



PS Series Storage Arrays

Choosing a Member RAID Policy

Abstract

This Technical Report introduces you to the member RAID policies for PS Series groups. 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
V3.0

Copyright © 2004 – 2006 EqualLogic, Inc.

October 2006

EqualLogic is a registered trademark of EqualLogic, 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.

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

PS Series Firmware Version 3.0

Table of Contents

Introduction to Member RAID Policies	1
PS Series Array Disk Layout	3
Member RAID Level Performance Characteristics	4
RAID Policy Disk Failure Protection	5
Changing the Member RAID Policy	9
Summary	10
Documentation and Customer Support	11

Introduction to Member RAID Policies

PS Series storage arrays include redundant, hot-swappable components—disks, 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 disks. 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, or RAID 5)
- Whether or not spare disks 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 disks are automatically configured according to the designated RAID level, with the appropriate number of spare disks.

Note: You can convert a member from one RAID policy to another only in certain circumstances, so be careful when choosing a RAID policy. For example, you can convert a member only to a RAID policy that requires *less* disk space than the current policy. For more information, see *Changing the Member RAID Policy*.

Although all RAID levels provide good performance, 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 group.

The following RAID policies can be set using the Group Manager GUI:

- **RAID-10.** Striping on top of multiple RAID 1 (mirrored) sets, with 1 or 2 spare disks (depending on the total number of disks in the array). RAID-10 provides good performance for random writes, in addition to the highest availability. However, since the disks are mirrored, RAID 10 provides the least capacity.
- **RAID-50.** Striping on top of two RAID 5 (distributed-parity) sets, with 1 or 2 spare disks (depending on the total number of disks in the array). RAID-50 provides a balance of performance, availability, and capacity.
- **RAID-5.** One RAID 5 set, with one spare disk. RAID-5 is similar to RAID-50, with more capacity (two additional disks) but lower availability and random write performance when compared with the other RAID types.

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

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

A no-spare RAID policy leaves the member vulnerable to data loss after a single disk failure. If you select a no-spare RAID policy, it is strongly recommended that the installation have on-site spare disks and 7x24 personnel for monitoring the member and replacing failed disks.

The following examples show using the Group Manager GUI to set the RAID policy for a member. See the *PS Series CLI Reference* manual for information about using the CLI.

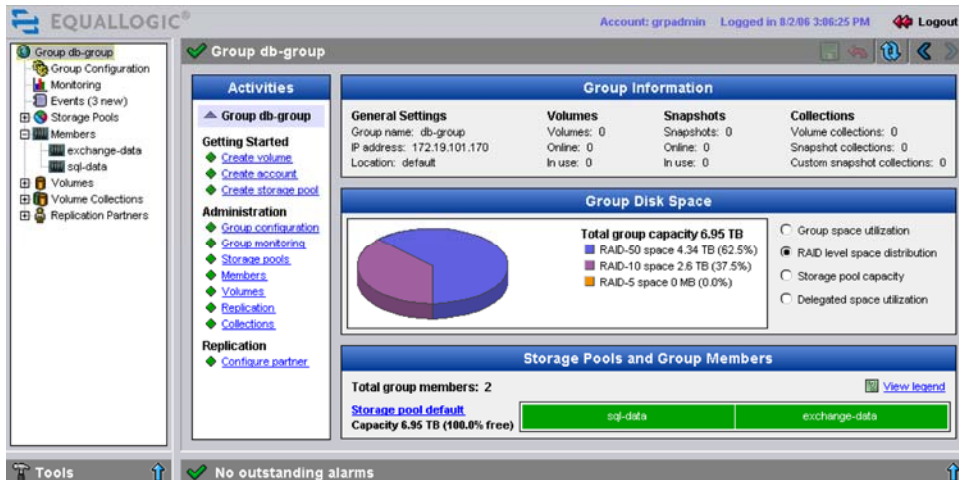
Once you select a RAID policy for a member you can perform only the following online RAID policy conversions:

- RAID-10 to RAID-50
- RAID-10 to RAID-5
- RAID-50 to RAID-5
- RAID-10, no-spare to RAID-50, no-spare (must use the CLI to convert)

If a member has a RAID policy that cannot be converted, you can change the policy only by removing the member from the group, re-adding it to the group, and then setting the desired RAID policy.

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



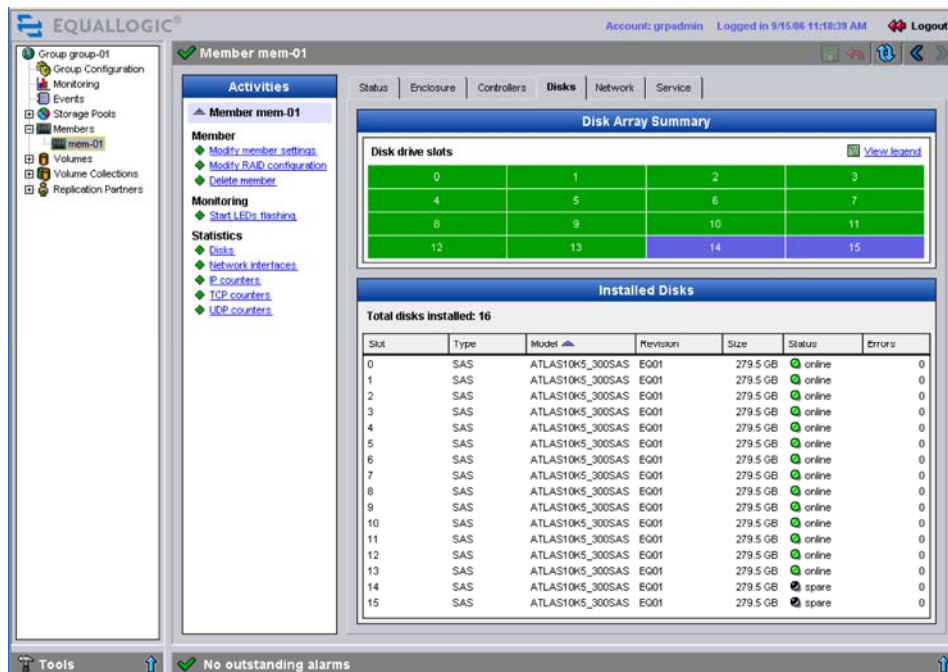
PS Series Array Disk Layout

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

Members → *member_name* → Disks tab

The Member Disks window appears, displaying 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 PS3000 array with 16 disks (including 2 spares).

Member Disks Window



If you choose RAID-10 as the RAID policy, disks are automatically organized into *multiple* RAID 1 (mirrored) sets, and data is striped across the sets. One or two disks are reserved as spares, depending on the total number of disks in an array. For example, in a PS3000 array:

- An 8-drive array configures three RAID 1 sets, with two disks in each set, and two spare disks.
- A 16-drive array configures seven RAID 1 sets, with two disks in each set, and two spare disks.

If you choose RAID 50 as the RAID policy, disks are automatically organized into *two* RAID 5 (distributed-parity) sets, and data is striped across the sets. One or two disks are reserved as spares, depending on the total number of disks in an array. For example, in a PS400E Series array:

- A 7-drive array configures two RAID 5 sets, with three disks in each set, and one spare disk.
- A 14-drive array configures two RAID 5 sets, with six disks in each set, and two spare disks.

If you choose RAID 5 as the RAID policy, disks are automatically organized into one RAID 5 (distributed-parity) set. One disk is reserved as a spare. For example, in a PS400E array:

- A 7-drive array configures one RAID 5 set with six disks in the set and one spare disk.
- A 14-drive array configures one RAID 5 set, with 13 disks in the set, and one spare disk.

Member RAID Level Performance Characteristics

A PS Series group is designed to provide storage to a broad range of environments and can deliver 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 workload (heavy or light) and how the application performs I/O (small random transfers or large sequential transfers). The read/write ratio is also a factor. 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.

A PS Series storage array operates under normal, non-failure conditions most of the time. However, when comparing performance RAID level characteristics, you must understand the array performance impact when an individual RAID 1 or RAID 5 set is:

- **Reconstructing.** When a set disk 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 1, 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/RAID 50, data reconstruction involves multiple disks, so the overall performance impact may be significant. For example, if you have a six-disk set, reconstruction involves six disks (five active and one spare). However, under a normal workload, reconstruction may not impair overall performance.
- **Degraded.** When a set disk fails and there is no spare disk to replace it, the set is still functional, but degraded. Another disk failure could result in a complete loss of data in the set.
 - With RAID 1, read performance decreases 50% *only* in the set that has experienced the failure. There is negligible impact on the overall read performance.
 - With RAID 5/RAID50, read performance decreases by 30% and random write performance decreases by 50% when a set is degraded.

RAID Policy Disk Failure Protection

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

Depending on the RAID policy and the total number of disks in each PS Series storage array, one or two spare disks are automatically configured and used in the event of a disk failure. If a spare disk is *not* available, each set can survive only one disk failure. In this case, data continues to be available, but the set is degraded. A disk failure in a degraded set results in loss of data.

When a disk in a RAID 1 or RAID 5 set fails, an array behaves as follows:

- **Spare disk is available.** Data from the failed disk is reconstructed on the spare and continues to be available. During reconstruction, the set that contains the failed disk is temporarily degraded. After reconstruction, there is no impact on performance.
- **Spare disk is not available; set becomes degraded.** Data continues to be available, but the set is degraded.
- **Failure in a degraded set.** When a disk fails in a degraded set, the member could be set offline, and any volumes and snapshots that have data located on the member are set offline.

When you replace a failed disk:

- **Spare disk was used.** Data has already been reconstructed on the spare disk, so the new disk becomes a spare.
- **Set was degraded.** Data is reconstructed on the new disk and after reconstruction, performance returns to normal.

Notes: The following figures show how a 14-disk, PS400E Series array responds to a sequence of 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 set.

The following figure, Surviving Disk Failures With RAID 10 as the Member RAID Policy, shows how using RAID 10 can increase (but cannot guarantee) the chance that an array can survive multiple disk failures in a 14-disk, PS400E array. This RAID level is appropriate for mission-critical data that requires the highest level of protection and rapid accessibility.

Surviving Disk 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 each of two individual sets

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



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

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




The following figure shows how RAID 50 provides more capacity than RAID 10 while continuing to provide protection from disk failure and good performance and availability for a 14-disk, PS400E array.

Surviving Disk 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 set

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



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

After reconstruction, the RAID 5 set is degraded, but the array is still functional.



The following figure shows how RAID 5 provides capacity and protection from disk failure but less availability for a 14-disk, PS400E 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.

Surviving Disk Failures with RAID 5 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 5 set

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



RAID protects data if a disk fails in the RAID 5 set (and a spare is not available)

After reconstruction, the RAID 5 set is degraded, but the array is still functional.



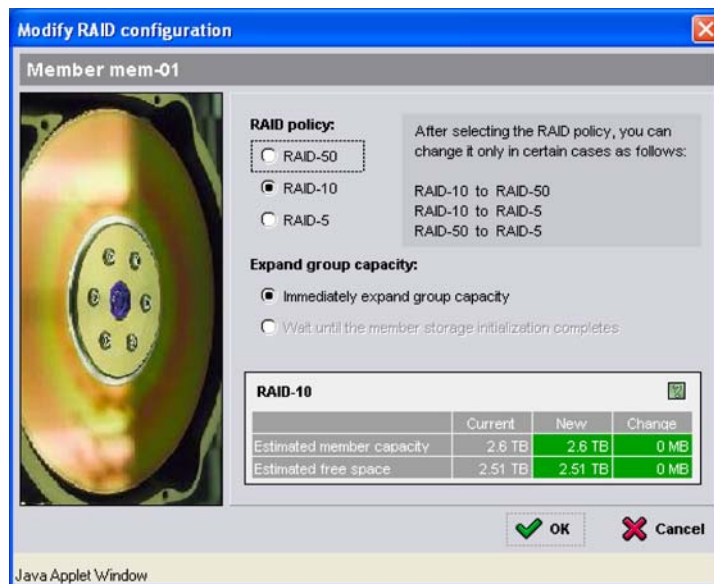
Changing the Member RAID Policy

To change 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.

Modify RAID Configuration Wizard

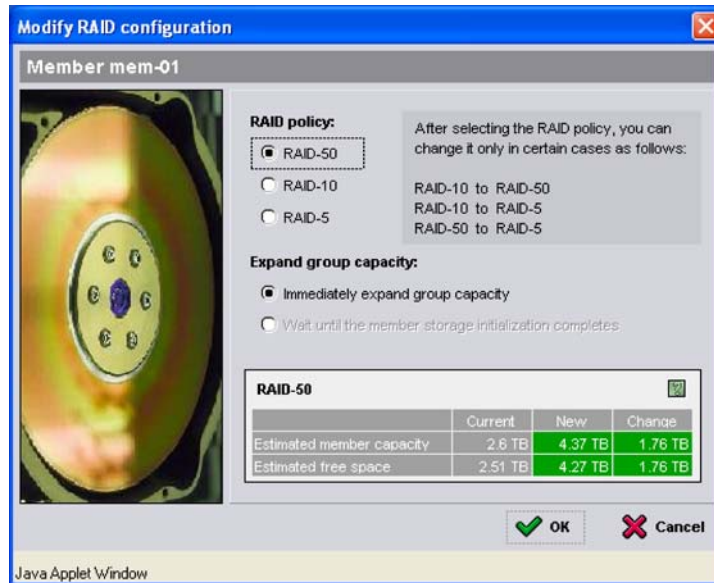


Notes: If you are not allowed a conversion, that option will appear grayed out.

You must use the CLI to change RAID-10, no-spares to RAID-50, no-spares.

Select the desired conversion. Note that the title of the table and the capacity listed in the table change to reflect the new RAID policy as shown in the following figure.

RAID Configuration Wizard Reflecting Changes



When you are satisfied, click OK.

Conversions can take a long time, so you may want to schedule it when the workload is light.

Summary

There are many factors to consider when choosing a RAID level for a PS Series member. Both RAID 10 and RAID 50 provide excellent reliability for your data, in addition to overall high performance. RAID 50 provides the most storage capacity, while RAID 10 provides the best performance for workloads that are mostly small random writes. RAID 5 provides the greatest capacity at the expense of performance. Most administrators choose RAID 50 because of its superior storage capacity.

When choosing a RAID level for a member, consider the factors summarized in the following table:

Workload Requirement	RAID-10	RAID-50	RAID-5
Capacity	Good	Better	Best
Sequential writes	Best	Best	Best
Random writes	Best	Better	Good
Availability	Best	Better	Good
Performance impact of drive failure/RAID reconstruction	Least	Moderate	Moderate – Heavy

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.
- *QuickStart*. Describes how to set up the storage array hardware and create a PS Series group. Be sure to 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. Be sure to use the manual for your array model.

Technical support on EqualLogic products is available for customers with arrays under warranty and customers with a valid support contract. You can obtain technical support in the following ways:

- Visit the EqualLogic Customer Support website to download the latest documentation and firmware. Go to the EqualLogic corporate website at www.equallogic.com. If you have a customer support account, click `Login` on the top right side of the page. In the login dialog, enter your user name and password.

If you do not have an account, click `Request a Customer Account` to register for an account.

- From the EqualLogic Customer Support website, you can submit a service request. Go to the EqualLogic corporate website at www.equallogic.com. If you have a customer support account, click `Login` at the top of the page. In the login dialog, enter your user name and password. Click `Customer Service Home` on the top left side of the page. Click `Log a Case` and follow the instructions.
- In the United States, call 877-887-7337. Outside the United States, call +00 1 919-767-5729. If the issue is urgent, ask to speak with a member of the EqualLogic Customer Support team.
- Send e-mail to supportnp@equallogic.com and clearly describe the issue or problem.