



# Oracle9i™ Database Using a PS Series Group for NFS Server Storage

---

## **Abstract**

This Technical Report describes how to use a PS Series group to provide NFS server storage space to servers running Oracle9i Database. It includes group and server requirements, describes configuration choices, and provides recommendations for high availability

Copyright © 2005 EqualLogic, Inc.

February 2005

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.

EqualLogic has successfully completed the qualification with Oracle<sup>®</sup> using the test suite provided by Oracle through the Oracle Certification Environment for Storage program.

PS Series Firmware Version 2.0 or later.

## Table of Contents

---

Introduction .....	1
Basic Steps .....	2
PS Series Group Requirements .....	3
Creating Volumes .....	3
Controlling Server Access to Volumes .....	4
Network Switch Configuration Recommendation .....	6
Oracle9i Database Server Configuration .....	6
NFS Server Requirements .....	6
Connecting to PS Series Group Volumes .....	7
Setting Up the Database Servers as NFS Clients .....	8
Installing and Configuring Oracle9i Database .....	8
Configuring the Environment for High Availability .....	8
PS Series Storage Arrays .....	9
Network Switches .....	11
Servers .....	11
More Information and Customer Support .....	13

## Table of Figures

---

Figure 1: PS Series Storage Array .....	1
Figure 2: Basic Configuration .....	2
Figure 3: PS Series Group Manager – Volume Window .....	4
Figure 4: PS Series Group Manager – Volume Access Window .....	5
Figure 5: Highly Available Configuration .....	9
Figure 6: Recommended Dual Control Module Array Network Configuration .....	11
Figure 7: Multiple Network Paths to a PS Series Storage Array .....	12



## Introduction

---

This Technical Report describes how to use a PS Series group to provide NFS server storage space to servers running Oracle9i Database. It includes group and server requirements, describes Oracle9i and NFS configurations, and provides recommendations for high availability operation.

A PS Series group is an iSCSI storage area network (SAN) comprised of one or more PS Series storage arrays (group members) connected to an IP network and managed as a single storage system. This self-managing iSCSI SAN is easy to use and affordable, regardless of its size, and has fully redundant hardware. By automating complex tasks like RAID configuration and performance optimization and providing fast, flexible scalability, a PS Series group can dramatically reduce the cost of storage acquisition and management.

**Figure 1: PS Series Storage Array**

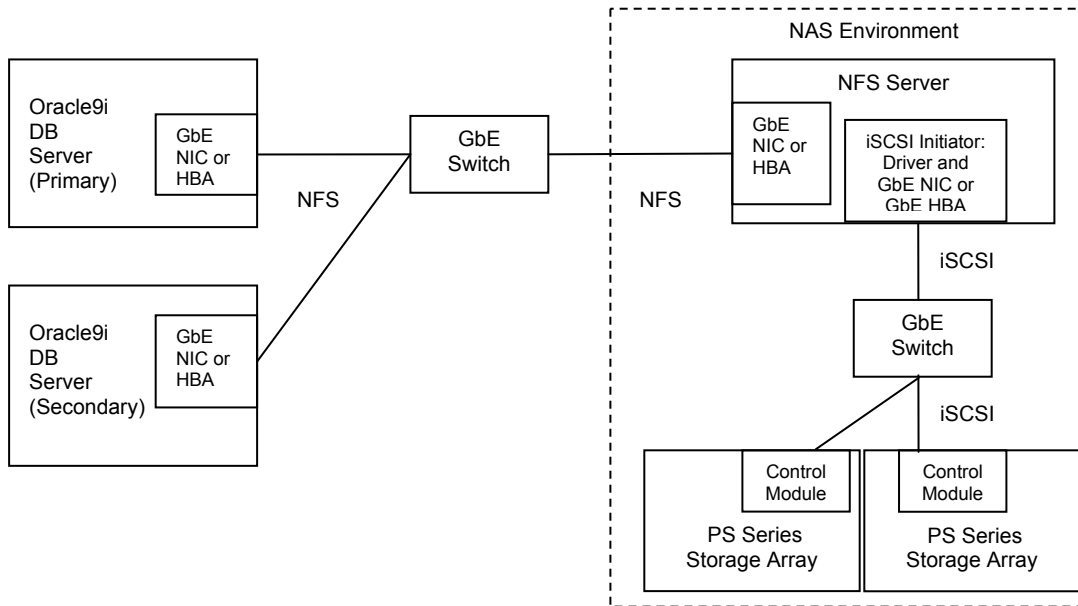


Oracle9i Database and a PS Series group can be used in a variety of environments. Oracle9i is supported by many operating systems, clustering configurations, and storage technologies. Because a PS Series group is based on IP, SCSI, and iSCSI protocols and provides block storage, it can be used in a wide variety of configurations. A server needs only an industry-standard iSCSI initiator to access data stored in the group.

This Technical Report was created to satisfy the requirements of the Oracle Certification Environment (OCE) for Storage testing. The testing was completed using a configuration that included Oracle9i Database running on Solaris NFS servers connected to a PS Series group through a standard Ethernet network. In addition, the testing focused on a basic Oracle9i Database configuration with a primary (active) server and a secondary (failover) server.

The following figure shows a basic configuration. See *Configuring the Environment for High Availability* for information about setting up a highly available environment using Oracle9i Database and NFS servers and a PS Series group.

**Figure 2: Basic Configuration**



## Basic Steps

Follow these basic steps to set up Oracle9i Database servers to access PS Series group volumes through an NFS server:

1. Set up the PS Series group and create the volumes for the Oracle9i database files and log files. For each volume, be sure to create at least one access control record that allows the NFS server access to the volume. See *PS Series Group Requirements* for more information.
2. Ensure that each network switch in the network path between servers and the PS Series group meets the recommendation described in *Network Switch Configuration Recommendation*.
3. Ensure that each server that will run Oracle9i Database meets the requirements described in *Oracle9i Database Server*.
4. Ensure that the NFS server meets the requirements described in *NFS Server Requirements*.
5. On the NFS server, connect to the volumes, create file systems on the volumes, and then mount the volumes. See *Connecting to PS Series Group Volumes* for more information.
6. Set up the database servers as NFS clients. See *Setting Up the Database Servers as NFS Clients* for more information.
7. Install Oracle9i Database on the database servers, specifying the NFS mount points as the database and log file locations. See *Installing and Configuring Oracle9i Database* for more information.

If desired, you can implement hardware and software options to ensure reliable access to data., as described in *Configuring the Environment for High Availability*.

## PS Series Group Requirements

---

For information about setting up and managing a PS Series group, including creating volumes and access control records, see the PS Series *QuickStart* or *Group Administration* manual.

PS Series group requirements are as follows:

- PS Series Firmware Version 2.1 or a higher version.
- For each group member (PS Series storage array), at least one active network connection. Configure multiple network connections for higher performance and availability.
- Volumes that will be used by the NFS server for storing Oracle9i database files and the log files. Although it is not required, Oracle recommends that you store database files and log files on *separate* volumes for fault isolation and security. See *Creating Volumes* for more information.
- For each Oracle9i database file volume and log file volume, at least one access control record that will allow the NFS server access to the volume. See *Controlling Server Access to Volumes* for more information.

In addition, it is recommended that you consult the *Network Connection and Performance Guidelines* Technical Report on the EqualLogic Customer Support website for information about the best way to connect arrays to the network and how to improve network performance.

## Creating Volumes

---

A volume in a PS Series group represents a logical portion of group storage space. Before creating volumes, be sure to fully understand the capacity requirements for the Oracle9i database files and log files so you can allocate a sufficient amount of storage space to each volume. You cannot decrease the size of a volume.

**Note:** Although it is not required, Oracle recommends that you store database files and log files on *separate* volumes for fault isolation and security.

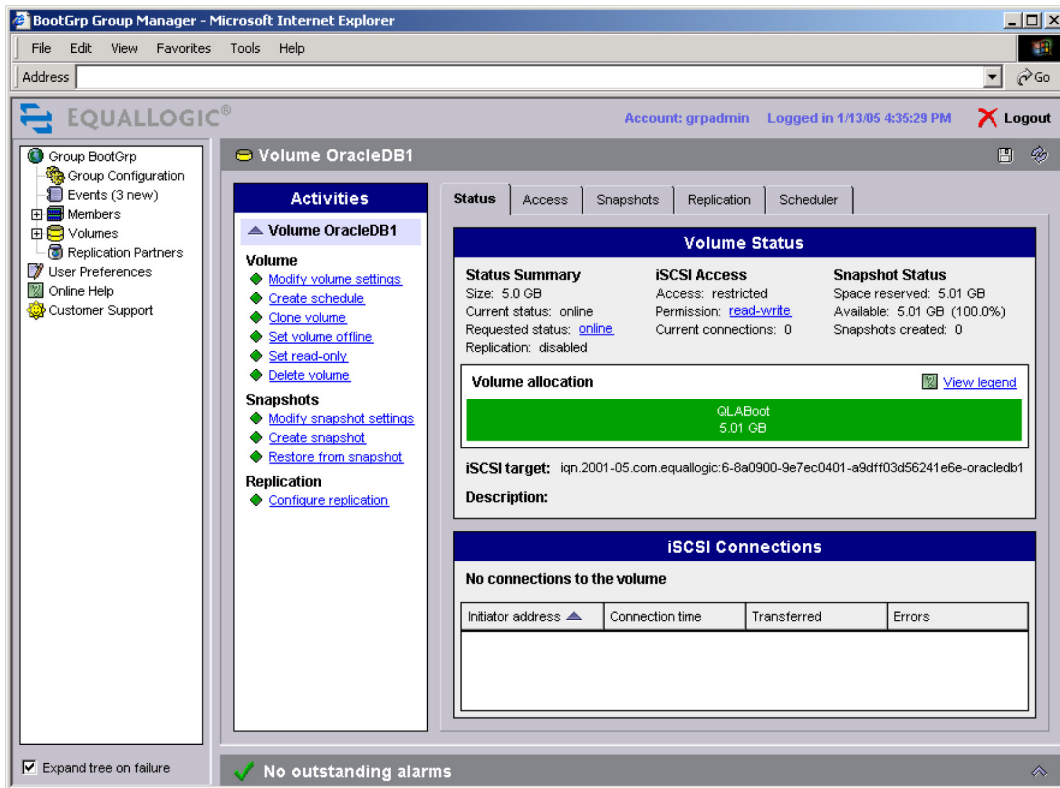
When you use the Group Manager GUI or CLI to create a volume, you must specify:

- Unique volume name, up to 64 alphanumeric characters.
- Size of the volume.
- Optionally, snapshot space to reserve, if you want to create snapshots of the volume.

The group will automatically generate an iSCSI target name for the volume; the volume name will be appended to the end of the target name. Servers connect to the volume through the target name.

The following figure shows the Group Manager GUI window that describes a volume named OracleDB1.

**Figure 3: PS Series Group Manager – Volume Window**



## Controlling Server Access to Volumes

A PS Series group uses access control records to prevent unauthorized access to volumes and snapshots by explicitly permitting access by some servers, while denying access to others. A volume and its snapshots share a list of access control records (sometimes called an access control list). A record can apply to the volume, the volume snapshots, or both. For example, you could create one record that allows a server access to the volume and the volume snapshots and create another record that allows a different server access only to the volume snapshots.

In each access control record, you can specify an IP address, iSCSI initiator name, CHAP user name, or any combination of the three requirements. An NFS server must match *all* the requirements in *one* record in order to access the volume or snapshot. For example, if a record includes an IP address and an iSCSI initiator name, a host must present both the IP address and the initiator name in order to match the record's requirements.

For each database file volume and log file volume, you must create at least one access control record that will permit volume access by the NFS server.

**Notes:** For troubleshooting purposes, you can create an access control record that puts no access restrictions on a volume or its snapshots. However, this is not recommended.

If you are using IP addresses or iSCSI initiator names to restrict access, be sure to create an access control record for each IP address or initiator name presented by a server. As an example, for each NIC that is supposed to handle iSCSI traffic on a server, create a record

and specify the IP address assigned to the NIC. This will ensure that the server can access the volume or snapshot, regardless of which NIC is used for the connection.

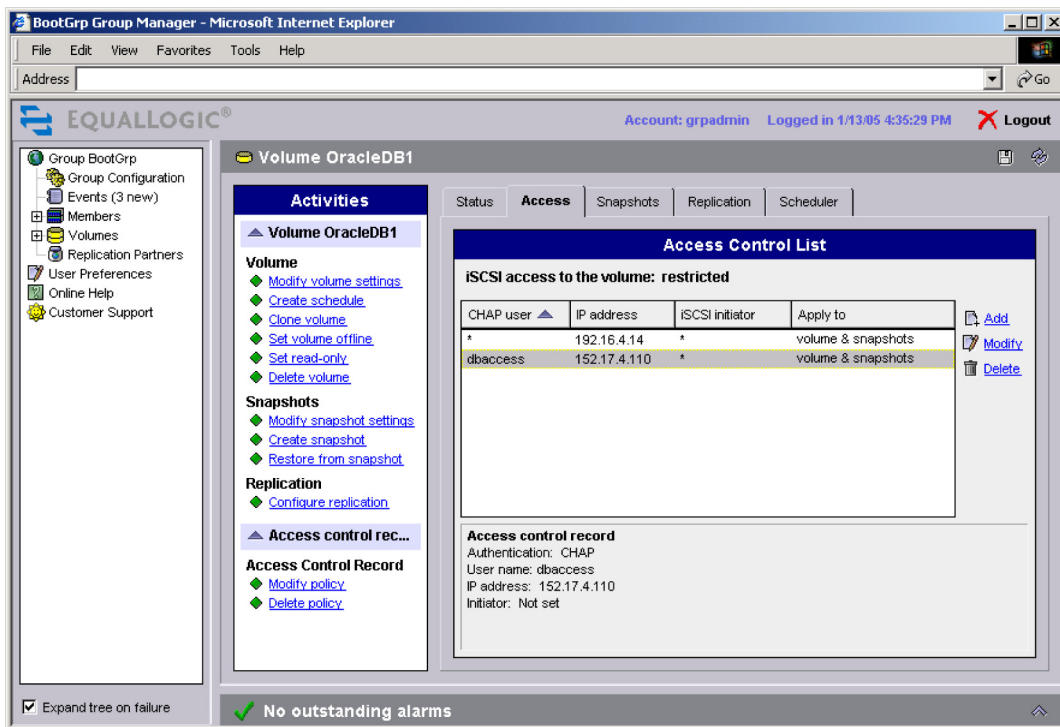
If you use CHAP to authenticate initiator access to a volume, it is recommended that you also specify an IP address in the access control record. If you only use CHAP, initiators that support discovery will attempt to log in to the target, even if they do not have the right CHAP credentials, resulting in a large number of events logged in the group and an inefficient use of resources.

If you are using DHCP to provide IP addresses to server network interfaces and also using IP addresses in access control records, server access to a volume will fail if the DHCP server changes the IP address assigned to a network interface to one that does not match an access control record.

When using the Group Manager GUI to create a volume, you are given the opportunity to create an access control record at that time. In addition, you can use the GUI or CLI to create, display, and modify records for a volume and its snapshots at any time.

The following GUI window lists the access control records for the volume OracleDB1.

**Figure 4: PS Series Group Manager – Volume Access Window**



## Network Switch Configuration Recommendation

---

It is recommended that you use Gigabit Ethernet network switches for the highest performance. The Oracle Certification Environment for Storage test on which this Technical Report is based was completed using the following switch:

Cisco Catalyst 3750G-48TS - 48 Ethernet 10/100/1000 ports

However, any standard Gigabit Ethernet switches can be used to connect the Oracle9i Database servers to the NFS server and to connect the NFS server to the arrays in a PS Series group.

When selecting a switch, performance is the most important consideration, because not all Gigabit Ethernet switches can run all ports at line speed, which is recommended.

In addition, consult the *Network Connection and Performance Guidelines* Technical Report on the EqualLogic Customer Support website for information about switch settings that can improve network performance.

## Oracle9i Database Server Configuration

---

For Oracle9i Database installation and configuration documentation, see:

<http://www.oracle.com/technology/documentation/oracle9i.html>

The Oracle Certification Environment for Storage test on which this Technical Report is based was completed using the following database server configuration, however, the concepts may be easily adapted to other operating systems and hardware:

- Oracle9i 9.2.0.1.0. See *Installing and Configuring Oracle9i Database* for more information.
- Sun<sup>®</sup> SunFire V250 running the Solaris<sup>™</sup> 9 operating system.
- Gigabit Ethernet NIC or HBA.
- Connection to PS Series group volumes. See *Connecting to PS Series Group Volumes* for more information.

The Oracle Certification Environment for Storage test currently supports the use of NFS only on certain operating systems, including AIX, HP-UX, HP Tru64 UNIX, Linux (Intel Only), and Solaris Operating System (SPARC).

In addition, consult the *Network Connection and Performance Guidelines* Tech Report on the EqualLogic Customer Support website for information about operating system and initiator settings that can improve network performance.

## NFS Server Requirements

---

You can utilize any standard NFS server with Oracle9i Database and a PS Series group.

For the purposes of this Technical Report, the NFS server is simply a server with multiple NICs, one connected to the IP LAN and accessed by the database servers through NFS, and the other connected to the IP SAN for access to the PS Series group.

Requirements for an NFS server include the following:

- One or more industry-standard iSCSI initiators.
- One or more Gigabit Ethernet NICs or HBAs.

There are various iSCSI initiators available, including both software and hardware initiators:

- Software initiators are often provided by the operating system in the form of a driver and can be used with a standard Ethernet NIC. For optimal performance, it is recommended that you use a Gigabit Ethernet NIC. Consult the driver documentation for installation instructions.
- Hardware initiators are iSCSI host bus adapters (HBAs). Vendors produce drivers for various operating systems. For optimal performance, use Gigabit Ethernet HBAs that offload TCP/IP, iSCSI, and other Ethernet functions, including out-of-order packet processing, resulting in decreased server overhead.

Consult the initiator vendor documentation for installation instructions. Always check the Knowledge Base on the EqualLogic Customer Support website for the latest initiator information.

## **Connecting to PS Series Group Volumes**

---

To connect the NFS server to the PS Series group volumes for the Oracle9i database and log files, follow the iSCSI initiator vendor instructions for discovering and connecting to iSCSI targets (volumes). If the volume uses CHAP for authentication, you must provide the correct CHAP user name and password to the initiator.

Once connected to a volume's iSCSI target, the volume will appear as a new disk in the NFS server's disk configuration utility. Then, use the normal operating system procedures to perform these tasks:

1. Format and build a new file system on each device
2. Mount each file system. Be sure to make the mount persistent across server reboots.
3. Test that you can read and write to the mounted file systems.
4. Use NFS to share the file systems with the Oracle9i Database servers (NFS clients). Be sure that the shares are persistent across server reboots.
5. Ensure that the NFS service is started. The NFS service must start after the iSCSI connection is established when the server boots.

If you cannot connect to a volume, be sure the server is allowed access through one of the volume's access control records, as described in *Controlling Server Access to Volumes*.

## Setting Up the Database Servers as NFS Clients

---

To set up a database server to be an NFS client, perform these steps on the database server:

1. Create a directory mount point for each database and log file NFS share.
2. Mount the file systems for the database files and log files. NFS offers a number of options for mounting using the `-o` switch with the `mount` command. It is recommended that you specify the following options:
  - `rw` – Mounts the directory read-write.
  - `hard` – Continues to retry requests until the server responds.
  - `suid` – Allows `setuid` execution.
  - `intr` – Allows keyboard interrupts to kill a process that is hung while waiting for a response on a hard-mounted file system.

In addition, ensure that the mounted shares are persistent across system reboots.

3. Test that you can read and write to the new database and log file NFS shares.

## Installing and Configuring Oracle9i Database

---

See the Oracle9i Database installation instructions for detailed information about pre-installation requirements, installation procedures, and post-installation requirements. Be sure to carefully follow these requirements on the database servers.

The Oracle9i Database installation requires at least one mount point for the software and at least one mount point for the database and log files. The installation program will display an option to create a sample database (recommended). When prompted for the “Directory for Database files,” enter the path to the mount point of an NFS share to store the database file.

## Configuring the Environment for High Availability

---

In a minimum configuration, each database and NFS server has only one NIC or HBA connected to a switch and each PS Series storage array has only one network interface connected to a switch, presenting several possible points of failure, as shown in Figure 2. Many datacenter environments need highly available data to meet their service level requirements.

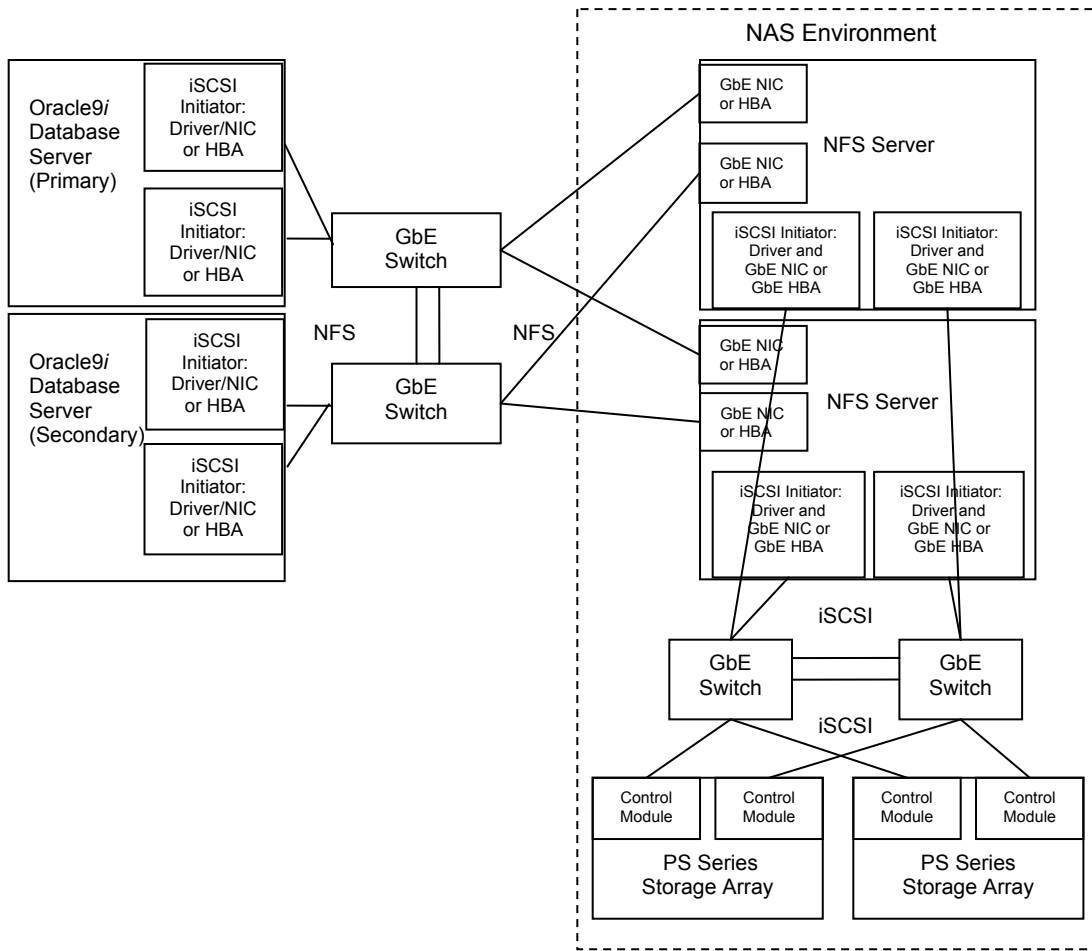
For high availability, you need complete redundancy for each component in the environment:

- PS Series storage array. Configure multiple network interfaces and connect them to different switches.
- Switch. Connect servers and arrays to redundant switches.

- NFS server. Configure multiple paths between the NFS server and PS Series storage arrays and between the NFS server and the database servers. You can also use clustering software on the NFS server.
- Oracle9i Database server. Configure multiple paths between the database servers and the NFS server. You can also use clustering software on the database servers.

See Figure 5 for an example of a highly-available configuration.

**Figure 5: Highly Available Configuration**



The following sections describe how to achieve high availability in an Oracle9i Database environment with a PS Series group.

## PS Series Storage Arrays

Each PS Series storage array is designed to meet and exceed the strenuous requirements of the data center. EqualLogic engineered fault tolerance into every aspect of the array hardware design. The

redundancy in each storage array eliminates any single point of failure, with a net result of greater than 99.999% availability.

Important high availability features integrated in each PS Series storage array include:

- **Hardware redundancy.** All hardware components are fully redundant and hot swappable, including control modules, fan trays, and power supplies.

Each hot-swappable control module features high-performance, dual-core 64-bit processors with a high-speed HyperTransport™ I/O bus and twin 64-bit DDR channels. A control module is equipped with one GB of high speed battery-backed DRAM. Control module caches are battery-backed and mirrored for ultimate protection without compromising performance.

Each hot-swappable disk drive is interconnected with its own independent, hot-swappable serial channel. Additionally, each disk drive is supported mechanically with an inertial dampening chassis that eliminates drive vibrations. Data on an array is protected by enterprise-class RAID technology (RAID 10 or RAID 50) and automatic hot-sparing. Automatic disk reconstruction, accelerated rebuild times, advanced stripe integrity algorithms, and automatic set expansion and geometry transformation are also provided, in addition to patent-pending predictive media-error detection and correction functionality

- **Network connection failover.** Each control module has three network interface ports, labeled `eth0`, `eth1`, and `eth2`. A dual control module array provides three *pairs* of network interfaces (for example, `eth0` on CM1 and `eth0` on CM2). In a dual control module array, the `eth0` and `eth1` network interface port pairs are redundant, so the active control module can access both ports in the pair.

For the highest availability, connect multiple network interfaces to multiple switches. If one network interface fails, iSCSI initiators that were connected to the failed interface can reconnect to the group IP address and be redirected to a functioning network interface, provided the active control module can access the port. I/O operations can then resume through that interface.

For example, in a dual control module array, if the `eth0` and `eth2` ports on CM0 and the `eth1` port on CM1 are connected to a network and `eth0` fails, initiators that were connected to `eth0` can be redirected to `eth1`. However, initiators will be redirected to `eth2` only if CM0 is the active control module.

- **Control module failover.** In a dual control module array, if the active control module fails, the secondary control module can automatically will take over and become active, depending on the type of failure. If a cable is connected to either port in a redundant pair (`eth0` or `eth1`), I/O may continue through that network interface without interruption, regardless of which control module is active. However, the `eth2` network interface will be usable only if a cable is connected to the `eth2` port on the active control module.

Control module failover is transparent to applications, but iSCSI initiators must reconnect to the group IP address.


- **Redundant port failover.** In a dual control module array, if both ports in a redundant pair (`eth0` or `eth1`) are connected to a network and a connection failure occurs (for example, the cable is disconnected, the port fails, or the network switch port fails), network I/O will fail over

to the functioning port in the pair.

Redundant port failover is transparent to iSCSI initiators.

The following figure shows how to optimally connect cables to a dual control module PS Series storage array.

**Figure 6: Recommended Dual Control Module Array Network Configuration**



Refer to the *Network Connection and Performance Guidelines* Tech Report for more details on how to connect network cables in a way that maximizes network availability and takes full advantage of the hardware redundancy in a storage array.

---

## Network Switches

Using redundant switches in your environment ensures that I/O between the Oracle9i Database servers, the NFS server, and the PS Series group continues even if a failure in one of the switches occurs. In addition, some switches support technologies like trunking for improved availability and performance.

See your switch documentation for information about highly available configurations.

---

## Servers

There are several ways to increase availability in an Oracle9i Database server and NFS environment:


- Configure multipath access between database servers and the NFS server and between the NFS server and the PS Series group.
- Configure clustering software on the database or NFS servers.

You can implement high availability network access by using an IP network multipathing solution. There are various solutions available. Some solutions leverage standard operating system IP failover techniques and do not require special Ethernet hardware, software, or switches.

If you configure a multipathing solution on the Oracle9i Database servers and the NFS server and set up multiple network paths to a device, and a path failure occurs (for example, a NIC, cable, or switch fails), network I/O will fail over to a viable path, without disrupting service.

The following figure shows multiple network paths from an NFS server to a PS Series storage array.

**Figure 7: Multiple Network Paths to a PS Series Storage Array**



Usually, you can configure network multipathing in the two ways:

- An *active-active* configuration, where all the paths can carry data simultaneously to the PeerStorage Array. If a path fails, the alternate path automatically handles all I/O.
- An *active-passive* configuration, where only one path is active at any time. If that path fails, the alternate path automatically assumes the role of the active Ethernet connection.

For information about implementing highly-available network access from a Sun Solaris 8 or Solaris 9 server to a PS Series group, see the *Sun™ Solaris™ Server: Implementing Highly Available Network Access* Technical Report on the EqualLogic Customer Support website.

EqualLogic also provides a multipathing solution for Windows® servers. See the EqualLogic Customer Support website for information.

In addition, for server redundancy, you can use clustering technology to ensure that a server is always available to clients. There are many clustering solutions available for a variety of operating systems and applications.

Oracle offers the Oracle Real Application Clusters (Oracle RAC) product, which allows Oracle Database servers to run any packaged or custom application unchanged across multiple connected (or “clustered”) servers. Oracle RAC provides the highest level of availability and flexible

scalability. If a clustered Oracle9i Database server fails, Oracle9i continues running on the remaining servers. If you need more processing power, simply add another server without taking users offline. With Oracle RAC, you can coordinate the use of large numbers of servers acting as one logical server on an enterprise grid, and your data center can deliver the highest quality of service on low-cost, modular hardware.

## **More Information and Customer Support**

---

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

Visit the EqualLogic Customer Support website, where you can download the latest documentation and firmware. You can also view FAQs, the Knowledge Base

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 Support website, from the EqualLogic website ([www.equallogic.com](http://www.equallogic.com)), click **Support** and log in to a support account. If you do not have an account, create one by clicking the link under the login prompt.

To contact customer 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 customer support team.