

# PS SERIES BEST PRACTICES DEPLOYING POOLS AND TIERED STORAGE IN A PS SERIES SAN

## ABSTRACT

This Technical Report describes the concept of tiered storage, some of the potential benefits, and helps you take advantage of storage pools and other features of Version 3 PS Series Firmware to implement tiered storage in a PS Series SAN.

Copyright © 2006 EqualLogic, Inc.

November 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.

PS Series Firmware Version 3.0 or later.

EqualLogic, Inc.  
9 Townsend West  
Nashua, NH 03063  
Tel: 603.579.9762  
Fax: 603.579.6910

## TABLE OF CONTENTS

Technical Report and Software Revision Information.....	iv
Introduction.....	1
What is tiered storage? .....	1
Why implement tiered storage?.....	1
How does a PS Series SAN support tiered storage?.....	1
Planning and Design Considerations .....	5
What business problem am I solving?.....	5
Are there any special organizational requirements?.....	5
What service agreements must be met?.....	5
Do I need a tiered solution?.....	5
What are the storage requirements of each application? .....	5
What are the characteristics of the available storage?.....	6
Which RAID levels are appropriate for each pool? .....	6
How will you migrate existing data into the new environment?.....	6
Some Examples of Tiered Storage Designs .....	7
Preparing for Deployment.....	9
Assign Applications to Tiers .....	9
Determine Resource Requirements for Each Tier.....	9
Plan Storage Pools.....	9
Plan Volumes and Assign them to Pools.....	9
Deploying.....	10
Add Members to, or Set Up a PS Series Group .....	10
Create Pools within the Group .....	10
Assign Members to Pools.....	12
Create or Modify Volumes and Assign Them to Pools.....	12
Summary .....	15
EqualLogic Documentation and Technical Support .....	16

## TECHNICAL REPORT AND SOFTWARE REVISION INFORMATION

The following table describes the release history of this Technical Report.

Revision	Date	Change
1.0	November 2006	Initial release.

The following table shows the software versions used for the preparation of this Technical Report.

Vendor	Model	Software Revision
EqualLogic, Inc.	PS Series Firmware	3.0 or later

The following table lists the documents referred to in this Technical Report.

Vendor	Document Title
EqualLogic	Choosing a Member RAID Policy Technical Report
EqualLogic	Network Connection and Performance Guidelines Technical Report
EqualLogic	Group Administration Manual
EqualLogic	Quick Start Guide

## INTRODUCTION

EqualLogic PS Series products support deployment in a tiered SAN architecture. In this technical report, you will learn the basics of tiered storage and how a SAN built from multiple PS Series arrays can be used to implement a tiered storage architecture. Subsequent sections of the report address planning and design considerations, preparing for deployment, and typical steps involved in an actual deployment.

By the end of this report, you should have a firm idea about how you can choose to use PS Series products to implement a tiered storage architecture that addresses your organization's unique needs.

This report assumes that you have an existing PS Series group that you are seeking to expand. If this is not the case, you will still find that most of the concepts and considerations in this report apply to the initial deployment of a SAN or to working with an existing multi-member SAN.

### What is tiered storage?

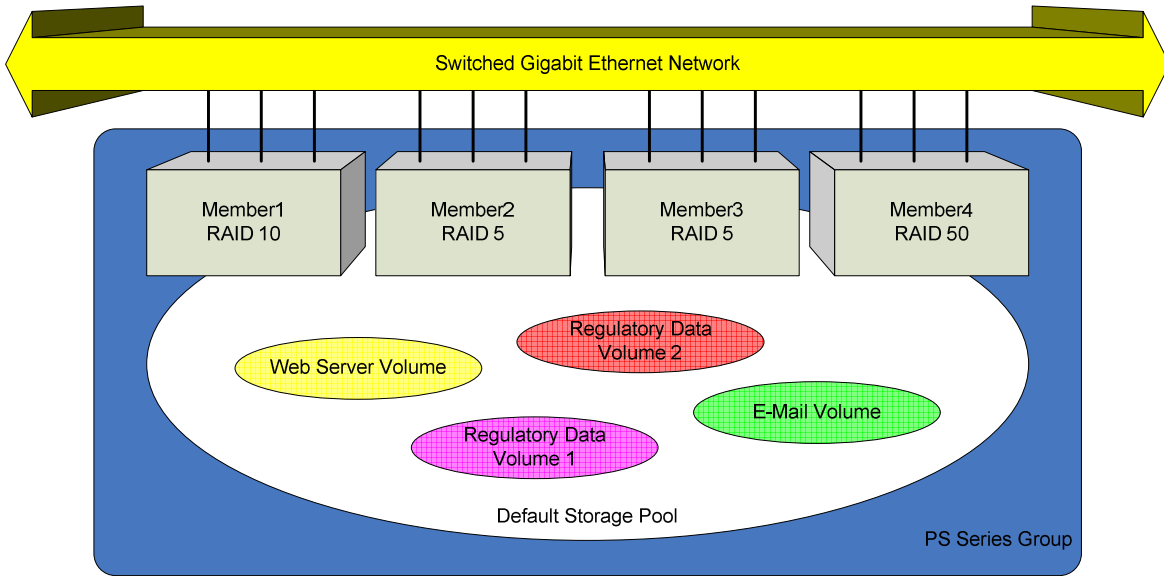
Tiered storage is storage that is differentiated and optimized for the type of data that is stored on it. Different types of data have different values to the company and different associated service level agreements. A tiered storage system provides different types of storage that are suited to the needs of the different types of data that are stored in the system. Instead of optimizing for maximum storage utilization, a tiered storage system allows you to optimize for maximum application performance and lowest cost.

### Why implement tiered storage?

Tiered storage is important because your IT resources are finite. In the past all data was usually stored on high-speed, high-cost disk. That is the simplest approach, but may not be the best. Some of your data supports mission-critical real-time applications. Other data needs to be retained (although seldom accessed) for years to meet regulatory requirements. Some of your data has requirements between these two extremes. Some of your data is covered by Service Level Agreements (SLAs) that impose other restrictions or requirements that are difficult to address in a one-size-fits-all storage environment. Implementing storage tiers that provide the appropriate levels of performance, availability, redundancy, fault-tolerance and SLA compliance for each type of data allows you to meet those requirements at the lowest cost.

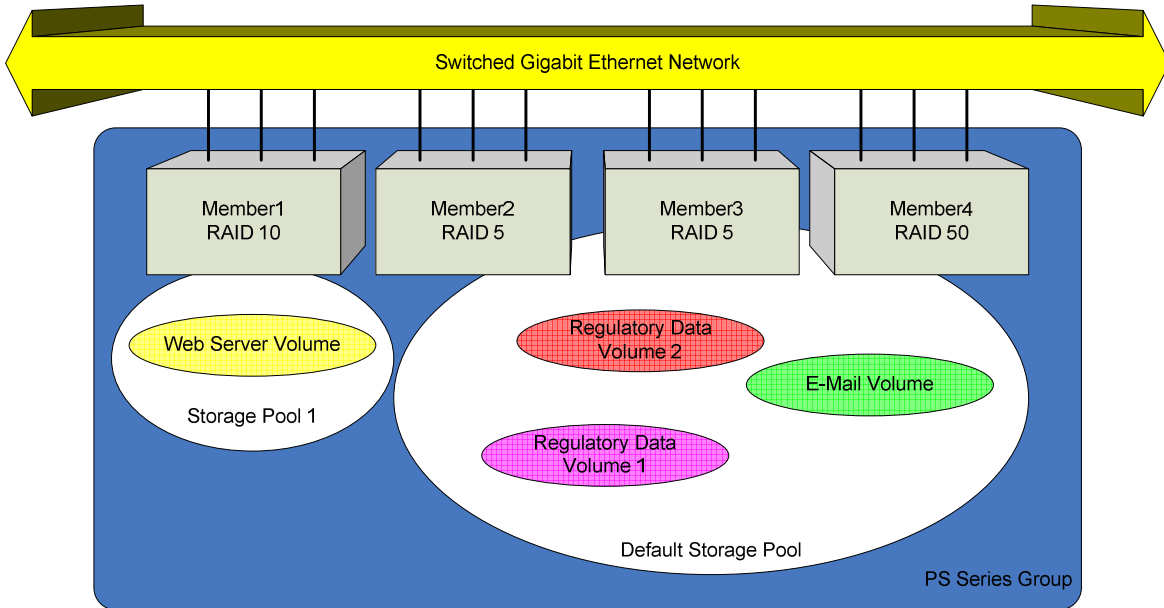
### How does a PS Series SAN support tiered storage?

A *PS Series SAN* (also known as a *PS Series group*) is a Storage Area Network (SAN) comprised of one or more PS Series Arrays connected to an IP network. Individual PS Series Arrays that are part of a group are called *members*. Each member is assigned to a storage pool (more about this important new concept later). When a group contains more than one member, those members can provide different capacity and cost levels, and be configured to use different RAID levels. These characteristics, along with the ability to move volumes into specific pools without any interruption in service, can be used to implement tiered storage without disrupting your current operations. Figure 1 shows a PS Series group with multiple members in a single pool.



**Figure 1: Multiple members in a single pool.**

Beyond this, when a PS Series group contains multiple members, the overall storage space can be divided into separate *pools*, providing an even greater level of control over storage resources in the SAN. Given enough member arrays (a pool must contain at least one member) a PS Series group can be divided into up to four pools. A pool acts like a SAN within a SAN, creating an isolated storage environment within the overall SAN. Figure 2 shows the same PS Series group divided into two pools.



**Figure 2: Two tiers of data storage implemented as two pools.**

Volumes are assigned to a pool and remain within that pool unless moved by the administrator. Load balancing is done automatically within the pool, using storage resources contained within the pool's members.

Whether or not you choose to deploy multiple pools, you still manage the entire SAN (including all pools) through the single administrative interface shown in Figure 3. This simplifies administration of the group as a whole. This also makes it possible to move volumes and even members between pools as your requirements change. Some examples will make the benefits of implementing tiered storage in a PS Series group clearer.

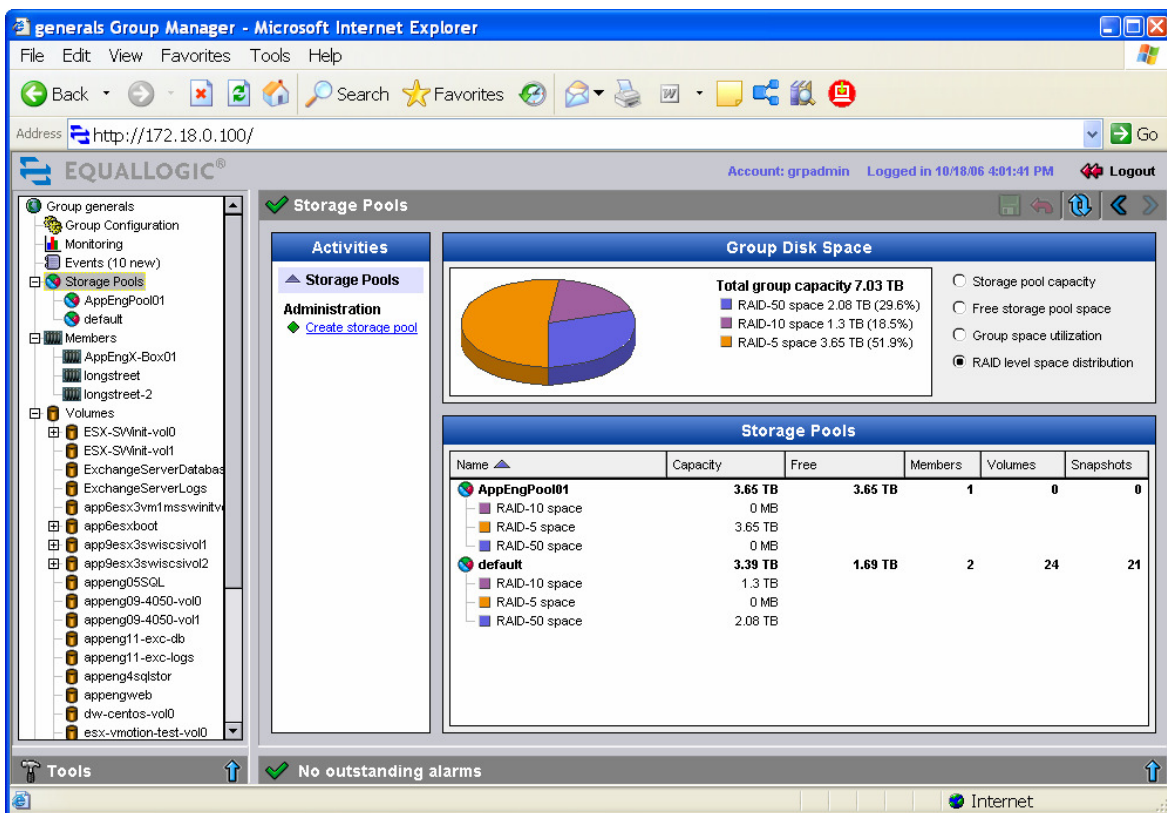


Figure 3: Manage all your pools from a single interface.

In the scenario described earlier (mission-critical, real-time data and bulk-stored data retained for years), we can now discuss how different PS Series arrays could be matched to these needs.

Mission-critical applications (particularly those that involve lots of small, random writes) might be best stored on PS Series arrays with 15,000 RPM drives configured as RAID 10. This results in storage optimized for high-speed random-access.

For many applications, this performance is not required for volumes that are providing bulk storage of less-frequently-needed data. These volumes might best be stored on PS Series arrays with 7,200 RPM drives configured as RAID 50. This results in storage optimized for greater storage capacity.

**Note:** For a more complete discussion of creating a SAN that meets your organization’s requirements; please see *Planning and Design Considerations* later in this report.

### Using Pools to Optimize Price/Performance

If you have mission-critical data to protect, as well as a need for bulk storage, you could choose to create two pools to meet these two very different needs. The pool for the mission-critical data would contain your highest performance arrays, configured with the RAID level that best supports the application’s read and write pattern. Only your mission-critical applications would have access to this pool.

The other pool could contain high-capacity arrays configured to maximize storage capacity and minimize cost for long-term storage.

This model works well when you want to add additional arrays to the existing group. You can add the new arrays to whichever pool makes the most sense, as well as migrating existing arrays from pool to pool, all without any downtime. For example, you might add a PS3800XV to the pool for mission-critical data, then reconfigure a PS300E array from the mission-critical pool with RAID 50 and move it to the long-term storage pool to increase that pool’s capacity. Again, all this happens online, without any downtime for the SAN.

## **Using Pools to Segment Data by Organization**

Perhaps each business unit is responsible for their own storage costs, and they understandably want to guarantee exclusive access to the storage they pay for.

You could meet this requirement by creating separate storage pools for each business unit. By placing the each unit's data in its own pool, and setting the access controls to ensure that only that unit has access to that pool, you can ensure that business units have full use of their own storage. Within each pool you still get the benefits of PS Series automatic load-balancing, and you retain the ability to manage all the storage in the group from a single interface.

## **Using Pools to Segment Data by Application**

Pools also help you to isolate crucial applications, while retaining the management benefits of a single SAN. By creating a separate pool for the application data, and setting the access controls so that only the servers associated with that application have access to that pool, you can physically separate the data while still keeping a single PS series group and managing all the storage through a single management interface. Within each pool you still get the benefits of PS Series automatic load-balancing and management.

## **Tiered Storage within a Single Pool**

It is possible to implement basic tiered storage within a single pool. By creating a heterogeneous pool (one containing multiple members with different RAID levels or performance characteristics), and letting the SAN automatically migrate volumes to the most appropriate members, you can get many of the benefits of multiple pools with minimal effort.

## **PLANNING AND DESIGN CONSIDERATIONS**

When thinking about a tiered storage solution, there are a number of factors to consider. These include, but are not limited to, the following:

- What business problem am I solving?
- Are there any special organizational requirements?
- What service agreements must be met?
- Do I need a tiered solution?
- What are the storage requirements of each application?
- What are the characteristics of the available storage?
- Which RAID levels are appropriate for each pool?
- How will you migrate existing data (if any) into the new environment?

The following sections provide more in-depth information about each of these considerations. Be sure to have a well-considered plan in place before beginning to implement tiered storage.

### **What business problem am I solving?**

What business problem are you trying to solve with this SAN deployment or upgrade? Is there a need to store data from additional applications in the SAN? Do business units require exclusive use of their own storage resources? Do you wish to take advantage of the new capabilities of PS Series SAN?

Whatever the specific details, the requirements of the business problem should drive your planning and design work.

### **Are there any special organizational requirements?**

Organizational requirements, such as a need to guarantee a business unit exclusive access to storage, can complicate the effort to solve the primary business problem. Be sure to factor in any special organizational requirements during this planning and design phase.

### **What service agreements must be met?**

In addition to any special organizational requirements, you need to consider any service agreements that apply to your solution. Formal commitments for levels of performance, availability, redundancy, and fault tolerance, for example, will affect many aspects of your design. Addressing these issues now ensures that your design will meet your service commitments in addition to solving the primary business problem.

### **Do I need a tiered solution?**

The high-level considerations discussed previously will help you determine whether you should implement a single-pool storage design, or create multiple tiers of storage by setting up separate pools. While you still manage multiple tiers of data from a single interface regardless of the number of pools in the group, a single-pool group is easier to implement and maintain than a multi-pool group. EqualLogic best practices suggest always using as few pools as possible and allowing the PS Series firmware to handle load balancing and assignment of volumes to the most appropriate storage.

### **What are the storage requirements of each application?**

Since pools utilize only their own resources, each pool must have the storage capacity and throughput to serve the applications that are assigned to it, regardless of the overall capacity of the group. Be sure to design

each pool to have enough capacity for all the data that will be stored on it, plus enough capacity for any snapshots or data replication that will be required.

### **What are the characteristics of the available storage?**

The specific types of arrays you have can affect how you construct pools. If you have arrays with varying drive speeds and storage capacity, you can optimize the performance of your overall SAN through careful consideration of these factors. Consider using arrays with the fastest drives for real-time or transactional data, while using those with the greatest storage capacity for bulk storage of less time-critical data.

### **Which RAID levels are appropriate for each pool?**

The ability to assign different RAID levels (5, 10, or 50) to each member of a pool increases your operational flexibility. The only constraint on RAID level assignments is that, to guarantee best performance, you should not assign the same RAID level to members of a pool that have different drive speeds.

For more information on RAID level options, see the EqualLogic Technical Report, *PS Series Storage Arrays Choosing a Member RAID Policy*, available at the EqualLogic website.

### **How will you migrate existing data into the new environment?**

If you are moving storage for an application onto the SAN, you will need to migrate that application's data. EqualLogic best practices call for migrating application data onto the SAN using a file-level program such as Robocopy for Windows systems or dd for UNIX, rather than using disk imaging tools.

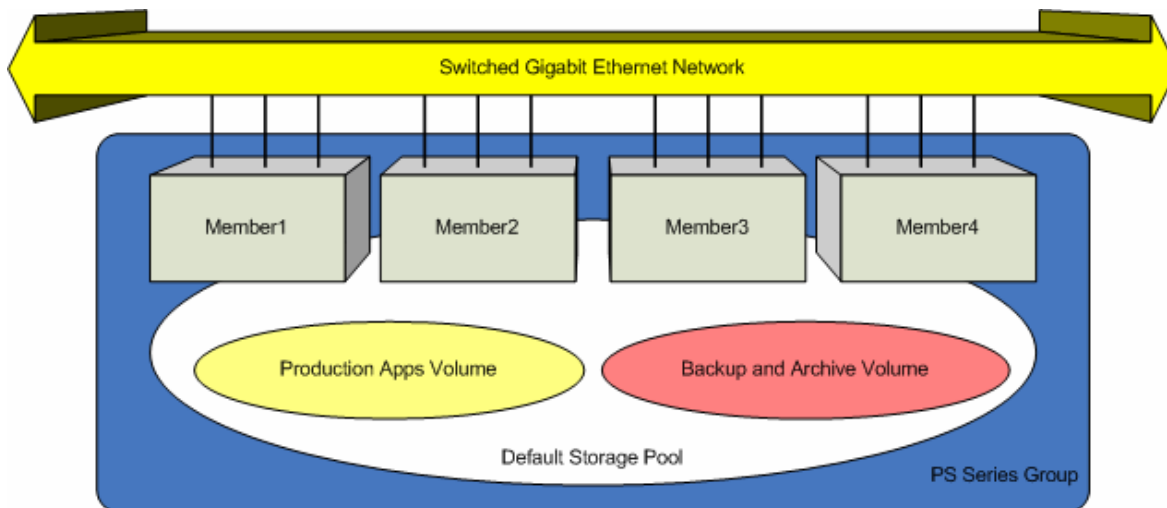
## Some Examples of Tiered Storage Designs

To aid you in your own tiered storage planning and design work, we have included these three examples of common tiered storage designs. When viewing these examples, keep in mind that (except when being migrated from one member to another) volumes are stored on only one member array, each member array must belong to one and only one pool, and all members belong to one pool.

**Note:** The number of members and volumes used in the following examples were chosen to simplify the explanation of each example, and are not meant to constrain your own design in any way.

### A Single Pool

In many cases, a single storage pool will meet your needs. Figure 4 shows a single-pool SAN with four member arrays and two volumes, one for production application data (a database, email stores, and file & print data) and one for backup and archive data.



**Figure 4: All data in a single pool.**

A design like this offers you a lot of flexibility without the use of multiple pools. In the most basic design, all members may be the same model and configured with the same RAID level. The PS Series firmware automatically monitors data usage patterns and over time would optimize performance by the way it allocates data across the members. This homogeneous pool design could be quite sufficient for your needs.

Thanks to the flexibility of the PS Series firmware, you can also implement a heterogeneous pool design, where the members are configured with different RAID levels and/or consist of different models of PS Series array. Each RAID level offers its own combination of performance and availability. The PS Series firmware can take advantage of this by automatically moving data within volumes to the member with the RAID level most appropriate for that data's usage pattern.

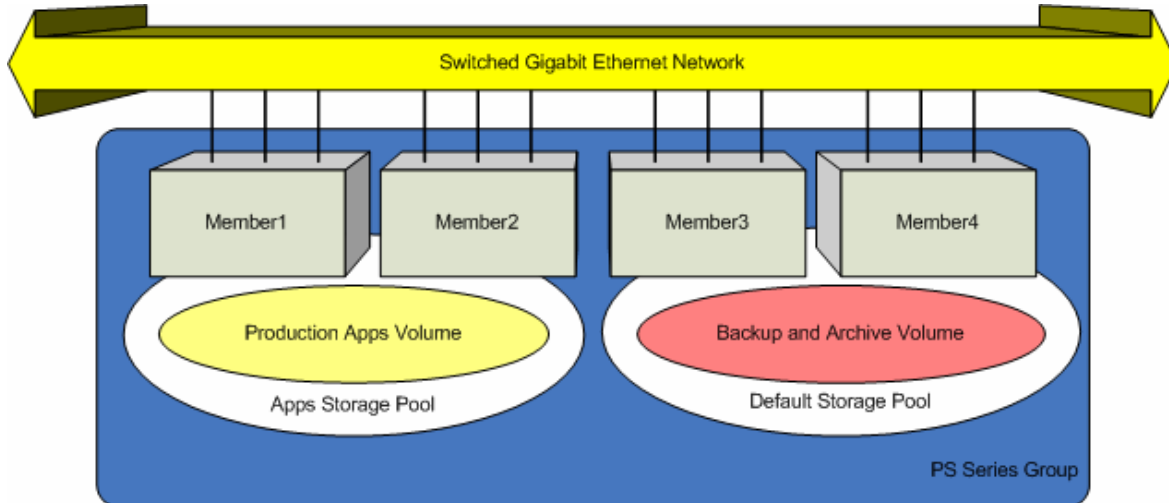
Similarly, if the array consists of PS Series arrays with differing RAID levels, the PS Series firmware will take that into account and move data to the RAID level most appropriate for that data's usage pattern.

PS Series firmware uses load balancing metrics to assign volumes to a RAID level appropriate for current usage patterns. By default, automatic load balancing handles the assignment of volumes to RAID levels. However, if you wish to specify a preferred RAID level for a volume, you can do so. Specifying a preference overrides automatic load balancing for that volume. If a member is available with the specified RAID level, the PS Series firmware places the volume in that RAID level.

You use the Advanced tab of the Modify volume settings dialog box to set a volume storage preference.

## Two Pools

Suppose you want more separation of application data from backup and archive data, or you want more control over which volumes run on which hardware. You might implement two pools, as shown in Figure 5. Again, we have four member arrays and two volumes, one for production application data and one for backup and archive data. The difference here is that each volume is in a different storage pool.



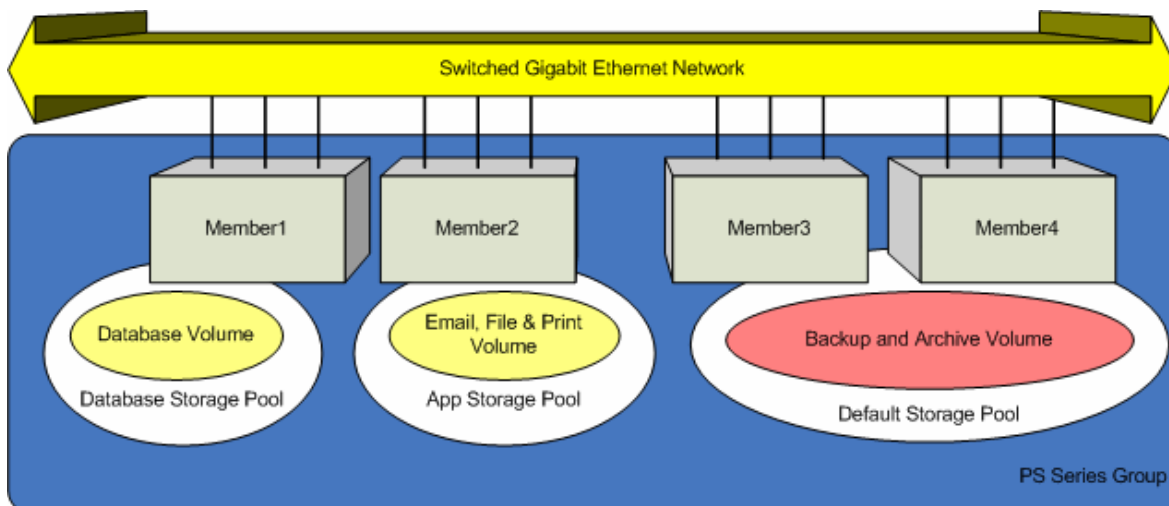
**Figure 5: Separating Application data from Backup & Archive data with two pools.**

Since aside from management, each pool is independent of the other, each pool gives you all the options that you had for a single pool. That is, you can design each pool to be homogenous (all members of the pool running PS300E arrays configured as RAID 10, for example) or heterogeneous (with differing RAID levels and/or array models within the same pool).

One additional constraint you must take into account with this configuration is the amount of storage available within each pool. Not only must the storage capacity of the PS Series group as a whole be large enough to contain both volumes, but the capacity of each pool must be large enough to contain the volume that is stored within the pool.

## Three Pools

Now suppose that you want to optimize the performance of the database application. One approach would be to implement three pools: one for the database, one for the other applications, and one for backup and archive data. It might look like Figure 6.



**Figure 6: A third pool allows for further Application Data tiering.**

In this design, you would likely put the member (or members) with the best random access performance (Member 1 in Figure 6) in the Database Storage Pool to optimize the performance of the database. As noted previously, you can make each pool homogenous and heterogeneous to best optimize the performance of the overall environment.

As with the two-pool design, you must be concerned with the storage capacity of each pool, in addition to that of the group as a whole.

## **PREPARING FOR DEPLOYMENT**

Performing these activities prior to beginning the actual deployment will allow you to catch any problems with the plan on paper. Depending on the size and complexity of your storage plan, it may take multiple iterations to come up with optimal tier, pool, and volume assignments. Deployment preparation tasks include:

- Assigning applications to tiers
- Determining resource requirements for each tier
- Planning storage pools
- Planning volumes and assigning them to pools

### **Assign Applications to Tiers**

After taking into account the design considerations discussed above, organize applications into between one and four tiers.

### **Determine Resource Requirements for Each Tier**

Given the storage and throughput requirements of the applications in each tier, determine the total resource requirements of each tier. Remember to plan adequate space for snapshots, replication, and expected near-term growth.

### **Plan Storage Pools**

Determine if your available stock of PS Series arrays can be divided into pools in such a way as to meet the resource requirements for each of the planned tiers. Be sure to consider the impact of RAID-level selection within each pool.

Remember that resources cannot be shared across pools. You must allocate resources to pools by allocating entire PS Series arrays. Consider acquiring additional arrays to facilitate the creation of pools that meet your requirements. The PS Series product line includes arrays with a variety of speed, capacity, and price points, enabling you to more easily meet your corporate storage needs.

### **Plan Volumes and Assign them to Pools**

Determine what volumes are required. Determine the appropriate size (including reserve space for snapshots) of each volume, then assign them to the appropriate storage pools.

Once you complete this process for all the volumes that will be included in a given pool, confirm that the members assigned to that pool are capable of providing the necessary space.

## DEPLOYING

This section describes a sequence of steps that you will typically complete when you are deploying a tiered storage environment. These steps are:

- Add members to, or set up a PS Series group
- Create pools within the group
- Create volumes and assign them to pools

### Add Members to, or Set Up a PS Series Group

If you have an existing PS Series group and are adding members to it, follow the instructions under, “Adding Members to an Existing Group” in the PS Series *Group Administration* manual.

If you are building out a SAN from scratch, you will need to install the hardware and connect the serial cable (only needed for the initial setup) and network cables on the array that will be the first group member. See the PS Series *QuickStart* for detailed installation and setup information.

If you want multipath access between servers and the group, connect and configure multiple network interfaces on each array that is a group member. See the Technical Report *Network Connection and Performance Guidelines* on the EqualLogic Customer Service website for more information.

See “Customizing the Group” in the PS Series *Group Administration* manual for instructions on applying CHAP and other Access Controls to the members of the group.

After you create the group, start the Group Manager in your browser by connecting to the group IP address and entering the group administration name and password when requested.

### Create Pools within the Group

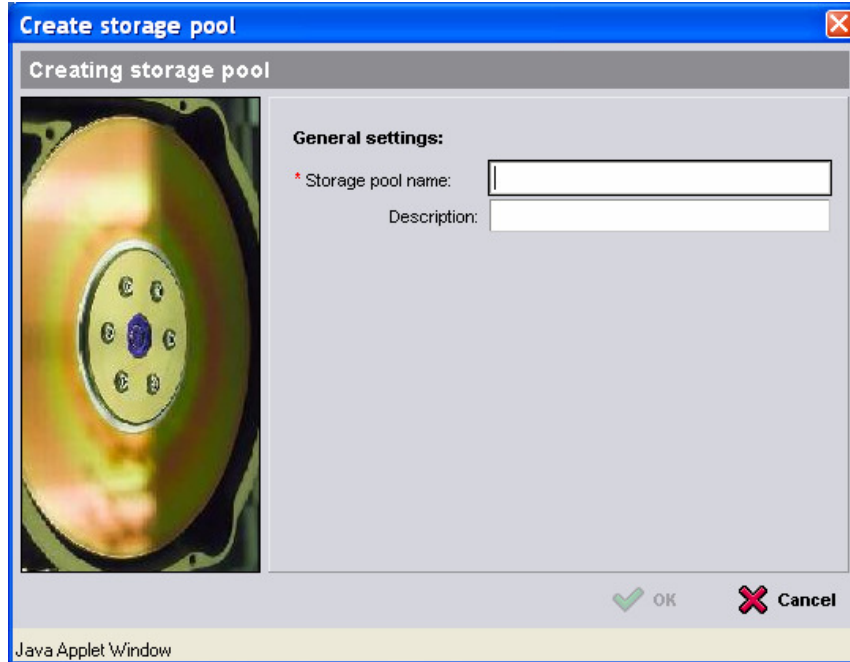
Once you have added members to the group, you are ready to create the storage pools that will contain the members. You should have already determined the number of pools you will create and their names. The following steps walk you through the process of creating one pool within the group.

**Note:** Additional information about creating and working with pools is available in Chapter 4 of the *Group Administration* guide.

## Creating a Pool


Follow these steps to create a pool within the group:

1. Open the GUI to the Group Summary window. In the Activities panel, click Create storage pool. This opens the Create Storage Pool dialog box.



**Figure 7: Create storage pools here.**

2. Enter a Storage pool name and Description, then click OK to create the pool. The new empty storage pool appears in the Storage Pools and Group Members panel of the Group Summary window.



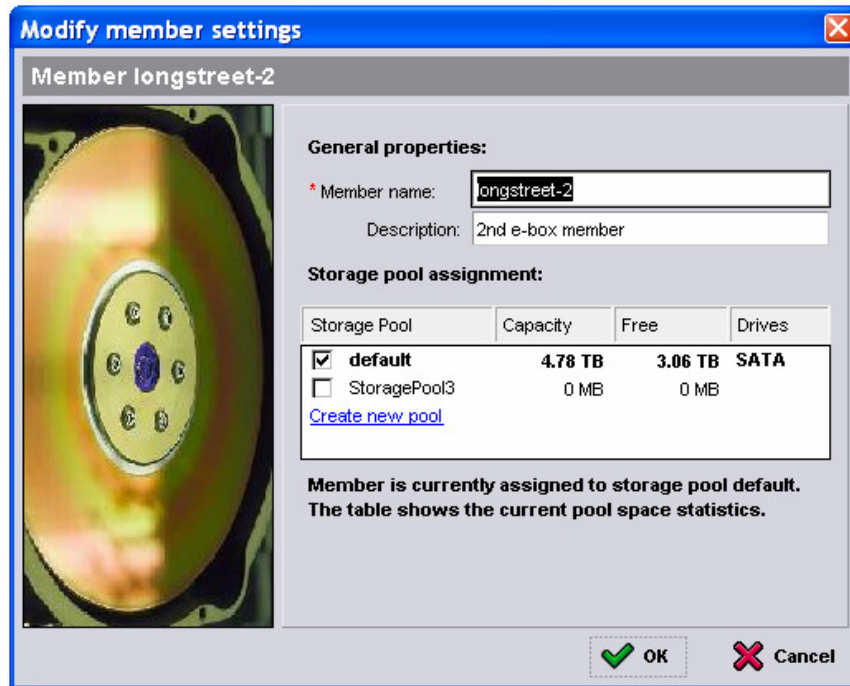
**Figure 8: The new empty pool appears in this pane.**

3. Repeat Steps 1 and 2 for each additional pool you wish to create.

## Assign Members to Pools

Once all the pools have been created, you are ready to assign members to the pools. You should have already determined which members will be assigned to which pools. The following steps walk you through the process of assigning one member to a pool.

1. Open the GUI to the Members window. Right-click the name of a member. In the shortcut menu that appears, click Modify Member Settings. This opens the Modify Member Settings dialog box.



**Figure 9:** Move a member into a new pool using this dialog box.

2. Modify the Member name and description if necessary, then select the storage pool you want to assign this member to.
3. Click OK. The member becomes part of the assigned pool.
4. Repeat Steps 1-3 for each member you wish to modify or assign to a pool.

## Create or Modify Volumes and Assign Them to Pools

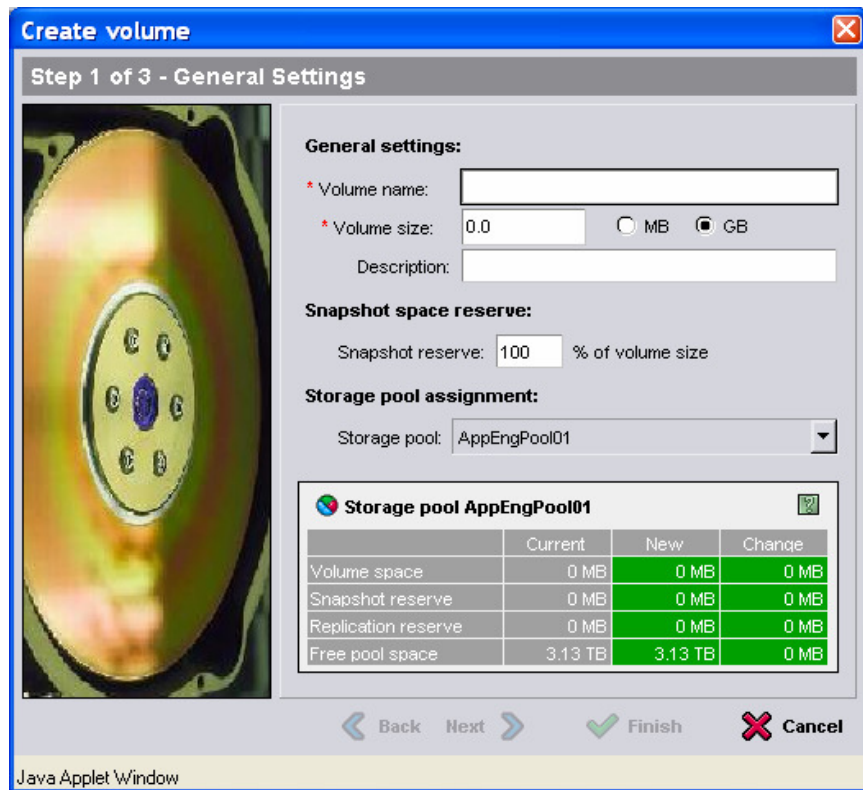
Now that the storage pools are created and have members assigned to them, you can create or modify the volumes you will need, assign them to the correct pools, and set the access control for each volume. The following steps walk you through the process for one volume.

**Note:** Additional information about creating and working with volumes is available in Chapter 5 of the *Group Administration* guide.

### Creating and Assigning a Volume

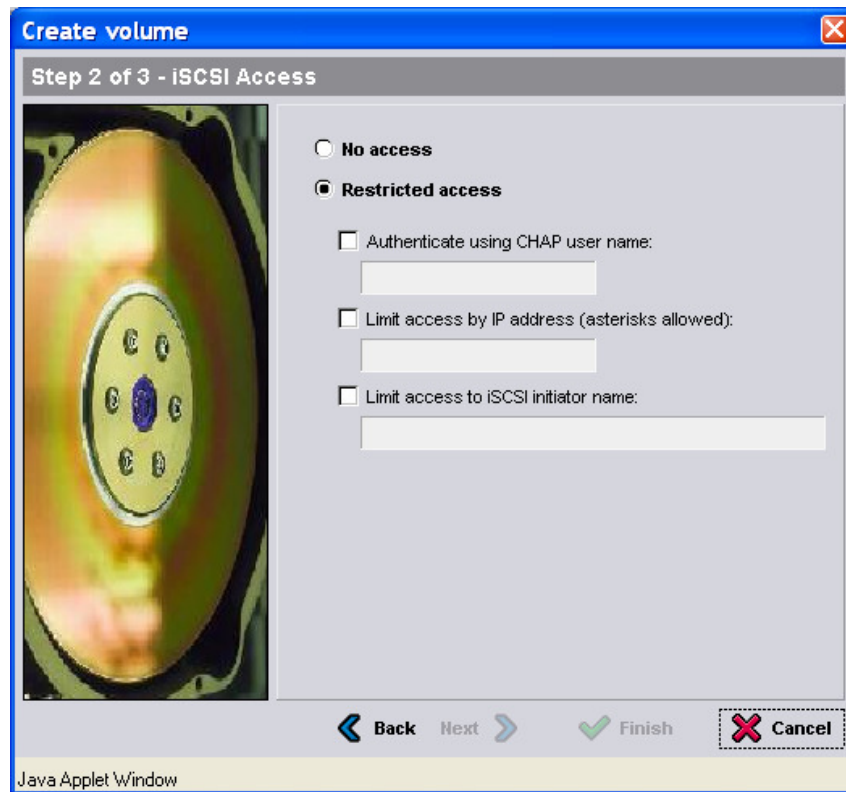
Follow these steps to create a volume and simultaneously assign it to a pool:

1. Open the GUI to the Group Summary dialog box. In the Activities panel, click Create volume. This opens the Create Volume dialog box to Step 1 of 3 – General Settings, as shown in Figure 10.



**Figure 10: Create a volume and assign it to a pool with this wizard.**

2. Enter the Volume name, Volume size, and optionally, a Description. Adjust the Snapshot space reserve percentage if desired.
3. In the Storage pool list, select the storage pool the volume will be assigned to. Click Next to open Step 2 of 3 – iSCSI Access in the Create Volume dialog box shown in Figure 11.



**Figure 11: Set Access Controls for the volume.**

4. Configure the Access Control method you will use for this volume. Click Next to open Step 3 of 3 – Summary in the Create Volume dialog box.
5. Click Finish to finish creating the volume.

## **SUMMARY**

A tiered storage solution provides different types of storage for each tier of applications and data, allowing you to optimize performance of the SAN while minimizing the overall cost of the implementation. An iSCSI SAN built around PS Series arrays allows you to implement a tiered storage system with multiple pools of storage, while still retaining all the benefits previously associated with the PS Series. These include automatic load balancing, single-screen management of the entire SAN, storage virtualization (within pools), and performance that increases as the SAN grows.

This report introduces tiered storage and the pooled storage capabilities of the PS Series arrays that implement storage tiers within the overall SAN. The report then discusses planning and design considerations that you should investigate before implementing your own tiered storage. After a discussion of the steps you need to take in preparation for deployment (such as assigning applications to tiers and volumes to pools), it provides a high-level guide to the deployment itself.

You should now have a good understanding of some of the uses of tiered storage, along with the ways in which EqualLogic's PS Series arrays can help you implement a tiered storage system that meets your organization's unique needs.

## EQUALLOGIC DOCUMENTATION AND TECHNICAL SUPPORT

Visit the EqualLogic Customer Service website, where you can download the latest documentation and firmware. You can also view FAQs, the Knowledge Base, and Technical Reports and submit a service request.

EqualLogic PS Series storage array documentation includes the following:

- *Release Notes*. Provides the latest information about PS Series storage arrays and groups.
- *QuickStart*. Describes how to set up the hardware and start using a PS Series storage array.
- *Group Administration*. Describes how to use the Group Manager 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 to manage a group and individual arrays.
- *Hardware Maintenance*. Provides information on maintaining the PS Series storage array hardware.

To access the Customer Service website, log in to the EqualLogic website ([www.equallogic.com](http://www.equallogic.com)), then click `Customer Service Home`.

To contact technical support, send e-mail to [supportnp@equallogic.com](mailto:supportnp@equallogic.com). If the issue is urgent, call 1-877-887-7337 to speak with a member of the support team.