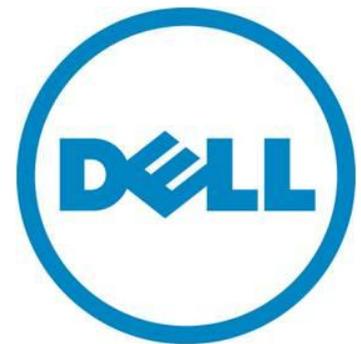


Dell EqualLogic™ iSCSI Optimization for Dell PowerConnect™ Switches

A Dell Technical Whitepaper

Reza Koohrangpour



Learn more at www.dell.com/networking

THIS TECHNICAL WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2011 Dell Inc. All rights reserved. Reproduction or translation of any part of this work beyond that permitted by U.S. copyright laws without the written permission of Dell Inc. is unlawful and strictly forbidden. For more information, contact Dell. Example: Dell, the DELL logo, and the DELL badge, PowerConnect, EqualLogic, and PowerVault are trademarks of Dell Inc. in the United States and other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

June 2011

Contents

Contents.....	1
Overview.....	2
Dell EqualLogic SAN Architecture.....	2
PowerConnect iSCSI Optimization.....	4
EqualLogic iSCSI Auto-Detection.....	5
PowerConnect Auto-Configuration.....	5
Case Study: Adding an iSCSI Array.....	7
Case Study: Moving an iSCSI Array.....	8
Case Study: iSCSI Array Controller Failover.....	9
PowerConnect iSCSI Monitoring.....	10
iSCSI Session Command.....	11
APPENDIX A.....	15
Enabling Jumbo Frames.....	15
Enabling Flow Control.....	15
Enabling Port-Fast.....	16
Disabling Unicast Storm Control.....	16
About Dell.....	16

Overview

Simple and efficient storage management is critical to addressing the increasingly complex IT challenges from trends such as server virtualization, and exponential growth of application data. Applications, whether they are running in physical or virtual environments, continually demand more from their storage systems—as performance, capacity, availability, backup, and disaster recovery needs intensify, so do expectations of the servers and storage that support them. As a result, IT professionals need a storage solution that integrates a full complement of high-end management capabilities, includes data protection, and operates intelligently. They need a reliable storage system that expands transparently and provides consistent data availability regardless of the server configuration, operating system, or application.

Internet Small Computer System Interface (iSCSI) is a TCP/IP-based protocol for establishing and managing connections between IP-based storage devices, hosts and clients. Dell™ EqualLogic™ PS Series iSCSI storage arrays offer a fundamental change in the way enterprises think about purchasing and managing storage. Built on patented peer storage architecture, these solutions offer enterprise-class performance and reliability, intelligent automation, and seamless virtualization of storage for simplified storage management. The PS Series combines an intelligent, automated management framework and a comprehensive set of enterprise data services with a fault-tolerant hardware architecture that supports many major operating systems. It delivers a modular and cost-effective solution that can be deployed in appropriate increments for small and medium businesses, while also being cost-effective for large enterprises requiring scalable capacity and high-end performance. With the Dell EqualLogic PS Series of storage arrays Dell provides a storage solution that delivers the benefits of consolidated networked storage in a self-managing, iSCSI storage area network (SAN) that is affordable and easy to use, regardless of scale. By eliminating complex tasks and enabling fast and flexible storage provisioning, these solutions dramatically reduce the costs of storage acquisition and ongoing operations.

The Dell EqualLogic SAN solution makes use of the SAN infrastructure to support inter-array communication and data load balancing on top of supporting data transfers between the hosts and the SAN. By combining these powerful storage systems with Dell's PowerConnect™ Ethernet switches with iSCSI Optimization, IT departments are able to reduce deployment time and management complexity using Dell's unique iSCSI Optimization features covered in this document.

Dell EqualLogic SAN Architecture

The Dell EqualLogic (EQL) PS Series has two controllers. One controller is active and the other is in passive failover mode, where it is not powered. There is no change on level of I/O performance if an EQL array configured with a single controller or dual redundant controllers. In order for controller failover to work, each Ethernet port¹ on the active controller that is connected to the SAN must have its corresponding port on the passive controller also connected to the same SAN network.

¹ Current EQL PS array controllers in the market have two, three, or four Ethernet ports.

In the event of controller failure, the standby controller gets activated and the following changes occur:

- The backup controller automatically enables each of the corresponding Ethernet ports that were enabled on the failed primary controller.
- The IP addresses that were assigned to each of the failed (previously active) controller Ethernet ports are reassigned to the corresponding ports on the passive controller. This has been shown in Figure 1 and Figure 2.

 A link failure on one or more of the Ethernet ports on the active controller does not constitute a controller failure. For this reason, it is important to ensure that ports from each controller are connected to at least two different switches.

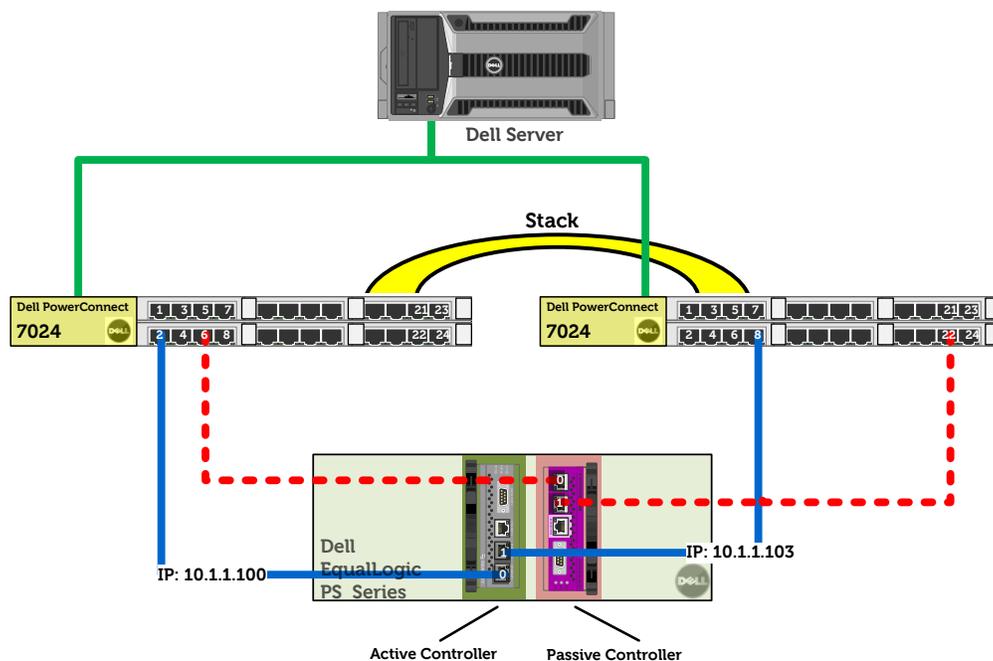


Figure 1. iSCSI Controller in Normal Operation

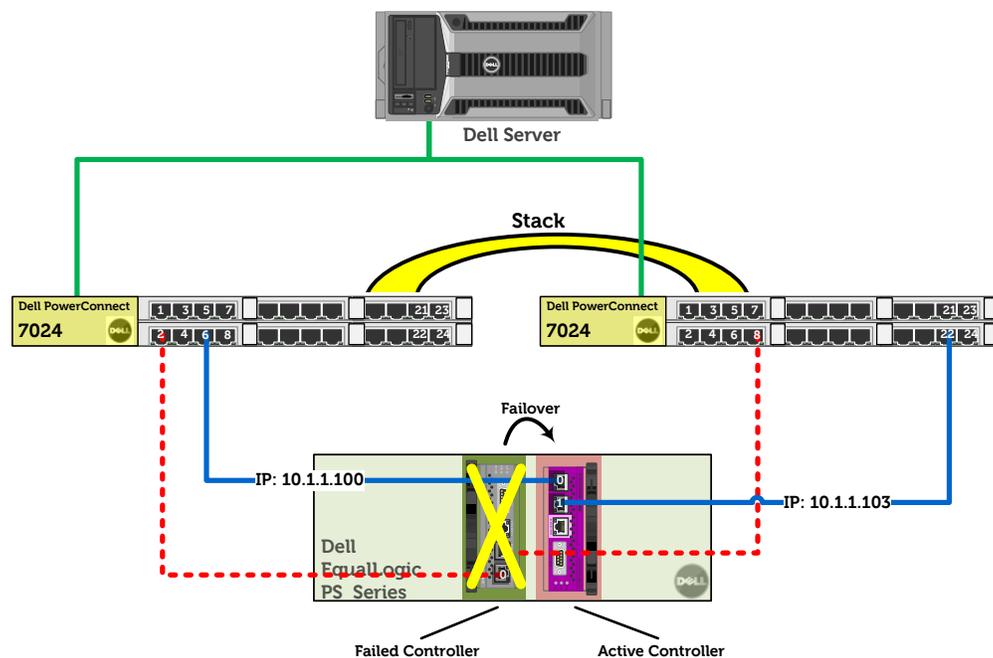


Figure 2. iSCSI Controller in the Failover Mode

PowerConnect iSCSI Optimization

The stacking PowerConnect 5500, 6200, and 7000 switch series and the non-stacking PowerConnect 8000 switch series are the world's first switches to automatically optimize themselves for iSCSI storage.

After enabling iSCSI feature, the above switches are able to automatically detect a directly connected EqualLogic array to any of their switch ports and to configure themselves for iSCSI environments, at both global and port levels. This will help customer save time configuring their iSCSI SAN solutions. The above switch series also provide information about all active iSCSI sessions to allow for easier management and optimization.

Some of the key iSCSI optimization features include:

- EqualLogic iSCSI Auto-Detection:
Ability of the switch to detect any active EqualLogic array attached directly to its ports.
- PowerConnect Auto-Configuration:
Automatic configuration of switch ports and global settings for iSCSI arrays, hosts, and switch links or stacking ports.
- PowerConnect iSCSI Monitoring sessions:
Ability of the switch to monitor, and track active iSCSI sessions on the switch, including port information and iSCSI session information

EqualLogic iSCSI Auto-Detection

When you connect a Dell EqualLogic iSCSI array to a Dell PowerConnect series switch and run “**iscsi enable**” command, the switch is capable to automatically detect a directly connected EqualLogic array to any of their switch ports.

The iSCSI Auto-Detection feature is, by default, **disabled** on Dell PowerConnect series switches. This feature will be **enabled** by default in later firmware releases. To enable the iSCSI on a switch connected to EQL Array, follow the below:

```
console>enable
console#config
console(config)#iscsi enable
console(config)#exit
console#copy running-config startup-config
console#exit
```

The switch recognizes the appearance and/or withdrawal of Dell Equal Logic array ports via Link Layer Discovery Protocol (LLDP) notifications². Whenever LLDP packets are received by the switch in one of its ports, the Chassis ID attribute is compared with EqualLogic OUI (00:09:8a). If the match is found, the switch recognizes the existing of an EQL array on that port. LLDP is globally enabled by default, on all Dell PowerConnect switches.

-  The “**iscsi enable**” command only works for the switch and it does not configure any server or host connected to the switch. This command also does not create any new VLAN for the iscsi traffic.

PowerConnect Auto-Configuration

Dell EqualLogic PS Series iSCSI arrays offer high levels of built-in intelligence and self-management. In combination with a Dell PowerConnect Switch, the two platforms enhance the self-management capabilities and reduce manual intervention and associated costs.

The main concept of auto-configuration is that the switch will automatically configure its interfaces to provide optimal support for iSCSI array.

² This feature only works when the EQL arrays have software version 5.0.2 or later installed and running.

After enabling iSCSI, the following procedures occur automatically:

- Jumbo Frames are enabled on all switch ports. Jumbo Frames enable transporting data in fewer frames to ensure less overhead, lower processing time, and fewer interrupts. By default, Jumbo Frames are **disabled** on all PowerConnect switches. To enable Jumbo Frames manually on a PowerConnect switch, please refer to [Appendix A](#).
 - ✎ The host/server has to be configured for Jumbo Frames manually. Jumbo Frames may not necessarily improve performance if packets are smaller than 1500 bytes.
- Spanning-Tree Port-Fast (STPF) is automatically enabled **ONLY** on the switch interfaces that are connected to Dell EQL array to allow fast connection to the EQL array. This is a port specific feature that happens only on the switch ports that receive active EQL LLDP messages. By removing the connection, the STPF is still enabled on the port unless it gets disabled manually. To configure a switch port, manually, to support Spanning-Tree port-fast, please refer to [Appendix A](#).
 - ✎ MTU 9216 is enabled on all ports and port channels when **"iscsi enable"** command is executed.
- Unicast Storm Control (USC) is automatically disabled **ONLY** on the switch interfaces that are connected to Dell EQL array. This is a port specific feature that happens only on the switch ports that receive an active EQL appearance LLDP messages. By removing the connection, USC is still disabled on the port unless it gets enabled manually. To manually disable Unicast Storm Control, please refer to [Appendix A](#).
 - ✎ Unicast Storm Control is globally disabled on all Dell PowerConnect switch ports. The **"iscsi enable"** command ensures that a user has not accidentally enabled this setting.
- Flow Control enables lower speed switches to communicate with higher speed switches by requesting that the higher speed switch refrains from sending packets. Transmissions are temporarily halted to prevent buffer overflows. Flow control is globally enabled on all Dell PowerConnect switches with firmware release 3.2.0.7 and higher. The **"iscsi enable"** command will automatically enable Flow Control on all switch ports regardless of the firmware version. To manually enable Flow Control on a switch, please refer to [Appendix A](#).
 - ✎ The host/server has to be configured for Flow Control manually.

Table 1 shows which of the above features are enabled or disabled by default on various PowerConnect Switch series.

<input checked="" type="checkbox"/> ON by Default <input type="checkbox"/> OFF by Default	PowerConnect Switches				
	5500 Series	6200 Series	7000 Series	8000 Series	
Firmware	4.x.x	3.x.x	4.x.x	3.x.x	4.x.x
Flow Control	✓	✓	✓	✓	✓
Spanning-Tree Fast Port	✗	✗	✗	✗	✗
Jumbo Frames	✗	✗	✗	✗	✗
Unicast Storm Control	N/A	✓	✓	✓	✓

Table 1. iSCSI Controller in the Failover Mode

When you connect a Dell EqualLogic iSCSI array to a Dell PowerConnect series switch and run **“iscsi enable”** command, the switch is capable to configure its ports automatically if one of the following cases occurs:

- Adding an iSCSI Array
- Moving an iSCSI Array
- iSCSI Array Controller Failover

Case Study: Adding an iSCSI Array

When you connect a Dell EqualLogic iSCSI array to a Dell PowerConnect series switches and run **“iscsi enable”** command, the switch is designed to automatically configure itself in the event an extra array is added to its interface.

Figure 3 shows an example of adding an additional EqualLogic iSCSI array to port 20 to an existing configuration. Since the **“iscsi enable”** command has already been executed at the original setup, all switch interfaces have been configured at both global and port levels before adding extra EqualLogic array.

The following steps will occur automatically:

- The switch initially receives only LLDP messages from its interface **2** where the first array is connected at the existing setup.
- Another iSCSI array gets added to the switch on its interface **20**.
- The switch starts receiving LLDP notifications from its interface **20**.
 - The switch automatically configures spanning-tree port-fast on its interface **20**.
 - The switch automatically disables Unicast Storm on its interface **20**.

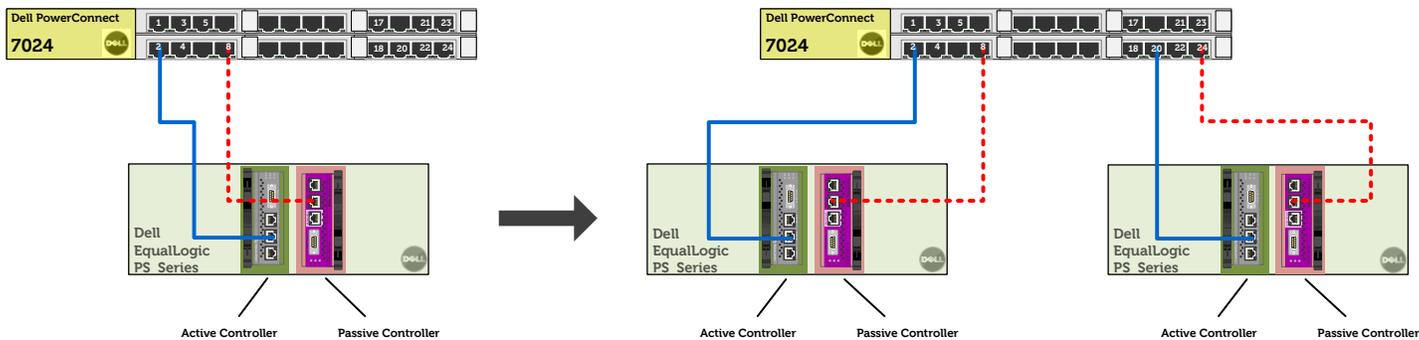


Figure 3. Graphical Representation of Adding an Array

Case Study: Moving an iSCSI Array

When you connect a Dell EqualLogic iSCSI array to a Dell PowerConnect series switches and run **"iscsi enable"** command, the switch is designed to automatically configure itself in case of moving an array from one interface to another.

Figure 4 shows an example of moving an EqualLogic iSCSI array from port 2 of a PowerConnect 7024 to port 5 on that switch. Since in the existing setup the **"iscsi enable"** command has already been executed, all switch interfaces have been configured at both global and port levels.

The following steps will occur automatically:

- The switch initially receives only LLDP messages from its interface **2** where the array is connected at the existing setup.
- Moving the array iSCSI from interface **2** to interface **5**.
- The switch does no longer receive any LLDP messages from its interface **2**.
 - Even though there is no connection on interface **2**, the spanning-tree port-fast is still enabled on this interface.
 - Even though there is no connection on interface **2**, the Unicast storm is still disabled on this interface.
- The switch starts receiving LLDP notifications from its interface **5**.
 - The switch automatically configures spanning-tree port-fast on its interface **5**.
 - The switch automatically disables Unicast Storm on its interface **5**.

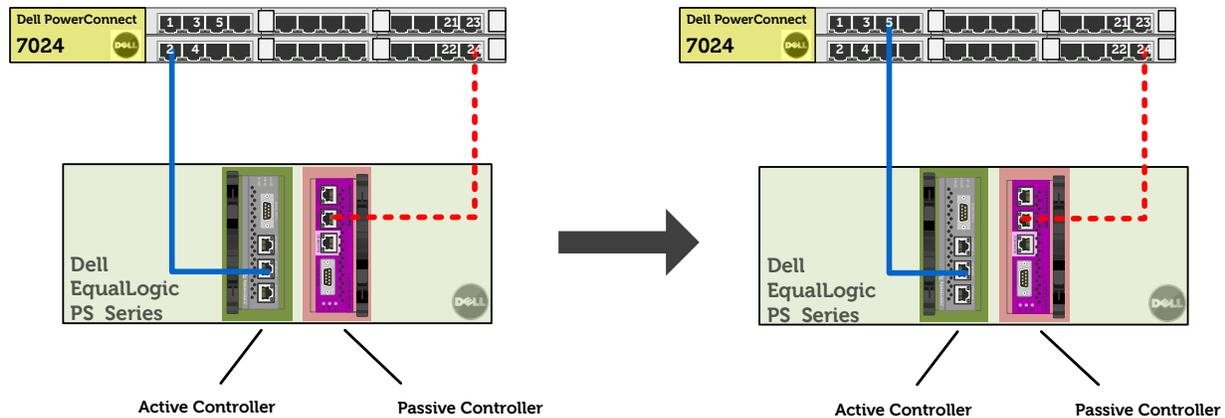


Figure 4. Graphical Representation of Moving an Array

Case Study: iSCSI Array Controller Failover

When you connect a Dell EqualLogic iSCSI array to a Dell PowerConnect series switches and run “**iscsi enable**” command, the switch is capable to automatically configure itself in case of any array controller failure.

Figure 5 shows the failover process. The following steps will occur automatically:

- The switch initially receives only LLDP messages from the iSCSI array through only array’s active controller **A**.
- The active controller **A** on the EQL array fails.
- The switch will no longer receive any LLDP messages from controller **A** interface.
- The passive controller **B** automatically takes over and becomes active.
 - The controller **B** automatically enables each of the corresponding Ethernet ports that were enabled on the failed primary controller.
 - The IP addresses that were assigned to each of the failed controller Ethernet ports are reassigned to the corresponding ports on the second controller.
- The switch receives LLDP notifications from controller **B** interface.
 - The switch automatically configures spanning-tree port-fast on its controller **B** interface.
 - The switch automatically disables Unicast Storm on its controller **B** interface.

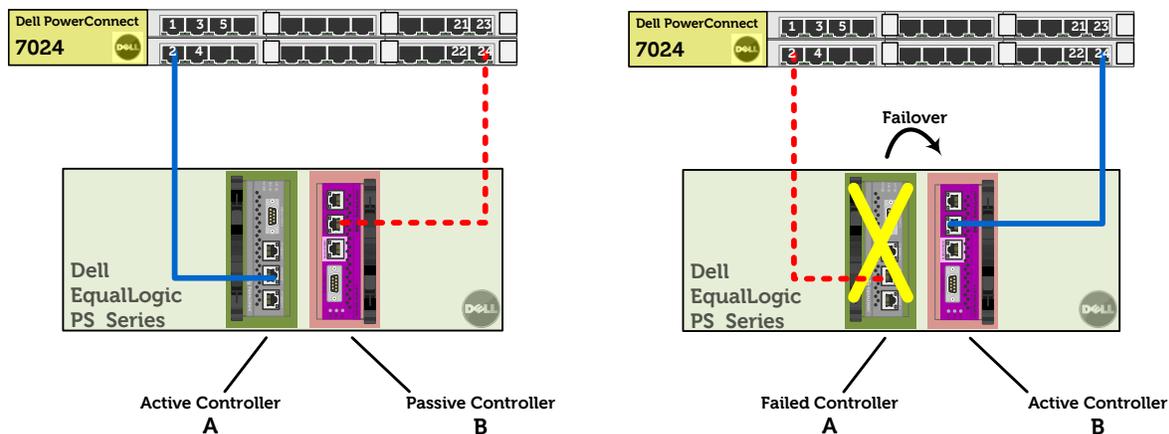


Figure 5. Graphical Representation of Controller Failover

PowerConnect iSCSI Monitoring

In an iSCSI SAN, devices are categorized as initiators or targets. An iSCSI initiator is generally configured on a host server (see Figure 6) and employs one or more Ethernet ports that are used for iSCSI communication. An iSCSI target is provided by a storage array and uses one or more Ethernet ports to present block storage devices to the initiator. The group of iSCSI TCP connections that link an initiator with a target is called an iSCSI session.

The switch uses well-known TCP ports 3260 or 860 to contact the iSCSI array. You can configure the switch to monitor traffic for additional port numbers or port number-target IP address combinations. You can also remove the well-known port numbers from monitoring and/or associate a target name with a configured target TCP port entry.

In addition to detection and monitoring of an iSCSI session, the Switch is able to track active iSCSI sessions and display key port and session information.

Follow the below steps in order to monitor an iSCSI session:

- Initiator logs in to Target via a GUI or a command line
- Enable iSCSI on the switch

```
console>enable
console#config
console(config)#iscsi enable
console(config)#exit
```

- Display iSCSI sessions

```
console>enable
console#show iscsi sessions [detailed]
```

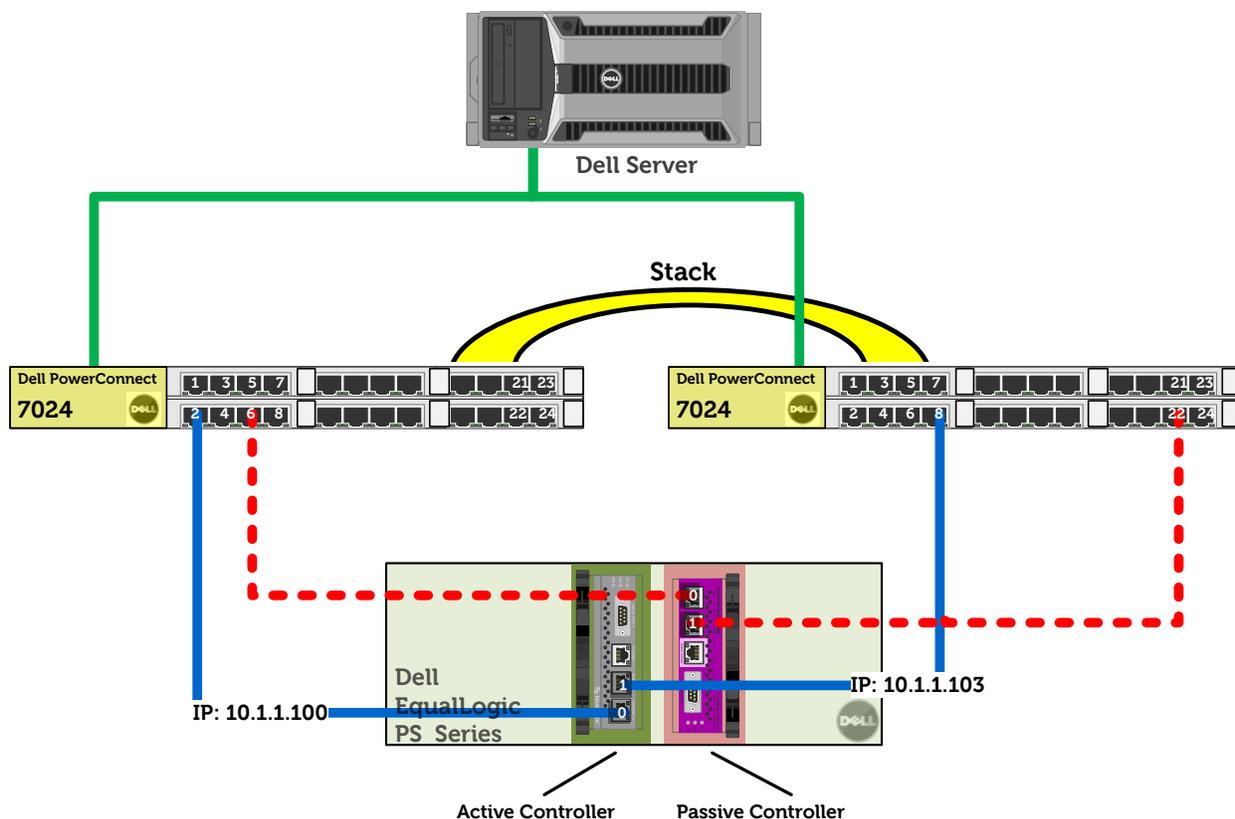


Figure 6. EqualLogic iSCSI attached to a PowerConnect Switch

 Monitoring the iSCSI session needs more than physical connection between initiator and target. The switch needs to identify IP packets to or from its ports for iSCSI traffic.

iSCSI Session Command

The **"show iscsi sessions"** command, in privileged switch mode, displays which initiator is connected to which target device. A user may have a single iSCSI Session for an iSCSI Target, but within that session there may be multiple connections. As mentioned before, devices that initiate iSCSI sessions generally use well-known TCP ports 3260 or 860 to contact targets. It is possible that more than one initiator be connected to any target device or one initiator to multiple target devices.

In order to display iSCSI sessions type the following commands into the switch.

```
console>enable
console#show iscsi sessions [detailed]
```

Adding "detailed" option to the command will result in showing more information on existing sessions.

The following is an example of what to expect to see when using the following command-line:

```
console#show iscsi sessions
```

Session 0:

 Target: iqn.2001-05.com.equallogic:0-8a0906-0e70c2002-10a0018426a48c94-iom010
 Initiator: iqn.1991-05.com.microsoft:win-x9l8v27yajg
 ISID: 400001370000

Session 1:

 Target: iqn.2001-05.com.equallogic:0-8a0906-0f60c2002-0360018428d48c94-iom011
 Initiator: iqn.1991-05.com.microsoft:win-x9l8v27yajg
 ISID: 400001370000

To view this info on a web interface (see Figure 7), please follow the steps below in Open Manage Switch Administrator interface:

- Select System → iSCSI → Sessions Table

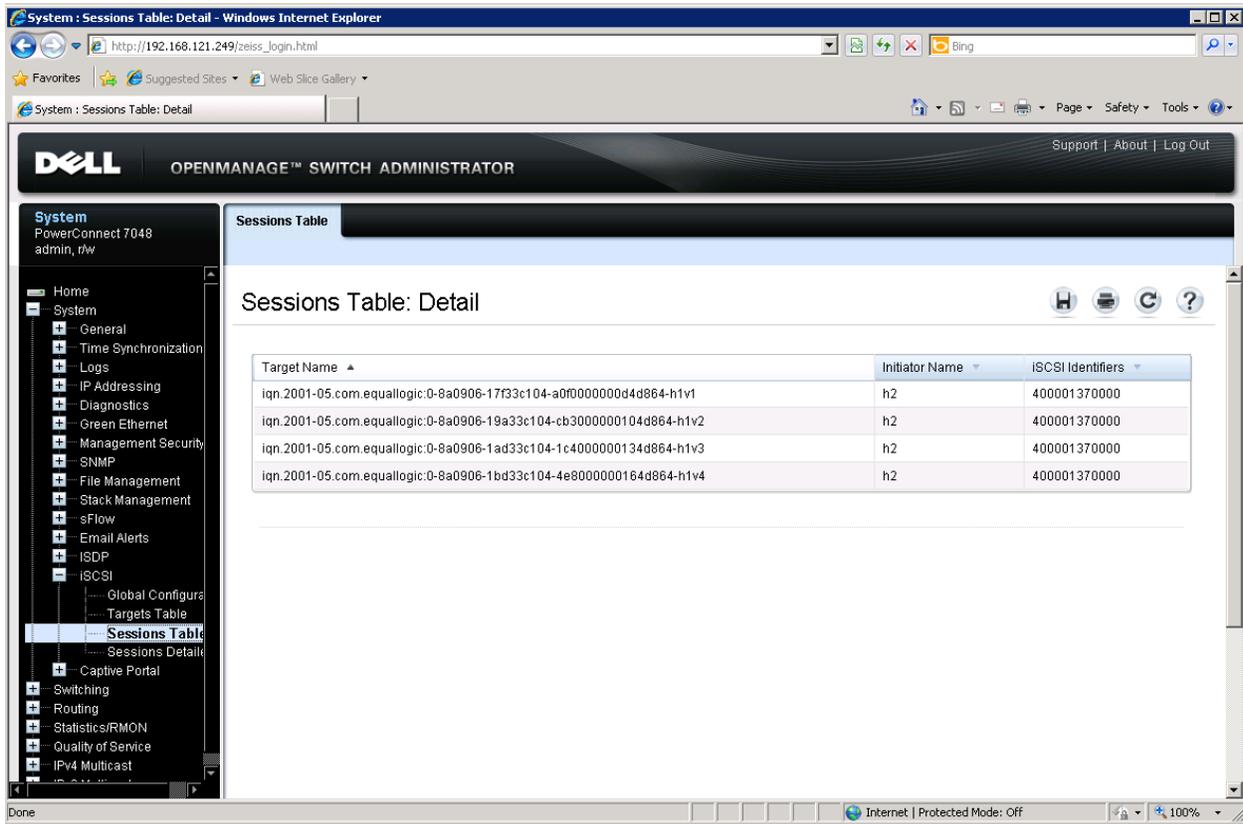


Figure 7. iSCSI Sessions Table via web Interface

The following is an example of what to expect to see when using the following command-line:

```
console#show iscsi sessions detailed
```

Session 0:

```
-----  
Target: iqn.2001-05.com.equallogic:0-8a0906-0e70c2002-10a0018426a48c94-iom010  
Initiator: iqn.1991-05.com.microsoft:win-x9l8v27yajg  
Up Time: 00:00:00:16 (DD:HH:MM:SS)  
Time for aging out: 44 secs  
ISID: 400001370000  
Initiator Initiator Target Target  
IP Address TCP Port IP Address TCP Port  
192.168.2.125 49272 192.168.2.20 3260
```

Session 1:

```
-----  
Target: iqn.2001-05.com.equallogic:0-8a0906-0f60c2002-0360018428d48c94-iom011  
Initiator: iqn.1991-05.com.microsoft:win-x9l8v27yajg  
Up Time: 00:00:00:13 (DD:HH:MM:SS)  
Time for aging out: 47 secs  
ISID: 400001370000  
Initiator Initiator Target Target  
IP Address TCP Port IP Address TCP Port  
192.168.2.125 49274 192.168.2.20 3260
```

To view this info on a web interface (see Figure 8), please follow the steps below in OpenManage Switch Administrator interface:

- Select System → iSCSI → Sessions Detailed

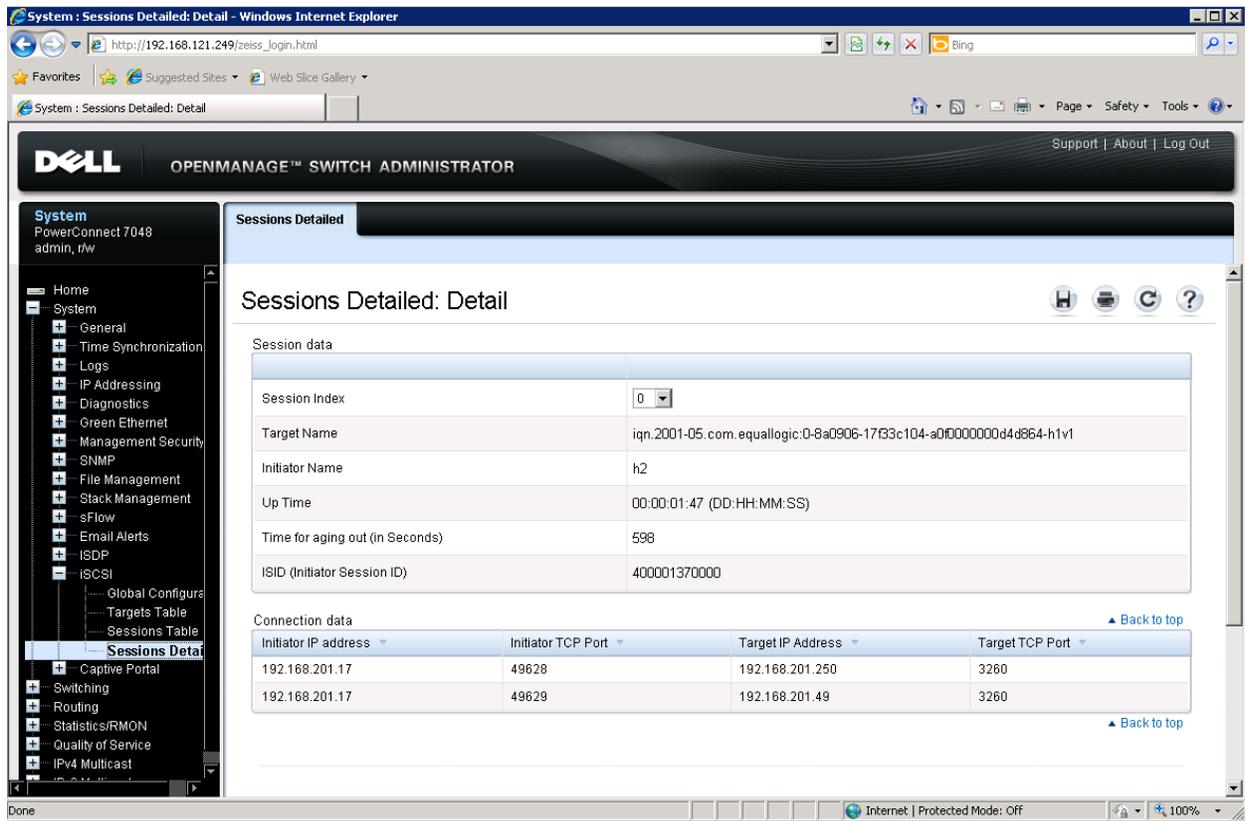


Figure 8. iSCSI Sessions Detailed via web Interface

APPENDIX A

Enabling 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. Jumbo Frames are not enabled by default. To enable jumbo frames on a PowerConnect switch, use the port jumbo-frame global configuration command. Jumbo Frames are enabled on all ports on a switch when enabled.

```
console#configure
console(config)#port jumbo-frame
console(config)#exit
console#copy running-config startup-config
console#exit
```

Enabling Flow Control

Flow control only works when the port is in full duplex mode, so be sure to enable full duplex on the port before enabling flow control. To enable flow control on all ports of a PowerConnect switch, enter the following commands:

```
console#configure
console(config)#interface range ethernet g(1-48)
console(config-if)#speed 1000
console(config-if)#duplex full
console(config-if)#flowcontrol on
console(config-if)#exit
console(config)#exit
console#copy running-config startup-config
console#exit
```

Enabling Port-Fast

To enable Port-Fast on a single port, the spanning-tree port-fast command must be used. Port-Fast should be enabled only on those ports being used to for inter-switch connections. The following steps are an example of using this command to enable port-fast on port 10 of a PowerConnect switch.

```
console#configure
console(config)#spanning-tree mode rstp
console(config)#interface ethernet g10
console(config-if)#spanning-tree portfast
console(config-if)#exit
console(config)#exit
console#copy running-config startup-config
console#exit
```

Disabling Unicast Storm Control

Three types of storms can occur on a network: broadcast, multicast, and Unicast³. 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.

```
console#configure
console(config)#interface range ethernet all
console(config-if)#no port storm-control unicast enable
console(config-if)#exit
console(config)#exit
console#copy running-config startup-config
console#exit
```

About Dell

Dell (NASDAQ: DELL) is a leading technology provider to commercial and public enterprises around the world.

³ The PowerConnect 5500 switch series does not set Unicast Storm Control.