares policy on a PS5500 or PS6500.

A no-spare RAID policy can leave the member vulnerable to data loss. If you select a no-spare RAID policy, Dell strongly recommends that you have on-site spare drives and 7x24 personnel for monitoring the member and replacing failed drives in a timely fashion.

RAID Policy Comparisons

Although all RAID levels provide good performance and data protection, there are some differences. When choosing a member RAID policy, you should identify the performance and availability needs of your workload and select a RAID policy that meets those needs. If your workload has mixed requirements in terms of performance and availability, you may want to 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 1. Both RAID 10 and RAID 50 provide excellent reliability for your data, in addition to overall high performance. RAID 6 provides high storage capacity, but at the expense of performance impact during data reconstruction. RAID 10 provides the best performance for workloads that are mostly small random writes.

Table 1: RAID Level Performance and Availability Comparison

Workload Requirement	RAID 10	RAID 50	RAID 5	RAID 6
Capacity	OK	Good	Excellent	Good
Availability	Excellent	Good	OK	Excellent
Sequential reads	Excellent	Excellent	Excellent	Excellent
Sequential writes	Good	Good	Good	Good
Random reads	Excellent	Excellent	Excellent	Excellent
Random writes	Excellent	Good	OK	OK
Performance impact of drive failure/RAID reconstruction	Minimal	Moderate	Moderate to Heavy	Heavy

RAID Level Performance Characteristics

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, the right member 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

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.

When comparing performance RAID level characteristics, you must understand the performance impact when an individual RAID set is:

- **Reconstructing.** When a drive in a RAID set fails and a spare drive 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 10, data reconstruction involves only two drives (one active and one spare) out of the total number of drives in the array, so the overall performance impact is minimal.
 - With RAID 5/RAID 50, data reconstruction involves multiple drives, so the overall performance impact may be significant. For example, if you have a six-drive set, reconstruction involves six drives (five active and one spare). However, under a normal workload, reconstruction may not impair overall performance.
 - With RAID 6, data reconstruction involves more overhead than a RAID 5 reconstruction because two parity calculations must be performed.
- **Degraded.** When a drive in a RAID set fails and there is no spare drive to replace it, the set is still functional, but degraded. Another drive failure (except in RAID 6) could result in a complete loss of data in the set.
 - With RAID 10, there is minimal impact on read performance *only* in the set that has experienced the failure. There is negligible impact on the overall read performance.
 - With RAID 5/RAID50, there is moderate impact on read performance and random write performance, due to the parity reconstruction overhead.
 - With RAID 6, there can be significant impact on read and write performance on a heavily loaded system, because two parity calculations must be performed to reconstruct the data. However, RAID 6 can survive the simultaneous failure of any two drives in a set.

RAID Policy Drive Failure Protection

It is important to consider the drive failure protection a RAID policy provides.

Depending on the RAID policy and the total number of drives in each PS Series storage array, one or more spare drives are automatically configured and used in the event of a drive failure. If a spare drive is *not* available:

- RAID 10, RAID 5, and RAID 50 are guaranteed to survive one drive failure. Data continues to be available, but the set is degraded. A second drive failure in a degraded set can result in loss of data.
- RAID 6 is guaranteed to survive the simultaneous failure of any two drives. Data continues to be available, but the set is degraded. A third drive failure in a degraded set can result in loss of data.

When a drive in a RAID set fails, an array behaves as follows:

- **If a spare drive is available:** The spare automatically replaces the failed drive. Data from the failed drive is reconstructed on the spare and continues to be available. During reconstruction, the set that contains the failed drive is temporarily degraded. After reconstruction, performance returns to normal.
- **If no spare drive is available:** Data continues to be available, but the set is degraded.
- **If another drive 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 also set offline.
- **If another drive fails in a degraded RAID 6 set:** The member is not set offline but continues to be degraded until the failed drives are replaced. Data is reconstructed on the replacement drives. If the first reconstruction is still underway, performance can be further reduced. After both drives are reconstructed, performance returns to normal.

When you replace a failed drive:

- **If a spare drive was used:** Data has already been reconstructed on the spare drive, so the new drive becomes a spare.
- **If a set is degraded:** Data is reconstructed on the new drive and after reconstruction, performance returns to normal.

Examples of Data Protection with Different RAID Policies

The following figures show how a 16-drive, PS5000 Series array responds to failures **in the same RAID set**, based on the member RAID policy.

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

RAID 10

The following figures show how using RAID 10 can increase (but cannot guarantee) the chance that an array can survive multiple drive failures in a 16-drive, PS5000 array. This RAID level is appropriate for mission-critical data that requires the highest level of protection and rapid accessibility.

In the following figure, the member RAID policy is RAID 10 with spares. Spares and RAID protect the data.

Surviving Drive Failures with RAID 10 as the Member RAID Policy


No failures

Data is protected by two spares and RAID.



Spare disks are automatically used if a disk fails in different RAID 1 sets

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



In the following figure, the member RAID policy is RAID 10 with no spares. RAID protects the data.

Surviving Drive Failures with RAID 10-No-Spares as the Member RAID Policy


No failures

Data is protected by RAID.



RAID protects data if a disk fails in each RAID 1 set (and spare disks are not available)

The RAID 1 sets are degraded, but the array is still functional.



RAID 50


The following figures show how RAID 50 provides more capacity than RAID 10 while continuing to provide protection from drive failure and good performance and availability for a 16-drive, PS5000 array.

In the following figure, the member RAID policy is RAID 50 with spares. Spares and RAID protect the data.

Surviving Drive Failures with RAID 50 as the Member RAID Policy


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.




In the following figure, the member RAID policy is RAID 50 with no spares. RAID protects the data.

Surviving Drive Failures with RAID 50-No Spares as the Member RAID Policy


No failures

Disks are protected by RAID.



RAID protects data if a disk fails in each RAID 5 set (and spare disks are not available)

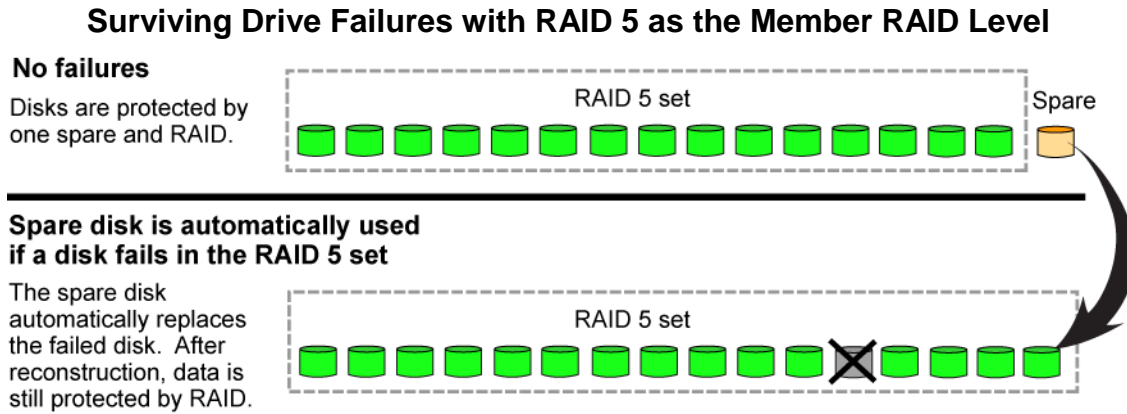
The RAID 5 sets are degraded, but the array is still functional.



RAID 5

The following figure shows how RAID 5 provides capacity and protection from drive failure but less availability for a 16-drive, PS5000 array. RAID 5 provides less availability because there is only one set and one spare. This level of RAID is best used for disaster recovery replicas, backups, and other types of data that is not accessed frequently.

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



Note: RAID 5 with no spares is not supported.

RAID 6

The following figures show how RAID 6 provides capacity and protection from *single or dual* drive failures for a 16-drive, PS5000 array. RAID 6 provides less availability because there is only one set and two spares. This level of RAID is best used for data that requires an additional level of data protection.

In the following figure, the member RAID policy is RAID 6 with spares. (There is only 1 spare in a 16-drive array with RAID 6.) RAID and the spare drive protect the data. Even after the spare is used, the RAID set can survive the failure of any two drives at the same time.

Surviving Drive Failures with RAID 6 as the Member RAID Level

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.



In the following figure, the member RAID policy is RAID 6 with no spares. RAID protects the data.

Surviving Drive Failures with RAID 6-No Spares as the Member RAID Level


No failures

Disks are 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.



SETTING THE MEMBER RAID POLICY

When you create a new group, or add a member to an existing group, you must set the RAID policy before the space on the member is available.

To set the member RAID policy in the GUI:

1. Click *Members* → *member_name*
2. A dialog box opens, warning you 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. However, you can choose 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
```

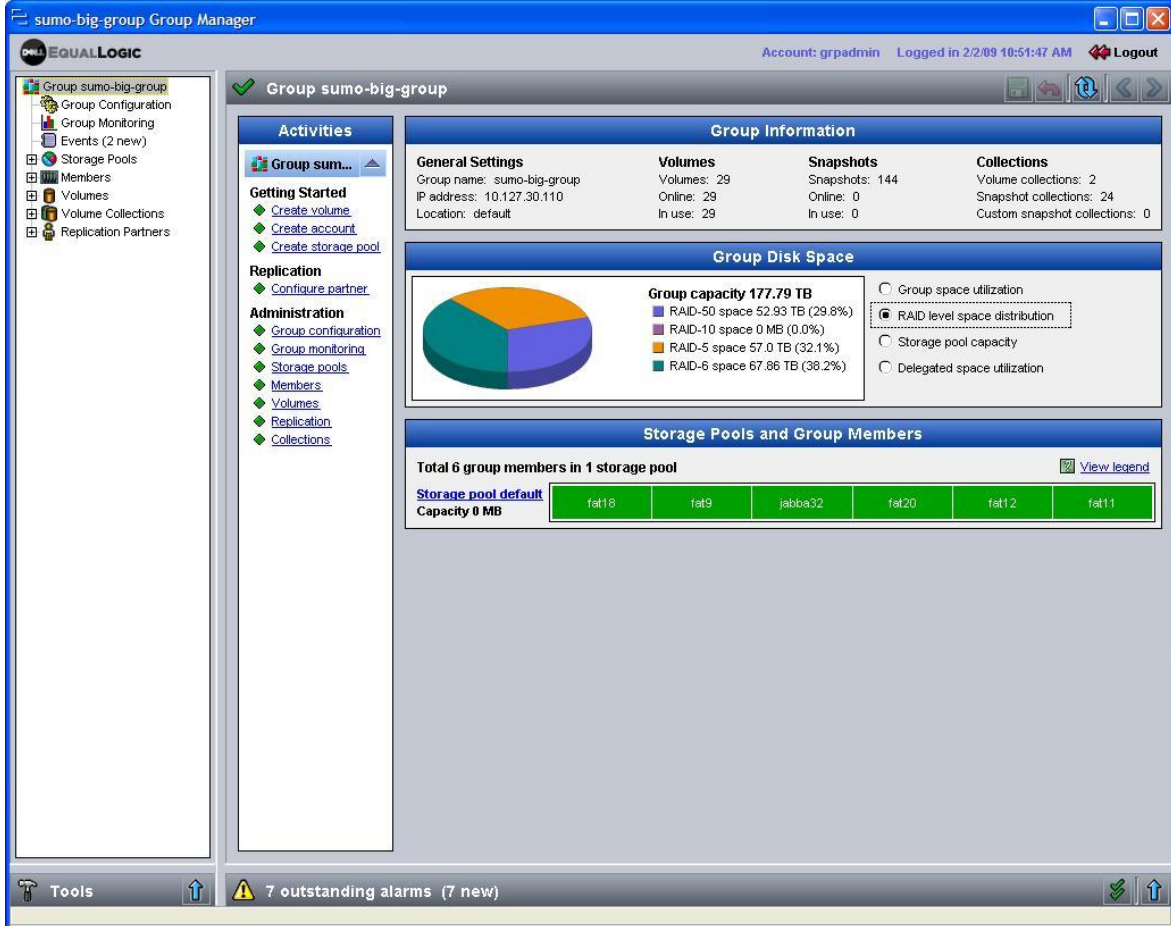
Specify *raid5*, *raid6*, *raid10*, or *raid50* for the *policy* variable.

Note: You can also optionally specify *raid6-nospares*, *raid10-nospares*, or *raid50-nospares*. However, you cannot specify (or convert to) a no-spares policy on a PS5500 or PS6500.

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 displaying the RAID level distribution appears.

PS Series Group RAID Level



PS SERIES ARRAY DRIVE LAYOUT

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

Members → *member_name* → Disks tab

The Member Disks window appears, displaying a graphical representation of the drives in an array and also the status of each drive. The following figure shows an example of the Member Disks window for a PS5500 array with 48 drives, using RAID 6, with one spare drive.

Member Disks Window

The screenshot shows the 'Member Disks' window for member 'fat18'. The interface includes a navigation pane on the left, a central menu with tabs for Status, Enclosure, Controllers, Disks, Network, Connections, and Service. The 'Disks' tab is active, displaying a 'Disk Array Summary' grid and an 'Installed Disks' table.

Disk Array Summary

Row	Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	Col 7	Col 8	Col 9	Col 10	Col 11
3	7	11	15	19	23	27	31	35	39	43	47
2	6	10	14	18	22	26	30	34	38	42	46
1	5	9	13	17	21	25	29	33	37	41	45
0	4	8	12	16	20	24	28	32	36	40	44

Installed Disks

Total disks installed: 48

Slot	Type	Model	Revision	Size	Status	Errors
30	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
31	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
32	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
33	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
34	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
35	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
36	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
37	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
38	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
39	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
40	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
41	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
42	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
43	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
44	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
45	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
46	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	online	0
47	SATA (HDD)	Hitachi HUA721010KLA330	GKAOAB0A	931.51 GB	spare	0

CONVERTING A MEMBER RAID POLICY

Once you select a RAID policy for a member you can convert it to a different RAID policy while the member remains online (in the group) only if the new RAID policy provides the same or more space as the current policy. Conversions can take a long time, so you may want to schedule it when the workload is light.

However, if the group has other members with sufficient free space, you can convert a member 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 online conversions are available in the Group Manager GUI and the CLI:

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

The following online conversions are available through the CLI only:

- RAID 10, no-spares to RAID 50, no-spares, RAID 5, no-spares, or RAID 6, no-spares
- RAID 50, no-spares to RAID 5, no-spares

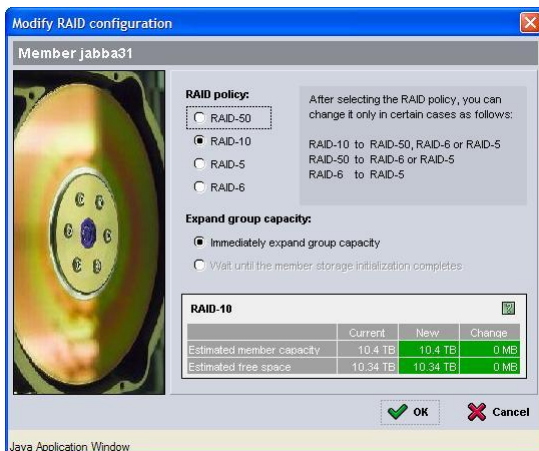
To convert the member RAID policy, open the Group Manager GUI and click:

Members → *member_name* → Modify RAID configuration

The Modify RAID Configuration wizard appears with the current RAID policy radio button selected.

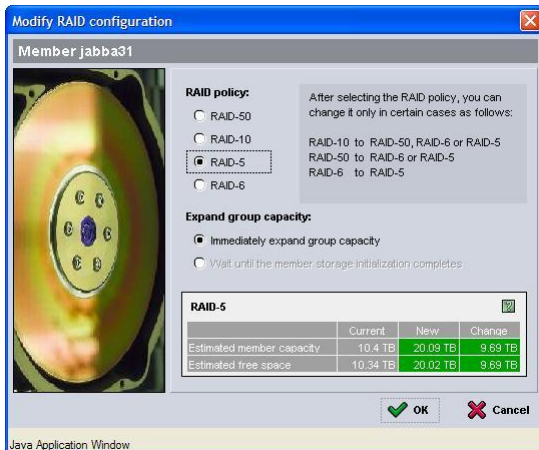
Note: If a certain conversion option is not allowed, it will appear grayed out. (You must use the CLI to change to or from a no-spares policy.)

Modify RAID Configuration Wizard



When you select the new RAID policy, the title of the table and the capacity listed in the table will change as shown in the following figure.

RAID Configuration Wizard Reflecting Changes



When you are finished, click **OK**.

DOCUMENTATION AND CUSTOMER SUPPORT

For detailed information about PS Series storage arrays, groups, and volumes, see the following documentation:

- *Release Notes*. Provides the latest information about PS Series storage arrays.
- *Installation and Setup*. Describes how to set up the storage array hardware and create a PS Series group. Use the manual for your array model.
- *Group Administration*. Describes how to use the Group Manager graphical user interface (GUI) to manage a PS Series group. This manual provides comprehensive information about product concepts and procedures.
- *CLI Reference*. Describes how to use the Group Manager command line interface (CLI) to manage a PS Series group and individual arrays.
- *Hardware Maintenance*. Provides information about maintaining the storage array hardware. Use the manual for your array model.

If you have a valid service contract or an array that is still under the first year of warranty, you can download documentation and firmware, search the Solutions database, and view FAQs and Technical Reports. You can also submit a support request (Log a Case) from the website. To access the customer support website:

- Open a browser and go to support.dell.com/equallogic, click **Login** and log in to a support account.
- If you do not have an account, create one by clicking **Request a Customer Account** under the login prompt.

CONTACTING DELL

For customers in the United States, call 800-945-DELL (800-945-3355).

Note: If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog.

Dell provides several online and telephone-based support and service options. Availability varies by country and product, and some services may not be available in your area. To contact Dell for sales, technical support, or customer service issues:

1. Visit support.dell.com.
2. Verify your country or region in the **Choose A Country/Region** drop-down menu at the bottom of the page.
3. Click **Contact Us** on the left side of the page.
4. Select the appropriate service or support link based on your need.
5. Choose the method of contacting Dell that is convenient for you.

Document Control and History Information

Revision	Change Author/Organization	Reason	Approved By/Organization	Date
A05	M. Cynewski / GCSS	ISO Compliance	M. Gauthier / GCSS	August 27, 2009