iSCSI Storage Solution Using Cisco Catalyst 4900 Series Switches and Dell™ EqualLogic™ PS Series SAN Arrays

What You Will Learn

For Ethernet-based server-storage connectivity, Dell™ EqualLogic™ PS Series SAN arrays offer best-in-class solutions for Small Computer System Interface over IP (iSCSI) storage in conjunction with their host-based data protection software.

The Dell EqualLogic PS Series arrays provide a broad set of enterprise-class features for both small and large enterprises. These features include a scale-out architecture that allows the SAN to grow as the needs of the customer grow. Virtualized storage simplifies overall storage management, providing advanced data protection features such as snapshots, clones, and remote replication with other Dell EqualLogic PS Series SAN groups. Host-based software included as part of the overall solution provides additional integration with leading operating systems and applications. Together these features provide a complete, virtualized storage solution without requiring the customer to purchase any additional licenses or software.

The Cisco Catalyst® 4900 Series Switches are designed to deliver high-performance, low-latency, wire-speed Layer 2 to Layer 4 services in a small form factor. The main benefits of these Cisco® switches are:

- **Architectural flexibility:** Offers increased link utilization without buffer overflow plus Layer 2 and 3 support at the access layer as needed
- **Investment protection:** Allows smooth transition from 10/100/1000-Gbps to 10 Gigabit Ethernet server interfaces
- **Simplified management:** Helps enable edge-switch autoprovisioning
- **Transport flexibility:** Supports a wide range of storage (network-attached storage [NAS], iSCSI, etc.), high-performance computing, and video applications

The Cisco Catalyst 4900 Series offers two models for this solution: the Cisco Catalyst 4948 10 Gigabit Ethernet Switch and the Cisco Catalyst 4900M Switch.

The Cisco Catalyst 4948 10 Gigabit Ethernet Switch is a wire-speed, low-latency, fixed-configuration switch designed for rack-optimized server switching. The switch offers 48 ports of wire-speed 10/100/1000BASE-T with 2 ports of wire-speed 10 Gigabit Ethernet (using X2 optics). The platform delivers exceptional reliability and serviceability through optional internal AC or DC 1+1 hot-swappable power supplies and a hot-swappable fan tray with redundant fans. These switches provide wire-speed connectivity for the Dell EqualLogic PS Series iSCSI storage arrays.

The Cisco Catalyst 4900M represents a revolutionary advance in data center top-of-rack Ethernet switching. It packs a 320-Gbps switching architecture into just two rack units (2RUs) of space and provides modularity for flexible deployment. The Cisco Catalyst 4900M provides flexibility in the transition from Gigabit Ethernet to 10 Gigabit Ethernet data center infrastructure by supporting 10 Gigabit Ethernet fiber-optic connections as well as Gigabit Ethernet fiber-optic and copper interfaces. These switches provide wire-speed Gigabit Ethernet and 10 Gigabit Ethernet connectivity for Dell EqualLogic PS Series storage arrays.

Together, the Dell EqualLogic PS Series arrays and the Cisco Catalyst 4900 Series switches provide an excellent, high-performance SAN solution with low total cost of ownership (TCO). This document describes how to integrate the Dell EqualLogic PS Series array into a network infrastructure consisting of Cisco Catalyst 4900 Series Switches and provides customer recommendations.

Basic Setup

Basic Setup of Dell EqualLogic PS Series Array

Basic setup of a Dell EqualLogic PS Series array consists of a few simple steps, some of which will need to be performed on a host computer that is also attached to the SAN. Depending on the operating system being used on the servers that will be attached to the SAN, the storage administrator will use either the Remote Setup Wizard included with the Dell EqualLogic Host Integration Toolkit for Microsoft Windows (if the server is running Microsoft Windows) or a simple terminal emulation program and a serial cable (for all other operating systems). Since the Cisco Catalyst 4900 Series Switch configuration will need to be performed using command-line operations through terminal emulation, this document assumes that the command-line interface (CLI) for will be used to configure the Dell EqualLogic arrays as well. For information about performing initial configuration from a host that is using the Remote Setup Wizard, please refer to the array documentation provided with the system or available online.

The basic steps for configuring a Dell EqualLogic array (and the Dell EqualLogic SAN group if this is the first array being configured) are presented here.

1. Install the arrays in a suitable data center rack.

   Depending on the type of rack in the data center, the actual steps required will vary. Dell EqualLogic PS Series arrays require a standard 19-inch data center or telecommunications rack. Hardware is provided to accommodate most varieties of rack design. Please review the installation instructions that ship with the array or are provided online for details regarding the specific rack type in the data center.

2. Connect the required cables.

   Each array typically is configured with two controllers and two power supplies to increase uptime of the arrays through redundancy. Each controller supports two to four Ethernet ports, depending on the model chosen, that can be configured to communicate using the iSCSI protocol. Also, depending on model, each controller either has a dedicated port or can have one of its ports configured as a dedicated management port that will not respond to iSCSI requests. Figure 1 illustrates the connections available on the back of the various array types offered by Dell.

Figure 1. Dell EqualLogic Gigabit Ethernet Array Models

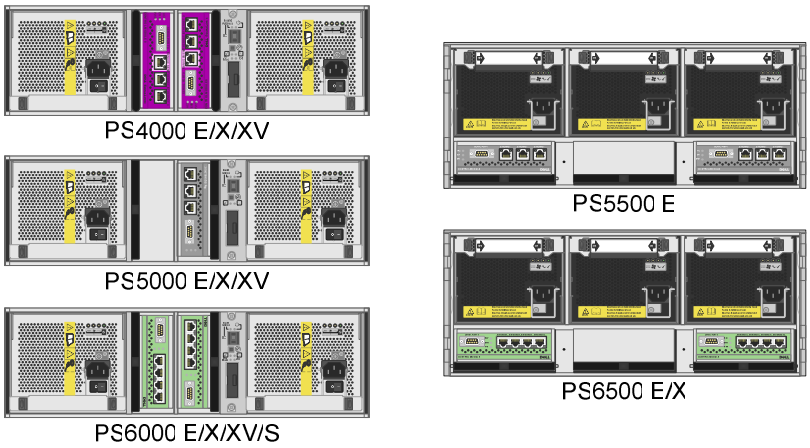

To properly cable an array to provide a fully redundant storage solution, the following connections are required:

- A minimum of two independent power sources for Dell EqualLogic PS4000, PS5000, and PS6000 Series systems; or a minimum of three independent power sources or two independent power sources and an automatic transfer switch for the Dell EqualLogic PS 5500 and PS 6500 Series systems
- A minimum of two switches
- A minimum of two active Ethernet ports from each controller in each array

See the *Dell EqualLogic Installation and Setup Manual* for detailed information about connecting network cables.

Figure 2. High-Availability Cabling Example Using the Cisco Catalyst 4900M

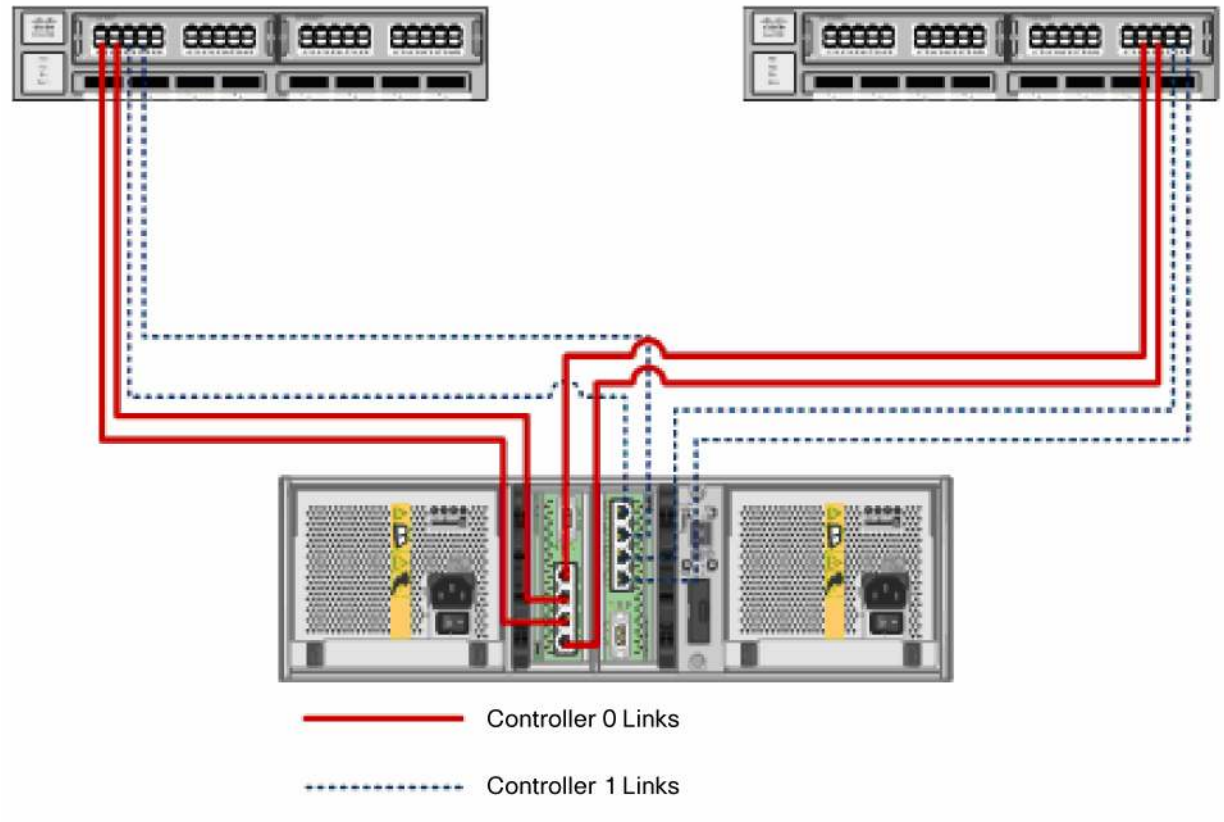

3. Initialize the arrays.

   Two methods for initializing a Dell EqualLogic array are available, depending on the operating system running on the host or management workstation connected to the SAN. If the host is running a version of the Microsoft Windows operating system, installing the Dell EqualLogic Host Integration Toolkit and subsequently running the Remote Setup Wizard from that host will start the initialization process.

   For environments without hosts running Microsoft Windows, Dell provides a comprehensive CLI. This CLI can be used directly on each array through the serial cable connection mentioned earlier. Please refer to the *Dell EqualLogic Installation and Setup Manual* for details.

4. Connect to the storage using the Dell EqualLogic Group Manager to start administering the SAN group.

   After the arrays have been initialized, either by creating a new SAN group or joining an existing SAN group, SAN management operations such as creating volumes, assigning volumes to servers and hosts, and managing data protection, can be easily accomplished using the Dell EqualLogic Group Manager application. Dell EqualLogic Group Manager is a Java Runtime Environment™ (JRE™) application that will run from any supported web

browser such as Internet Explorer® or Mozilla® Firefox®. Please review *the Dell EqualLogic Group Administration Guide* for more information about management tasks.

Basic Setup of Cisco Catalyst 4900 Series Switches

Basic setup of Cisco Catalyst 4900 Series Switches involves assigning an IP address and configuring some basic settings for the switch to communicate with local routers and the Internet. The minimal configuration described here does not cover all features but simply configures basic settings using a Telnet connection from your management network.

IP Settings

To configure the Cisco Catalyst 4900 Series Switch, obtain the following network configuration information from the network administrator:

- Switch IP address
- Subnet mask (IP netmask)
- Default gateway (router)
- Password for the **enable secret** password command
- Password for the **enable** password command
- Telnet password

To perform initial basic configuration on the Cisco Catalyst 4900 Series Switch, follow these steps:

1. At the terminal prompt, enter the **enable** command to enter privileged execution mode.
```
Switch> enable
Password: password
Switch#
```

2. Set the system time using the **clock set** command in privileged execution mode.
```
Switch# clock set 20:09:01 3 Apr 2006
```

3. Verify the change by entering the **show clock** command.
```
Switch# show clock
20:09:06.079 UTC Thu Apr 3 2006
```

4. Enter the `configure terminal` command to switch to global configuration mode.
```
Switch# configure terminal
```

5. Enter configuration commands, one per line. End with CTRL-Z.
```
Switch (config)#
```

6. Configure a host name for the switch and press Return.
```
Switch (config)# hostname Switch1
```

7. Configure the system prompt for the switch and press Return. To remove the new prompt and return the prompt to its default, use the **no prompt** command.
```
Switch (config)# prompt Switch1>
```

8. Use the **banner motd** global configuration command to set location information in the login banner. You can also set a system contact using this command.

```
Switch1(config)# banner motd c 170 West Tasman Drive, San Jose, CA c
```

or

```
Switch1 (config)# banner motd c 170 West Tasman Drive, San Jose, CA;
Tech Support 408 123 4567 c
```

9. Configure an **enable secret** password and press Return.

   The password can be from 1 to 25 alphanumeric characters, can start with a number, is case sensitive, and allows spaces but ignores leading spaces. The **enable secret** password is encrypted, and the **enable** password is in plain text.

```
Switch1 (config)# enable secret SecretPassword
```

10. Configure an **enable** password and press Return.

```
Switch1 (config)# enable password EnablePassword
```

11. Configure a virtual terminal (Telnet) password and press Return.

    The password can be from 1 to 25 alphanumeric characters, is case sensitive, and allows spaces but ignores leading spaces.

```
Switch1 (config)# password terminal-password
Switch1 (config)# line vty 0 15
```

12. Configure the interface that connects to the management network. (The IP address and subnet mask shown are for example purposes only; use an address appropriate for your network.)

```
Switch1 (config)# ip routing
Switch1 (config)# interface gigabitethernet 24
Switch1 (config-if)# no switchport
Switch1 (config-if)# no shutdown
Switch1 (config-if)# ip address 10.4.120.106 255.0.0.0
Switch1 (config-if)# exit
```

13. Exit the global configuration mode.

```
Switch (config)# exit
Switch #
```

14. View the configuration you just created and confirm that it is what you want.

```
Switch1# show run
!
hostname Switch1
!
banner motd ^C
170 West Tasman Drive, San Jose, CA ^C
!
!--- Output suppressed.
```

15. Verify the IP information by using the **show ip interface brief** and **show ip route** commands.

```
Switch1# show ip interface brief
Interface    IP-Address      OK?    Method     Status      Protocol
Vlan1        172.16.1.2      YES    manual        up       up
FastEthernet1     unassigned     YES    unset         up       up
!--- Output suppressed.
Switch1# show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B -
BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - ISIS level-1, L2 - ISIS level-2, ia - ISIS
inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
Gateway of last resort is 172.16.1.1 to network 0.0.0.0
     172.16.0.0/24 is subnetted, 1 subnets
C       172.16.1.0 is directly connected, Vlan1
S*   0.0.0.0/0 [1/0] via 172.16.1.1
Switch1#
```

16. Save the running configuration.

```
Switch1# copy system:running-config nvram:startup-config
```

SAN Network Recommendations

Overview

The data center network design is based on a proven layered approach that has been tested and improved over the past several years in some of the largest data center implementations in the world. The layered approach is the basic foundation of the data center design that seeks to improve scalability, performance, flexibility, resiliency, and maintenance. Figure 3 shows the basic layered design.

Figure 3. Basic Layered Design

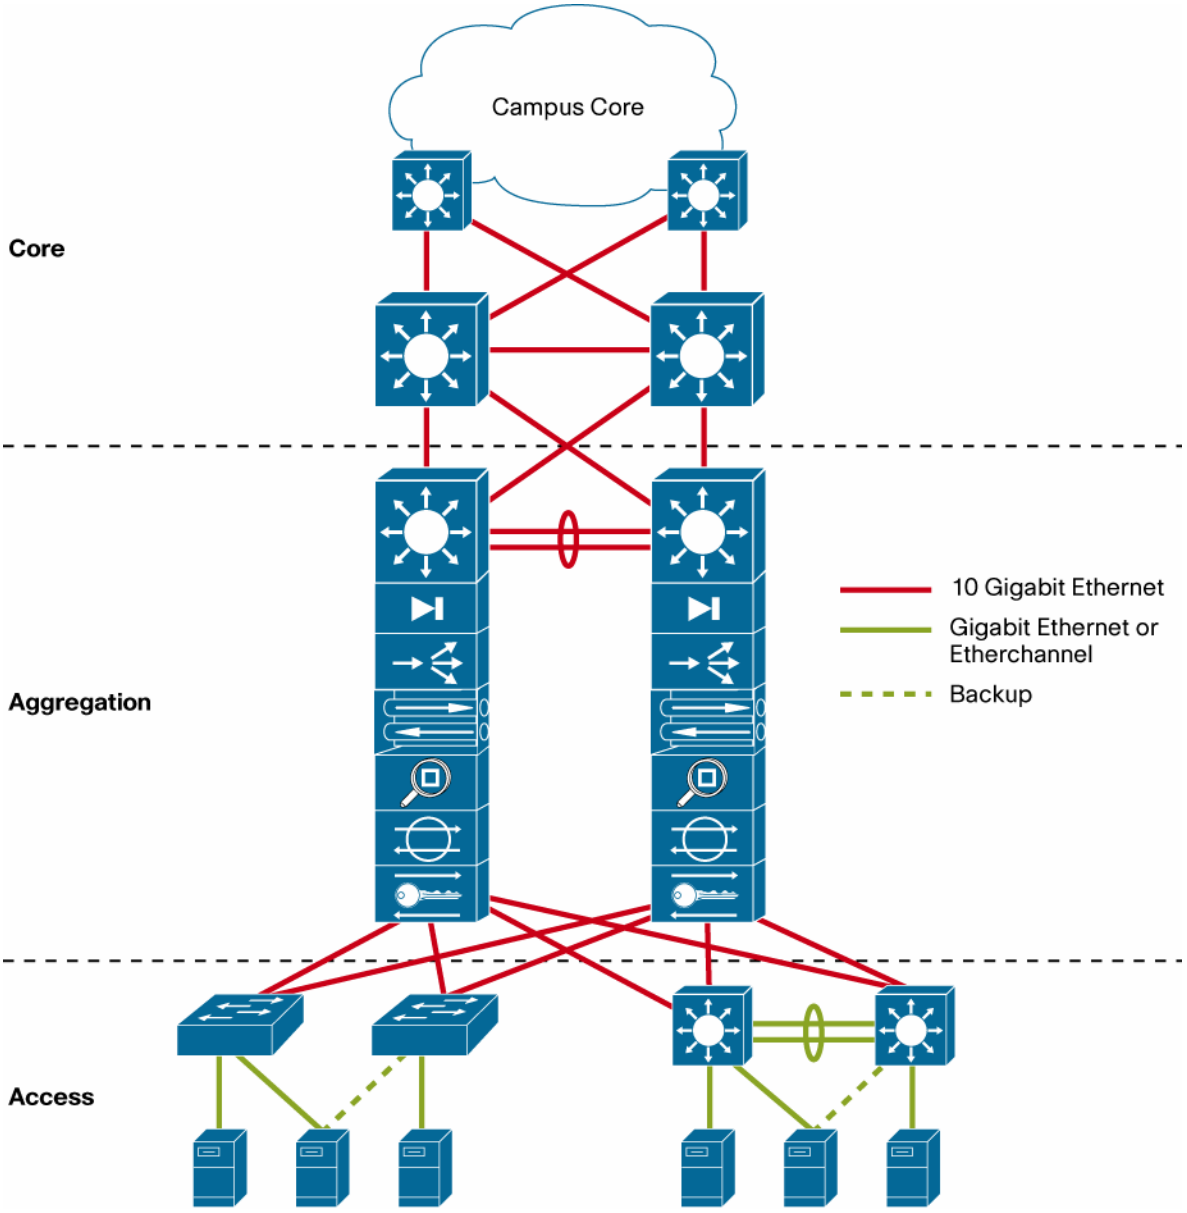

The layers of the data center design are the core, aggregation, and access layers:

- **Core layer:** The core layer provides the high-speed packet switching backplane for all flows going in and out of the data center. This layer provides connectivity to multiple aggregation modules, and it provides a resilient Layer 3 routed fabric with no single point of failure. The core layer runs an interior routing protocol, such as Open Shortest Path First (OSPF) or Enhanced Interior Gateway Routing Protocol (EIGRP), and load balances traffic between the campus core and aggregation layers using Cisco Express Forwarding (CEF/dCEF) hashing algorithms.
- **Aggregation layer:** Aggregation layer modules provide important functions, such as service module integration, Layer 2 domain definitions, spanning tree processing, and default gateway redundancy. Server-to-server multi-tier traffic flows through the aggregation layer and can use services, such as firewall and server load balancing, to optimize and secure applications. The smaller icons within the aggregation layer

switch in Figure 3 represent the integrated service modules. These modules provide services, such as content switching, firewall, SSL offload, intrusion detection, and network analysis.

- **Access layer:** The access layer is where the servers physically attach to the network. The server components can consist of rack-mount and tower servers, blade servers with integral switches, blade servers with pass-through cabling, clustered servers, and mainframes with open systems adapters (OSAs). The access layer network infrastructure consists of modular switches, fixed configuration 1 or 2RU switches, and integral blade server switches. Switches provide both Layer 2 and 3 topologies, fulfilling the various server broadcast domain and administrative requirements.

Figure 3 shows the three-tier model. In a small server and storage environment, a two-tier model consisting of core and access layers is another option. For more information about the features offered by different models, refer to Chapter 3, "Server Cluster Designs with Ethernet," in the *Cisco Data Center Infrastructure Design Guide*: http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_Infra2_5/DCInfra_3.html.

This document describes the recommended configuration for the access layer switches, the Cisco Catalyst 4948 10 Gigabit Ethernet Switch and Cisco Catalyst 4900M Switch, when connecting with Dell EqualLogic PS Series SAN arrays. Figure 4 shows an example of how to connect the Dell EqualLogic PS6000 to the access layer switches. Both the Cisco Catalyst 4948 10 Gigabit Ethernet Switch and Cisco Catalyst 4900M offer line-rate forwarding performance and can achieve various oversubscription levels.

Figure 4. Access Layer Connection with a Dell EqualLogic PS6000 Array-Based SAN

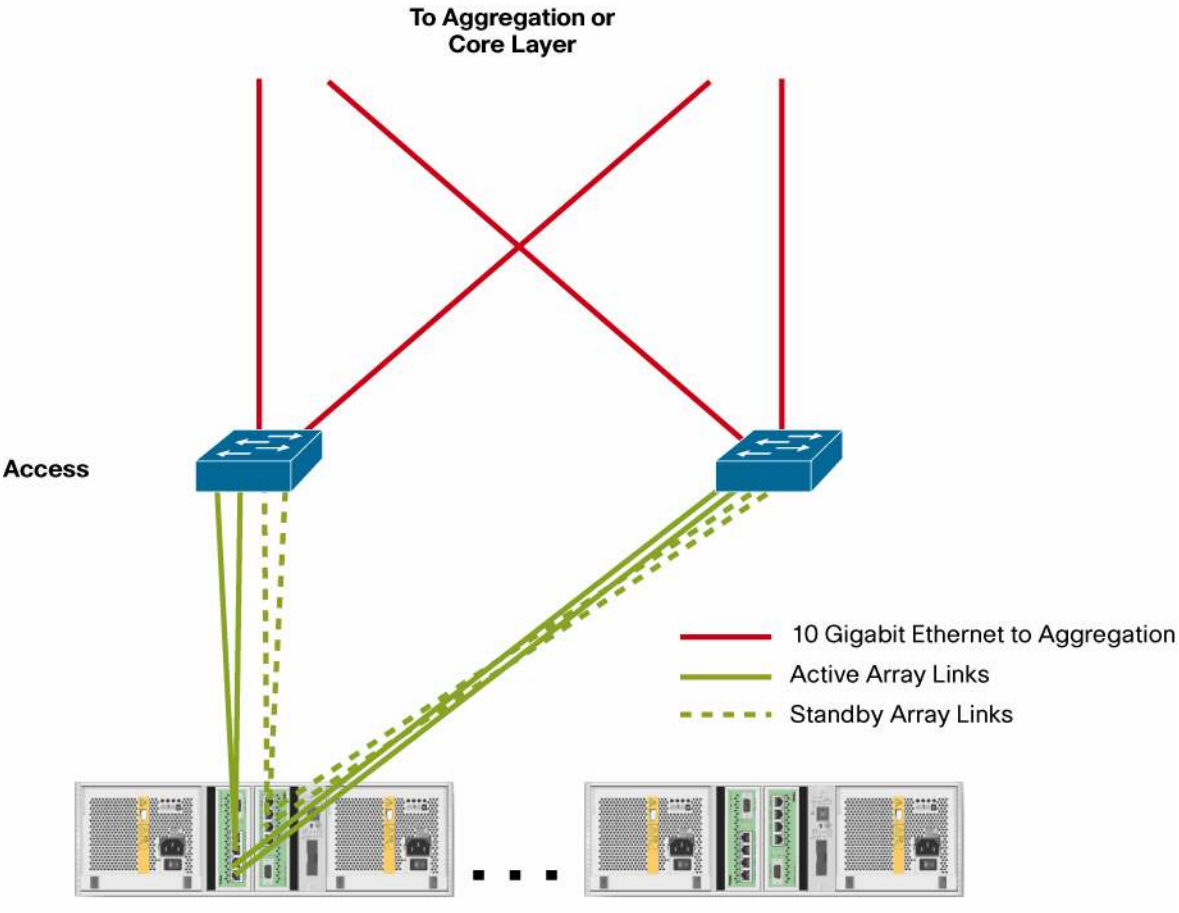

Cisco Catalyst 4948 10 Gigabit Ethernet Switch

With two Cisco Catalyst 4948 10 Gigabit Ethernet Switches, there are 96 10/100/1000 Ethernet ports that can be used to connected to the Dell EqualLogic PS6000 arrays. Each Dell EqualLogic PS6000 array has four active and four standby links, so a pair of Cisco Catalyst 4948 10 Gigabit Ethernet Switches can be connected to 96/8 = 12 Dell EqualLogic PS6000 arrays. With the 10 Gigabit Ethernet uplink, the oversubscription ratio is 2.4:1. If a lower oversubscription ratio is desired, you can connect to fewer Dell EqualLogic PS6000 arrays or use the Cisco Catalyst 4900M to achieve a lower oversubscription ratio.

Cisco Catalyst 4900M

The configuration to connect to a Dell EqualLogic PS6000 array would use two 20-port Gigabit Ethernet half-cards. That configuration yields a total of 40 Gigabit Ethernet ports per Cisco Catalyst 4900M. With two Cisco Catalyst 4900M Switches, there are 80 Gigabit Ethernet ports that can be used to connect to Dell EqualLogic PS6000 arrays. A pair of Cisco Catalyst 4900M Switches can be connected to 80/8 = 10 Dell EqualLogic PS6000 arrays. Each Cisco Catalyst 4900M still has eight 10 Gigabit Ethernet ports that can be used for uplinks. With one 10 Gigabit Ethernet uplink per Cisco Catalyst 4900M, the oversubscription ratio is 2:1, and with two 10 Gigabit Ethernet uplinks, the oversubscription ratio is 1:1.

Switch Setup Recommendations

Use the following recommendation to set up the Cisco Catalyst 4948 10 Gigabit Ethernet Switches and Cisco Catalyst 4900M Switches that will handle iSCSI traffic between servers and the Dell EqualLogic PS Series group members (storage arrays):

- Do not enable Spanning Tree Protocol on switch ports that connect end nodes (iSCSI initiators or storage array network interfaces). If Spanning Tree Protocol or Rapid Spanning Tree Protocol (RSTP; preferable to Spanning Tree Protocol) is used, enable the Cisco PortFast option on each switch port that connects end nodes. PortFast reduces network interruptions that occur when devices restart.

Note: Use of RSTP for a single-cable connection between switches is encouraged, as is the use of trunking for multiple-cable connections between switches.

- Configure flow control on each switch port that handles iSCSI traffic. Dell EqualLogic PS Series storage arrays will correctly generate and respond to flow control if it is enabled on a switch. If the server is using a software iSCSI initiator and network interface card (NIC) combination to handle iSCSI traffic, flow control must also be enabled on the NICs to obtain the performance benefit.
- Configure jumbo frames on each switch that handles iSCSI traffic. If a server is using a software iSCSI initiator and NIC combination to handle iSCSI traffic, enable jumbo frames on the NICs to obtain the performance benefit and help ensure consistent behavior.

Note: Be sure to enable jumbo frames on all SAN components (NICs, iSCSI host bus adapters [HBAs], and switches). Also, to simplify troubleshooting of initial deployments, make sure that NICs, switches, and Dell EqualLogic PS Series storage arrays are fully operational before enabling jumbo frames on the switches and NICs.

- If sharing physical switch resources with other networks (such as the corporate LAN or management network), use VLANs to separate iSCSI SAN subnets from other network subnets.

Understanding Switch and Server Recommendations

Spanning Tree Protocol

For switch ports that connect end nodes (iSCSI initiators or storage array network interfaces), Spanning Tree Protocol should not be used. However, if Spanning Tree Protocol or RSTP (preferable to Spanning Tree Protocol) is to be used for switch ports that connect end nodes, enable the Cisco PortFast option on each switch port. PortFast reduces network interruptions that occur when devices restart.

Note: The use of Spanning Tree Protocol for a single-cable connection between switches or the use of trunking for multiple-cable connections between switches is encouraged.

Spanning Tree Protocol (IEEE 802.1d) is a link management protocol that prevents loops in an Ethernet network by helping ensure that only one active path exists between switches. Upon connection, a switch performs a 30- to 50-second Spanning Tree Protocol calculation to transition ports to a forwarding or blocking state. During this time, no user data passes through the ports. The result is that links take longer to come online. This behavior can cause long failover times when recovering from a network switch or Dell EqualLogic PS Series storage array control module failure.

RSTP (IEEE 802.1w) allows a switch port to bypass the Spanning Tree Protocol listening and learning states and quickly enter the Spanning Tree Protocol forwarding state. To achieve fast convergence on a port, RSTP relies on two variables: the edge port and the link type. Switch ports connecting end nodes (iSCSI initiators and storage array network interfaces) should be configured as edge ports to take advantage of RSTP features. Some switches have port settings that allow the immediate transition of the port to the Spanning Tree Protocol forwarding state upon connection. For example, Cisco enhanced the original IEEE 802.1d STP specification with a proprietary feature called PortFast, which immediately transitions a port to the Spanning Tree Protocol forwarding state upon connection. However, do not connect other hubs, switches, concentrators, or bridges to these ports, because doing so may cause infinite packet loops.

Flow Control

On many networks, there can be an imbalance in the network traffic between the devices that send network traffic and the devices that receive the traffic. This is often the case in SAN configurations in which many servers (initiators) are communicating with storage devices. If senders transmit data simultaneously, they may exceed the throughput capacity of the receiver. When this occurs, the receiver may drop packets, forcing senders to retransmit the data after a delay. Although this will not result in any loss of data, latency will increase because of the retransmissions of dropped packets, and I/O performance will degrade. Flow control allows the receiver to instruct the sender to pause transmission of additional data when the receiver senses that it is being overwhelmed. The receiver does this by sending pause frames to the sender, which causes the sender to stop packet transmission for a short period of time. This pause allows the receiver to process its backlog so it can later resume accepting input. While latency is introduced into the SAN by flow control, it is much smaller when using flow control than when flow control is disabled and packets must be retransmitted.

For a Dell EqualLogic iSCSI SAN, flow control should be enabled on all switch ports that handle iSCSI traffic. In addition, if the server is using a software iSCSI initiator and NIC combination or HBA to handle iSCSI traffic, flow control must be enabled on the NICs and HBAs to obtain the performance benefit. The Cisco Catalyst 4900 Series Switches will respond to pause frames generated by targets and initiators, allowing them to control the overall data flow between servers and storage volumes.

Storm Control

An Ethernet traffic "storm" occurs when an excessively large number of data packets are transmitted, causing excessive network traffic. Many switches have traffic storm control features that prevent ports from being disrupted. Three types of storms can occur on a network: broadcast, multicast, and unicast. Storm management features typically work by disabling the ports that are exhibiting this behavior. Because iSCSI hosts and arrays can at times use more than 80 percent of the available network bandwidth, the switch management software commonly will interpret the high utilization as a unicast traffic storm; therefore, if the switch has the capability to manage *unicast* storm behavior, this feature must be disabled.

The Cisco Catalyst 4900 Series does *not* require any special configuration to manage unicast traffic due to its high-performance characteristics. Therefore, no settings need to be changed. If the network administrator chooses to configure multicast and broadcast storm control settings on Cisco Catalyst 4900 Series Switches, such settings will not affect the iSCSI SAN.

Jumbo Frames

The most important aspect of any SAN is the capability of that SAN infrastructure to store or retrieve as much information as possible in the shortest time possible. Latency incurred due to data storage activities will adversely affect an application's performance. For an iSCSI-based storage infrastructure, depending on the nature of the workload, use of jumbo frames may be the difference between a storage solution with average data throughput and high latency and one with high performance and low latency.

To take advantage of jumbo frames, all devices in the network path between the servers and the Dell EqualLogic PS Series group, including the switches and the NICs and HBAs used to access volumes, must have jumbo frames enabled. Switches configured for jumbo frames will support both standard Ethernet frames and jumbo frames. However, if a NIC or HBA is configured for jumbo frames but the switch is not, inconsistent behavior will be experienced, including connection failure. The switch will function properly when the frames are small, but if the NIC attempts to send frames larger than a standard Ethernet frame (1518 bytes), the switch will not be able to process the frames and will drop them. Also, if some switches are configured for jumbo frames but others are not, similar inconsistent results will occur.

A Dell EqualLogic PS Series storage array supports all frame sizes up to 9000 bytes (9018 including the TCP header). Dell EqualLogic PS Series storage arrays also support path maximum transmission unit (MTU) specification, enabling automatic detection of the maximum frame size between TCP/IP endpoints.

Note: To take advantage of jumbo frames, all devices in the network path between the servers and the Dell EqualLogic PS Series group, including the switches and the NICs and HBAs used to access volumes, must have jumbo frames enabled.

VLANs

VLAN capability gives a switch and the network infrastructure additional flexibility. Using VLANs, the network administrator can create multiple, virtual networks on a single physical switch. Each VLAN can be independently configured to meet the needs of the specific application for which it was created. This capability is especially important for iSCSI SAN implementations, which may need to share physical switch resources with other types of Ethernet traffic, such as the corporate LAN. Traffic isolation enables a number of important benefits, including security features such as port blocking and address filtering and higher performance.

Typically, VLANs are separate subnets in TCP/IP networks, so connectivity between VLANs usually requires Layer 3 routing, which is a capability of the Cisco Catalyst 4900 Series Switches. VLANs also support a variety of network settings, including jumbo frames. See the discussion of jumbo frames earlier in this document for more information.

Note: The Dell EqualLogic PS Series storage arrays do not support VLAN tagging. Therefore, each VLAN will require its own subnet to allow the use of standard Layer 3 routing to provide connections between two VLANs.

Configuring Cisco Catalyst 4900 Series Switches

This section provides information about configuring Cisco Catalyst 4900 Series Switches.

Cisco Switch Interfaces

For Cisco Catalyst 4948 10 Gigabit Ethernet Switches and Cisco Catalyst 4900M Switches, an interface name is composed of three parts that are formatted as follows:

slot_number /port_number

The variables are as follows:

- *slot_number*: The slot number is the number of the switch slot in which the interface module is installed. In a fixed-port configuration such as for a Cisco Catalyst 4948 10 Gigabit Ethernet Switch, the slot number is always 1. Module slots are numbered 1 to 3, from bottom to top and left to right. The Cisco Catalyst 4900M has one fixed slot, which is slot 1. The top-two half-slots are slots 2 and 3, left to right, respectively.
- *port_number*: The port number is the number of the port on the interface module or the switch.

Enabling the Cisco PortFast Option to Configure Spanning Tree Protocol Edge Ports

To configure Spanning Tree Protocol edge ports on Cisco Catalyst 4948 10 Gigabit Ethernet Switches and Cisco Catalyst 4900M Switches, enable the PortFast option. PortFast reduces the time that ports must wait for spanning tree to converge and should be used only on ports connected to end nodes (end stations), including iSCSI initiators and storage array network interfaces.

Caution: Connecting hubs, concentrators, switches, or bridges to a port that has PortFast enabled can cause temporary spanning tree loops. Do not enable PortFast on these ports.

The following example shows how to enable PortFast on Gigabit Ethernet interface 1/1 on the switch:

```
Switch> enable
Switch# config t
Switch(config)# interface gigabitEthernet 1/1
Switch(config-if)# spanning-tree portfast
Switch(config-if)# end
Switch# write
```

To view or confirm PortFast status for a port, use the following command:

```
Switch# show spanning-tree interface gigabitEthernet1/1
```

Note: The spanning-tree portfast default global configuration command can be used to globally enable the PortFast feature on all nontrunking ports.

This example shows how to globally enable PortFast by default on all access ports:

```
Switch# config t
Switch(config)# spanning-tree portfast default
Switch(config)# end
Switch# write
```

Configuring Flow Control

This section describes how to configure flow control on Cisco Catalyst 4948 10 Gigabit Ethernet Switches and Cisco Catalyst 4900M Switches. Flow control must be enabled on each switch port that handles iSCSI traffic.

The following commands shows how to configure Gigabit Ethernet interface 1/1 on the switch to autonegotiate the correct flow control setting with the device to which it is connected:

```
Switch> enable
Switch# configure terminal
Switch(config)# interface gigabitethernet1/1
Switch(config-if)# flowcontrol receive desired
Switch(config-if)# end
Switch# write
```

To view or confirm the flow control status on a port, use the following command:

```
Switch# show flowcontrol interface gigabitethernet1/0/1
```

Note: The **desired** directive is not a flow control option on 10 Gigabit Ethernet interfaces. Use the on directive for 10 Gigabit Ethernet interfaces.

Configuring Jumbo Frames

On Cisco Catalyst 4948 10 Gigabit Ethernet Switches and Cisco Catalyst 4900M Switches, you can set the MTU size for an individual interface, with a maximum value of 9198.

Use the following commands to configure jumbo frames, which sets the maximum packet size to 9198 bytes on Gigabit Ethernet interface 1/1:

```
Switch> enable
Switch# config t
Switch(config)# interface gigabitEthernet 1/1
Switch(config-if)# mtu 9198
Switch(config-if)# end
Switch# write
```

To configure a range of interface ports with jumbo frames, use the following commands:

```
Switch> enable
Switch# config t
Switch(config)# interface range gigabitEthernet 1/1-20
Switch(config-if-range)# mtu 9198
Switch(config-if-range)# end
Switch# write
```

To view or confirm the MTU size for a port, use the following command:

```
Switch# show interface gigabitEthernet 1/1
```

Configuring the Dell EqualLogic PS Series Array

After the basic configuration of the Dell EqualLogic PS Series arrays has been completed, each array will have a single Ethernet port from each controller configured for iSCSI and management functions. To take full advantage of the array and network performance capabilities, you must enable additional Ethernet ports on each array. This step allows each array to use additional bandwidth by using additional Gigabit Ethernet ports on each array's active

controller. Depending on the Dell EqualLogic PS Series array model, each array can be configured to use from to two to four total Ethernet ports on the array controllers. This configuration allows each array to gain up to 3 Gbps additional bandwidth (up to 4 Gbps total per array), and if the Dell EqualLogic SAN group contains more than one array, each volume also can be distributed across multiple arrays, increasing the overall performance of the volume by balancing volume access across multiple controllers.

Note: Each array port enabled must have a different IP address on the SAN. A single Dell EqualLogic PS6000 can have four active ports and will require four IP addresses to enable all the active ports. Each IP address must be unique and must be different from the Dell EqualLogic group IP address.

A second advanced configuration option is to configure one of the ports on each array as a management-only port. This port will not participate in any iSCSI transactions, but will respond only to standard TCP/IP traffic. This configuration option allows administrators to provide out-of-band management capabilities to the SAN. However, using this configuration option may reduce the overall number of ports—and thus the overall bandwidth—available for storage requests. Also note that the Dell EqualLogic PS4000 array family has a dedicated management port.

This section provides the CLI steps to enable these additional ports and to configure the last port on each array as a management-only port. For detailed information about using the Dell EqualLogic Group Manager web-based SAN management tool to make these changes, please consult the *Dell EqualLogic Group Administration Guide*.

Three methods are available for attaching to an array that has already been initialized:

- Use Telnet to attach to an array through TCP/IP.
- Use a serial communications cable to directly connect a workstation or server to each array.
- Use the CLI scripting function for Microsoft Windows or Red Hat Linux.

Note: To use the host scripting tools, install Perl on the host system. On Microsoft Windows platforms, obtain Perl from ActiveState (http://aspn.activestate.com/ASPN/Downloads/ActivePerl/) or as part of the cygwin distribution (http://www.cygwin.com).

Configuring Additional Ethernet Ports for iSCSI

CLI Commands

After connecting to an array in the SAN group and logging in, use the following steps to configure each additional port for iSCSI traffic. (For more information, please read the *Dell EqualLogic PS Series Storage Arrays CLI Reference* available on the Dell and EqualLogic websites.)

1. Configure the second Ethernet port (eth1) on Dell EqualLogic PS6000XV array “m5” with IP address 192.168.137.102 with subnet mask 255.255.255.0 and then activate the port with the following commands:

```
> member select m5 eth select 1 ipaddress 192.168.137.102 netmask 255.255.255.0
> member select m5 eth select 1 up
```

2. Perform these steps for each primary controller port in the array (standby controller ports automatically inherit the settings of the same port on the primary controller) that is to be activated.

EqualLogic Group Manager GUI

Using the web-based Dell EqualLogic Group Manager, you can easily configure or modify a network interface:

1. Choose Members > *member_name* and click the Network tab.
2. Select the interface to be configured and choose “Modify IP settings” from the Activities menu (Figure 5).

Figure 5. Selecting the Ethernet Interface on the Dell EqualLogic Array

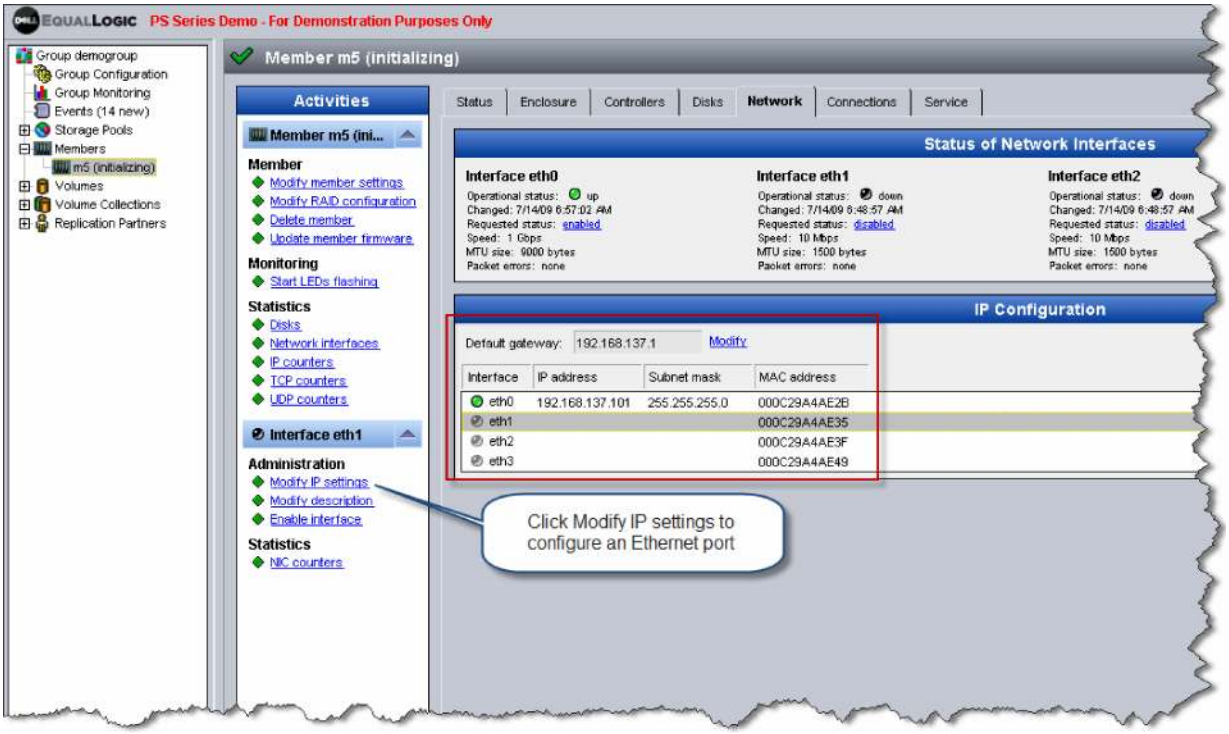

3. Enter the requested information for the IP address and subnet mask (Figure 6).
4. Check the “Enable this interface” check box.
5. Click OK.

Figure 6. Entering Ethernet Interface Information for the Dell EqualLogic Array

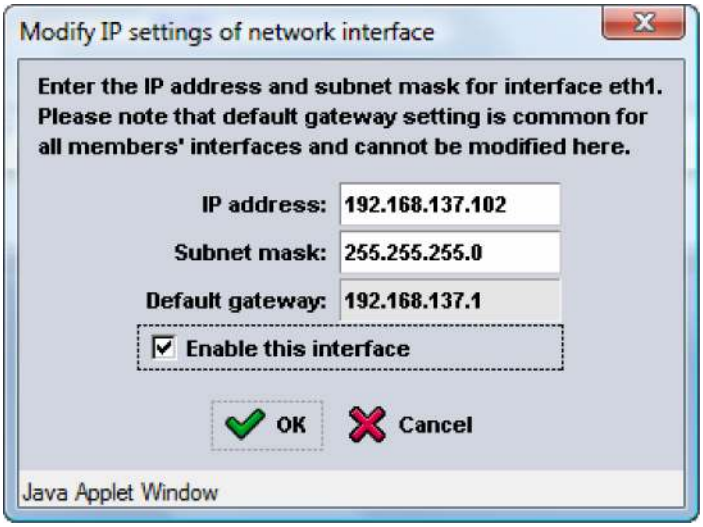

Configuring a Dedicated Management Port

On all Dell EqualLogic PS Series arrays with Version 4.0 or later firmware, the administrator can configure the highest-numbered Ethernet interface as a dedicated management interface that will not respond to iSCSI requests. This section lists the steps required to configure this port as a dedicated management port.

Note: The Dell EqualLogic PS4000 array has a dedicated management port already set aside. Simply configure this port with the appropriate IP address and subnet mask.

Note: The management network should be on a separate IP subnet from that being used for the SAN: typically, a subnet dedicated to systems management or the public LAN network.

CLI Commands

After connecting to an array in the SAN group and logging in, use the following steps to configure the last Ethernet interface in each array as a dedicated management port. (For more information, please read the *Dell EqualLogic PS Series Storage Arrays CLI Reference* available at the Dell and EqualLogic websites.)

1. Configure the eth3 Ethernet port on Dell EqualLogic PS6000XV array "m5" with IP address 192.168.138.104 with subnet mask 255.255.255.0 and then activate the port with the following commands:
   ```
   > member select m5 eth select 3 ipaddress 192.168.138.104 netmask 255.255.255.0
   ```

2. Enable the interface and set it as a management-only interface:
   ```
   > member select m5 eth select 3 up
   > member select m5 eth select 3 mgmt-only enable
   ```

3. Perform these steps for each array (standby controller ports automatically inherit the settings of the same port on the primary controller).

EqualLogic Group Manager GUI

Using the web-based Dell EqualLogic Group Manager, you can easily configure a network interface as a dedicated management interface. To configure the eth3 Ethernet port on Dell EqualLogic PS6000XV array m5 with IP address 192.168.137.104 with subnet mask 255.255.255.0 and then activate the port, enter the following commands:

1. Click Members > *member_name* > Network tab.
2. Select the last interface (typically eth2 or eth3) to be configured and choose "Modify IP settings" from the Activities menu (Figure 7).

Figure 7. Selecting the Ethernet Interface on the Dell EqualLogic Array

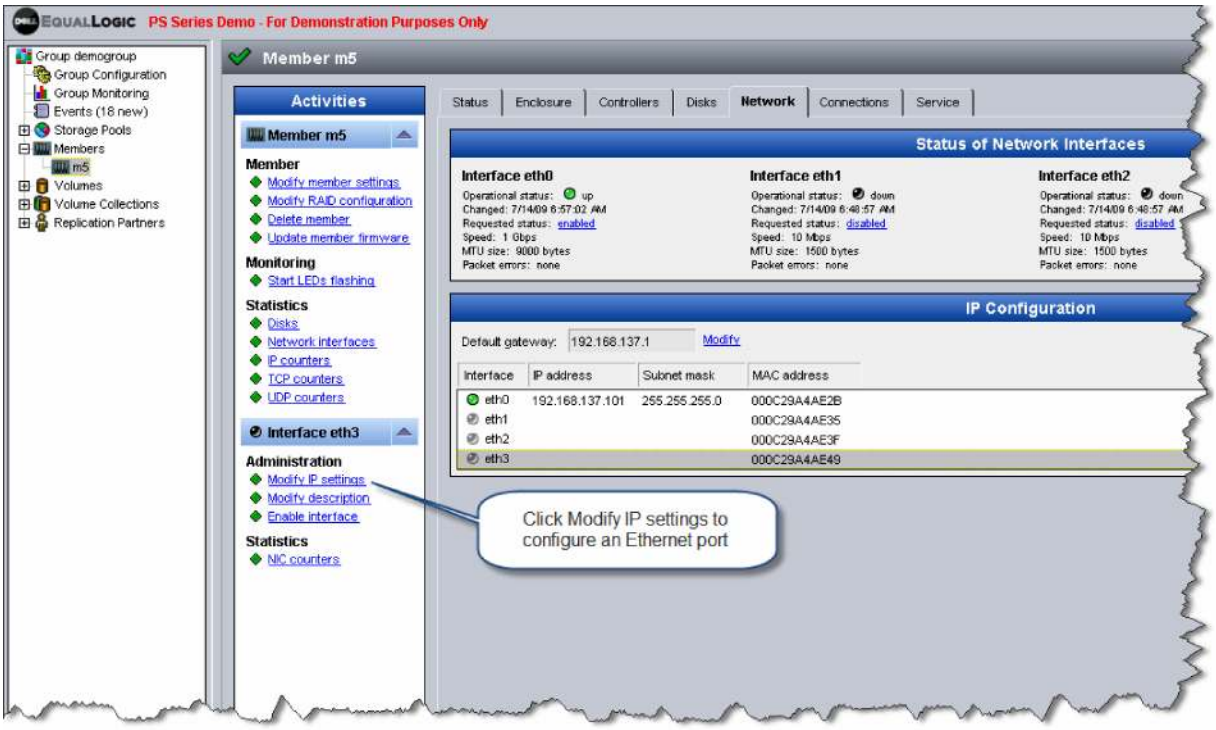

3. Enter the requested information for the IP address and subnet mask (Figure 8).
4. Check the “Enable this interface” check box
5. Check the “Restrict to management access” checkbox

Note: When the “Restrict to management access” check box is enabled, the “Default gateway” option will be disabled in the dialog box.

6. Click OK

Figure 8. Entering Ethernet Interface Information on EqualLogic Array

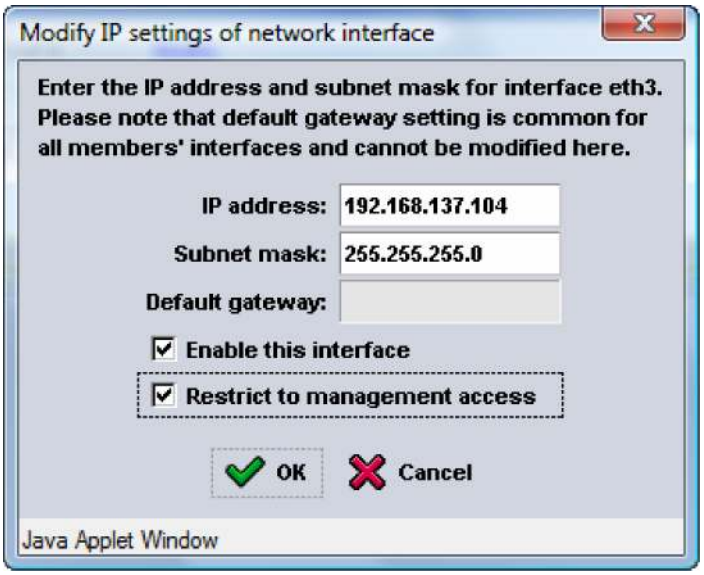

Conclusion

Application performance is one of the most important factors when considering a storage solution. Selecting the right storage solution, that is properly designed and configured, can have a huge effect on that performance. Choosing high-performance SAN components such as the Cisco Catalyst 4900 Series Switches for iSCSI storage solutions such as Dell EqualLogic PS Series arrays helps ensure that each enterprise application achieves its best performance. The Dell EqualLogic PS Series array, with its unique scale-out architecture, provides a storage solution with highly scalable capacity and performance. As application storage and performance needs increase, you can add arrays to the Dell EqualLogic SAN group to increase both capacity and performance.

As the SAN grows, organizations must have a SAN infrastructure with the features needed to support increased performance. Cisco Catalyst 4900 Series Switches provide line-speed throughput along with scalability and the capability to transition to 10 Gigabit Ethernet now. By using the configuration best practices provided in this document, you can implement a Dell EqualLogic PS Series and Cisco Catalyst 4900 Series solution that can provide the performance and scalability required for today's enterprise-class applications.

For More Information

Cisco Catalyst 4948 10 Gigabit Ethernet Switch and Cisco Catalyst 4900M Switch Software Configuration Information

- http://www.cisco.com/en/US/products/ps6021/products_installation_and_configuration_guides_list.html

Configuration Examples and Technical Notes for the Cisco Catalyst 4900 Series

- http://www.cisco.com/en/US/products/ps6021/prod_configuration_examples_list.html
- http://www.cisco.com/en/US/products/ps6021/products_installation_and_configuration_guides_list.html

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C11-563477-00 11/09