



PS Series Best Practices

Configuring and Deploying the Dell EqualLogic™ Multipath I/O Device Specific Module (DSM) in a PS Series SAN

ABSTRACT

This Technical Report describes how to use the Dell EqualLogic™ Multipath Device Specific Module, (DSM) to enable and deploy multipath I/O, (MPIO) for highly available access to a PS Series SAN.



Copyright © 2008 Dell, Inc. All Rights Reserved.

Dell EqualLogic is a trademark of Dell, 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.

Dell will not be held liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change.

February 2008

Table of Contents

Revision Information 4

Technical Resources 5

Introduction..... 6

Configuring and Managing the Dell Equallogic MPIO DSM 7

 Managing the Networks Running MPIO 7

 Adding Access Control to Volumes 8

 Cluster Connections..... 10

Setting up the Host / PS Group Connection..... 10

 Connect to iSCSI Target..... 11

 Verify MPIO connections..... 12

 Set MPIO Load-Balance Policy 14

Summary 15

Appendix A – Defaults and known issues 16

For More Information 19

Technical Support and Customer Service 19

REVISION INFORMATION

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

Revision	Date	Document Info
1.0	February 2008	Initial Release

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

Vendor	Model	Software Revision
Microsoft®	Windows® Server™ 2003 R2	Service Pack 2
Microsoft®	Microsoft iSCSI Initiator	Version 2.04 or higher
Dell EqualLogic™	PS Series Firmware	Version 3.2 with L2 Patch or higher
Dell EqualLogic™	Host Integration Tools	V3.0.0 or higher*
QLogic	iSCSI Firmware for QLA405x	V3.0.1.24
QLogic	STOR Miniport Driver for QLA405x	V2.1.4.17
QLogic	BIOS for QLA405x	V1.13

* Host Integration Tools V3.0.3 or higher is required for cluster configurations.

TECHNICAL RESOURCES

The following table lists the documents and resources referred to in this Technical Report. All listed resources are available on the Customer Support site at: support.dell.com/EqualLogic by logging into the Dell EqualLogic Customer User Account.

Vendor	Document or Resource Title
Dell EqualLogic™	Technical Documents
Dell EqualLogic™	Product Documentation
Dell EqualLogic™	Deploying Windows Server 2003 in an iSCSI SAN Technical Report
Dell EqualLogic™	PS Series Best Practices Deploying Microsoft Multipath I/O in an iSCSI SAN
Dell EqualLogic™	Host Integration Tools User Guide – V3.0.0 and later
Dell EqualLogic™	Host Integration Tools Release Notes – V3.0.0 and later
Microsoft	Microsoft Storage Technologies - Multipath I/O
Microsoft	Microsoft iSCSI Software Initiator Users Guide – V2.x

INTRODUCTION

High availability of data is a requirement for many businesses today to ensure protection against system/device faults as well as increasing performance for workload demands. Achieving this level of availability is sometimes a challenging and daunting task for server and storage administrators.

Redundant hardware and RAID technologies help but when using a storage area network, SAN, the paths between the servers and the storage are vital for data transfer and availability.

Microsoft provides software to support multiple I/O paths to SAN storage called Multipath I/O. MPIO allows servers to have multiple connected paths to the same SAN volume for increased redundancy and performance. For more information on deploying Microsoft MPIO with PS Series SANs see the following technical report: PS Series Best Practices Deploying Microsoft Multipath I/O in an iSCSI SAN. Using Microsoft MPIO, users can configure each server network port to connect to the SAN one at a time.

With many servers and multiple NICs per server, this process can be a lengthy and tedious job. To simplify this process, Dell EqualLogic has developed a Device Specific Module (DSM) to wrap around the Microsoft's MPIO capabilities. The Dell EqualLogic MPIO DSM takes the confusion out of connecting all available I/O paths from the server to the storage and automatically connects these paths based on designated subnets.

The Multipath I/O DSM Delivers:

- Automatic connection management
- Automatic failure detection and failover
- Automatic load balancing across paths
- Support for multiple connections to a single iSCSI target
- Increased I/O bandwidth
- Increased volume bandwidth
- Reduced network latency
- Easy installation and management
- Support for a variety of iSCSI initiators

The MPIO DSM is easy to install as part of the Host Integration Tool Kit V3.0.0 and above. The Host Integration Tool Kit is included with the purchase of a PS Series array or downloadable from support.dell.com/EqualLogic at no extra cost.

Dell EqualLogic MPIO Solution

The Dell EqualLogic DSM consists of two components, a kernel mode component that works in conjunction with the Microsoft MPIO driver to route I/O to the desired path, and a user mode service that manages connections. The connection manager automatically makes the appropriate connections to the target volume when a single connected is made. Together they allow administrators to easily install and configure multipath I/O for iSCSI networks.

You can manage specific subnets to use MPIO on the host by modifying the lists through the Remote Setup Wizard. By including and excluding specific subnets for MPIO you can control what networks you want to use for iSCSI traffic more easily. The number of connections per target depends on the number of arrays in the group and the number of NIC's being used for MPIO. See Appendix A for default PS group connection values.

The Dell EqualLogic MPIO DSM requires the Microsoft iSCSI Initiator to make the appropriate connections to the target volumes. You may use offload cards such as HBAs but the connections must be made through the Microsoft initiator for the DSM to work correctly.

When using the Microsoft iSCSI service, additional sessions are established once the first login to the target is made. By viewing the details on a connected target you can see how many sessions have been established. Allow some delay for additional sessions to be established after the initial login.




Figure 1: Basic MPIO Configuration

CONFIGURING AND MANAGING THE DELL EQUALLOGIC MPIO DSM

As mentioned the Dell EqualLogic MPIO DSM is included with the Host Integration Tool Kit V3.0.0 and above. For installation guidelines please refer to the *Host Integration Tools User Guide*.

Once the MPIO DSM has been installed, you can:

- Turn on and manage the networks to run MPIO
- Make connections to volumes
- Change the policies to load balancing I/O

Managing the Networks Running MPIO

You can turn on and include or exclude specific IP subnets to use for MPIO. To manage your MPIO subnets you must start the Remote Setup Wizard installed with the Host Integration Tools Kit.

On the Remote Setup Wizard's first screen you will see the *Configure MPIO settings for this computer* radial button at the bottom of the screen. Check that button and hit **Next**. Any configured IP subnets on the host will automatically appear on the left hand window in the *Subnets included for MPIO* box.

You can exclude a subnet by simply moving that subnet to the *Subnets excluded from MPIO* box as in figure 2. By including and excluding specific networks you can control where the MPIO traffic is handled. Be sure to propagate these changes throughout the hosts in your network.

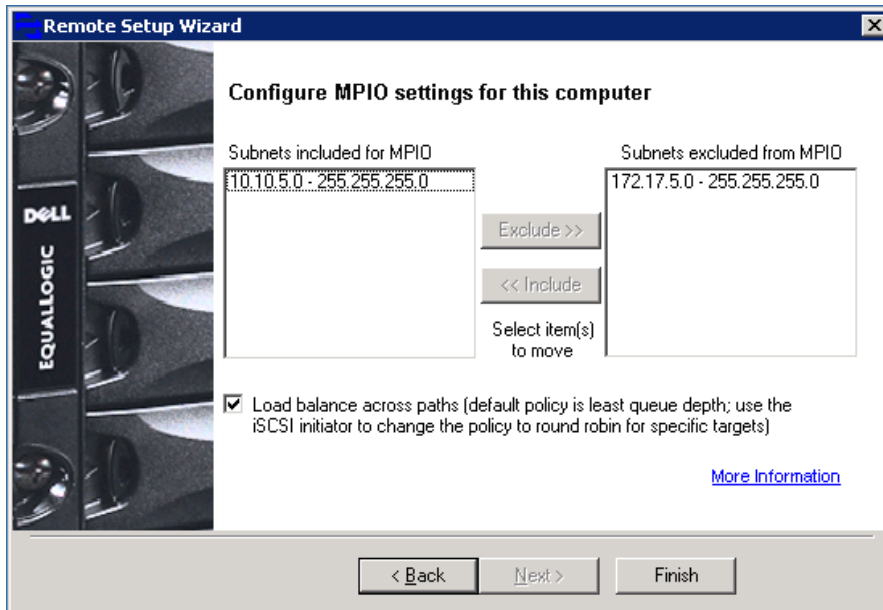


Figure 2: Remote Setup Wizard – MPIO Configuration

Additionally check the “**Load balance across paths...**” checkbox at the bottom of the screen to enable load balancing using a default policy of Least Queue Depth. If the checkbox is not checked the Failover Only load balancing policy will be used.

Adding Access Control to Volumes

You add access control to your volume through the PS Group Manager, either at the time of volume creation or afterwards. ACL records can be created using initiator name (Figures 3 and 4), IP addresses (Figure 5), or CHAP credentials. The examples show how to set up an ACL using the initiator name as well as the host or port IP address.

Note: It is important to note that the use of CHAP ACLs for QLogic HBAs is not permitted when using the Dell EqualLogic MPIO DSM.

Note: As a best practice it is not recommended to prevent unauthorized hosts from discovering targets.

A single ACL record for each host is typically sufficient for enabling volume connections from the MPIO DSM. However, if using IP addresses in the ACL, it is recommended to give a range of addresses or add an ACL record for each interface the host will use to connect to the volume, as shown in figure 5.

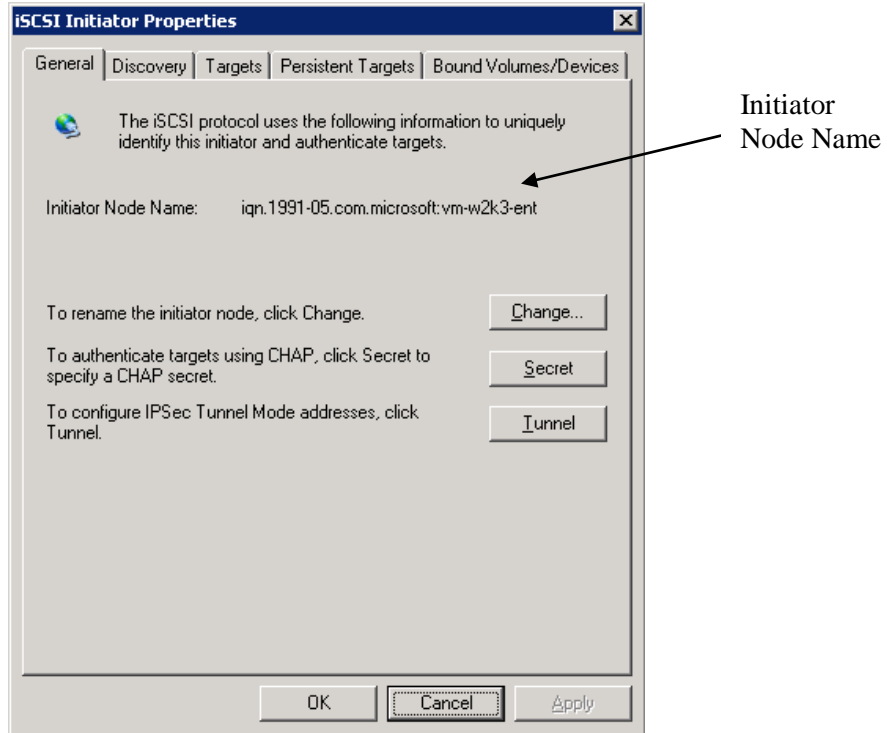


Figure 3: Microsoft Initiator Node Name

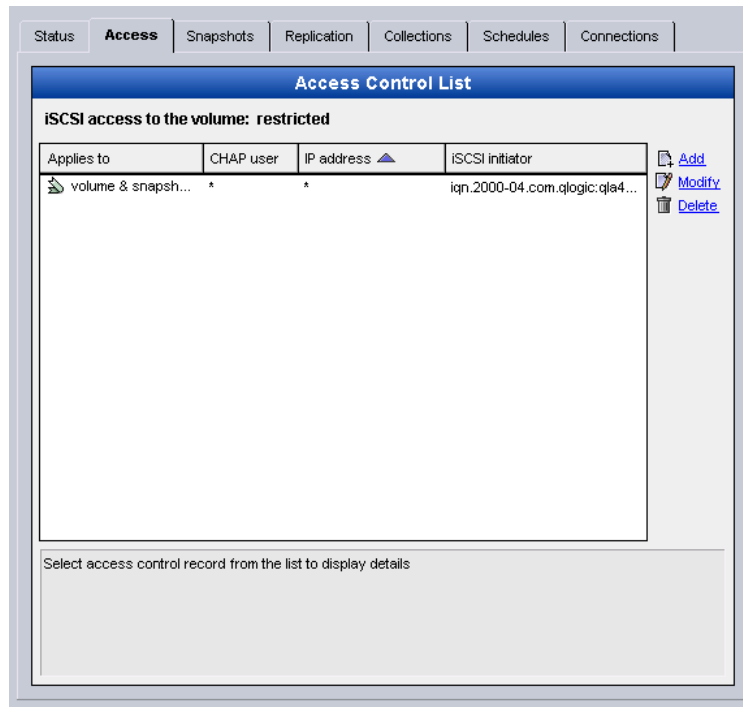


Figure 4: Adding the Initiator Node Name to the Volume ACL

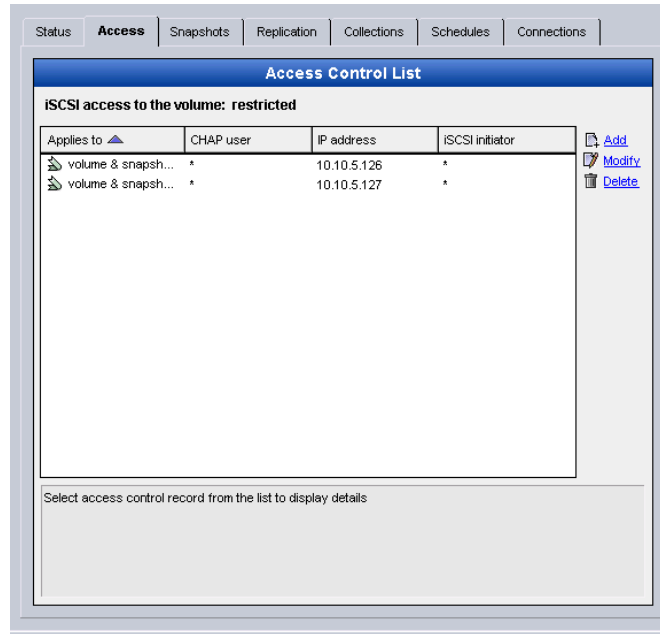


Figure 5: Creating Multiple ACLs for Persistent Connections

The MPIO DSM connection manager will not manipulate persistent connections. Setting up ACLs via each interface in the host will enable a persistent connection to the volume for each host interface. In the event of a link failure and reboot of the host, another persistent interface link will automatically connect to the volume at startup. With only one ACL, if that interface link is down and the host is restarted it will not be able to connect to the volume at startup.

Cluster Connections

If you are connecting volumes for cluster configurations there are some additional steps that need to be taken to assure proper connection. You will want to add an ACL record for each cluster host adapter connecting to the volume. Like figure 5 you will have multiple ACL records depending on the number of host in the cluster for volume access.

Important Note: Cluster configurations require using the MPIO DSM installed with the Host Integration Tools V3.0.3 and above.

Additionally you will have to make some registry changes to assure each host in the cluster is using persistent reservations. [Appendix A](#) has information on how to make these changes.

SETTING UP THE HOST / PS GROUP CONNECTION

This section will discuss how to make the proper connection to a PS Series SAN including adding the Target Portal and connecting to volumes from the host.

Using the Microsoft iSCSI Initiator Service, add a Target Portal to the PS Series group IP address to allow the host to discover available targets. This process is done in the **Discovery** area of the initiator properties window, as shown in figure 6.

Note: If you must use CHAP as a discovery method you'll have to click **Advanced** and check the CHAP logon information box and add the Target secret to complete the connection.

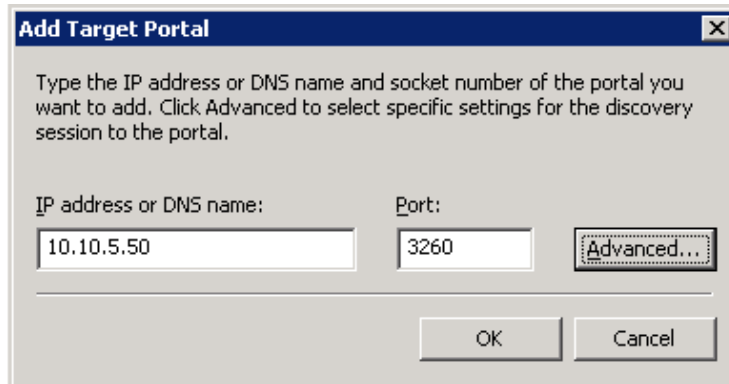


Figure 6: Adding the PS Series Group IP to the Target Portal

The Discovery window should now show the IP address of the PS SAN as in figure 7.

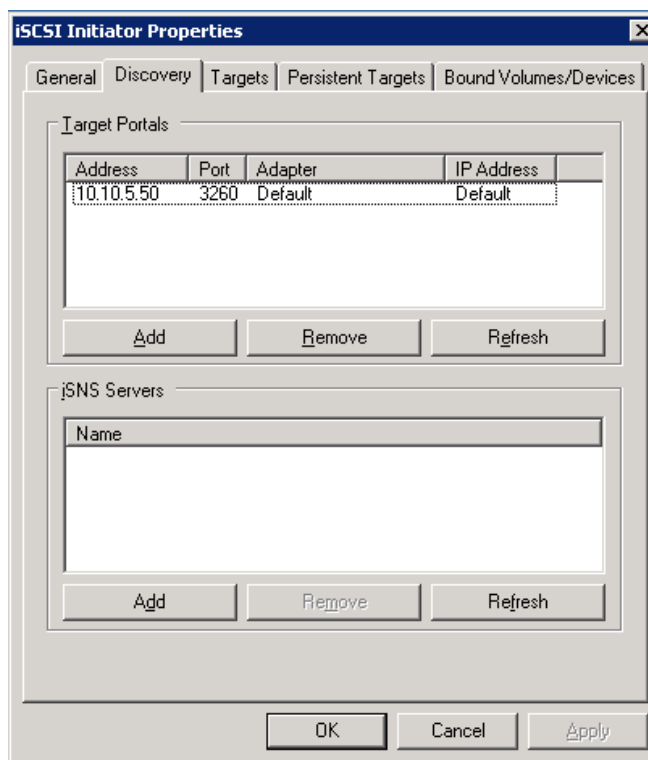


Figure 7: Target Portal – Discovery Tab

Connect to iSCSI Target

To enable MPIO on a target, you will need to select the "Enable multi-path" checkbox in the Microsoft initiator when logging into your target. Optionally if you plan to use persistent connections to the target you can select the "Automatically restore this connection when the system boots" checkbox as in figure 8.




Figure 8: Target Log On

Note that you only have to login to your target once, the MPIO DSM will handle creating all other appropriate connections. You may need to hit the “Advanced” button to specify your initial connection, if it's not immediately obvious to Windows.

For persistent connections to volumes it is important to repeat this process and log into the target portal (PS Group) via the IP address of each adapter, so that even if there is a failure in an adapter or network path at boot time, the host will attempt to reconnect through other available adapters.

Verify MPIO connections

Once you have made your initial connection, the MPIO DSM connection manager will make all of the connections it deems appropriate, and replace any original connections if necessary. If you have multiple arrays and multiple NICs, you may see a number of connections. For instance, two NICs and three arrays would result in six connections per volume. After making your initial connection, allow a minute or two for the DSM to communicate with the array group and get sufficient information to create the multiple connections. You can change the number of connections the DSM connection manager will manage by modifying some registry keys. For a list of registry keys and how they are modified refer to [Appendix A](#).

To verify that you have multiple connections to the array for your volume, there are two places to look. First, look at the MS Initiator control panel. Select your connected volume, and click on the Details button. Note that you now have more than one connection as figure 9.

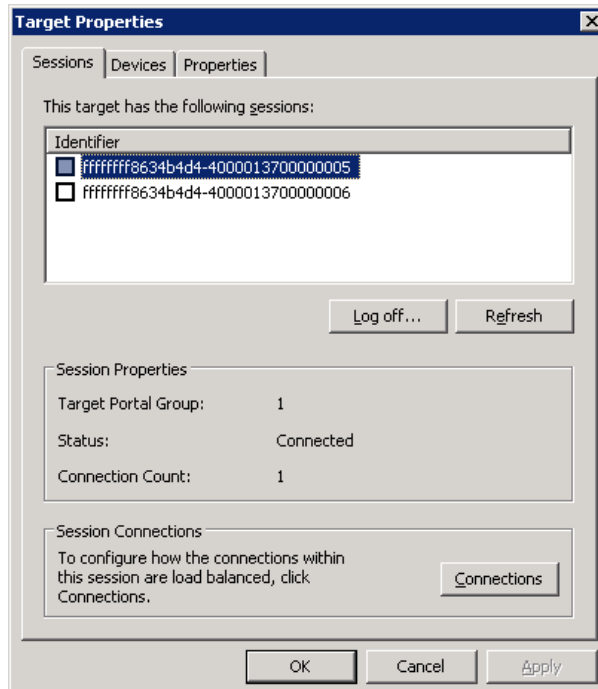


Figure 9: iSCSI Target Details

You can further identify the server interface for each path by selecting it and clicking on the Connections button, and noting the Source Portal column. In figure 10, the specific path is going out the 172.23.250.113 NIC, to the array whose group IP address is 172.23.10.60. Note that the "Load balance policy" shown in the Connections dialog is not the MPIO load-balance policy, and should be ignored.

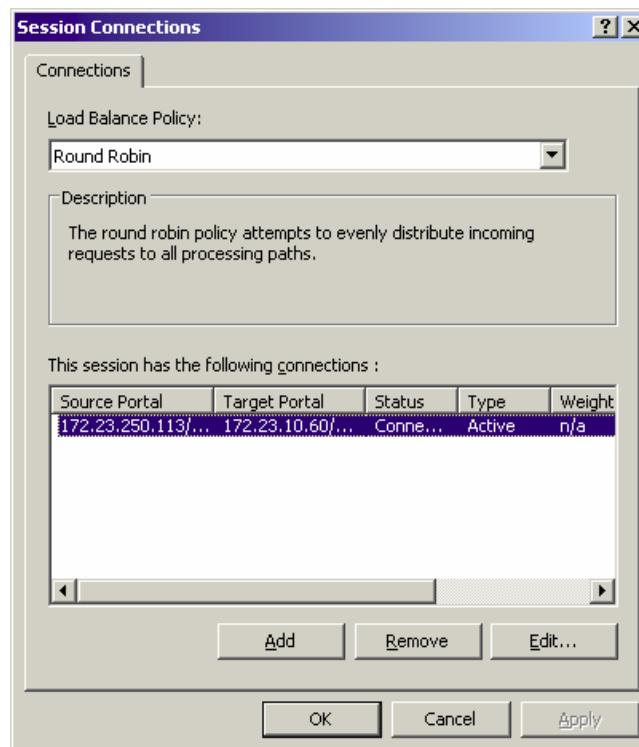


Figure 10: Connection Detail Window

Set MPIO Load-Balance Policy

Once you have multiple connections, it is appropriate to configure MPIO for how you want them used. When the MPIO DSM is enabled, the default load balancing policy is Least Queue Depth. When the MPIO DSM is disabled, the default load balancing policy is Fail Over Only. In most cases it is best practice to accept the default load balancing policy because it is tuned to run with PS Series SANs.

The three supported options are:

- **Fail Over Only:** one connection is used for all data until it times out or fails, and then one of the other available paths is chosen. If this load balancing scheme is selected, the MPIO DSM will not manage connections to the target.
- **Round Robin:** Data is sent out to each available connection in a rotating sequence, effectively utilizing all available ports fully.
- **Least Queue Depth:** Data is sent out to each available connection, with preference given to the connection which is least busy at the time the I/O is requested.

Additional unsupported options include round robin with subset, weighted paths, and least blocks. Although these options are still present and seem available, an error message will show up if you choose one of these options.

To configure the MPIO load balance policy, you must go to the MS Initiator control panel, select a target, click on the Details button, and select the Devices tab as in figure 11.

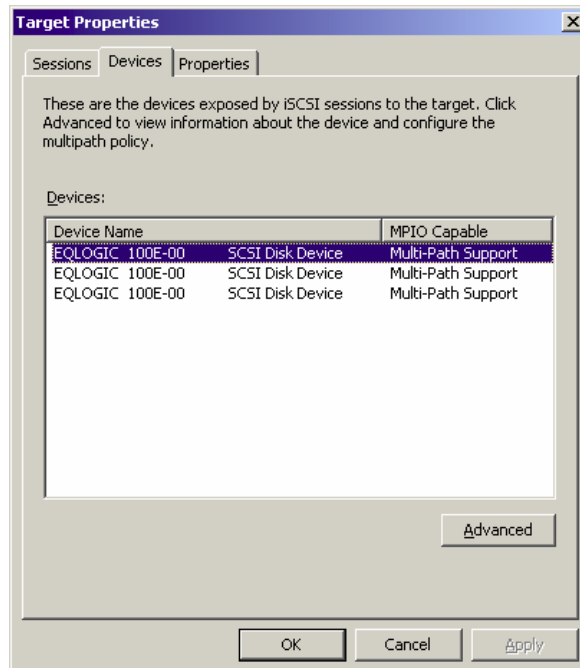


Figure 11: Target Devices View

Verify that the "MPIO Capable" column says "Multi-Path Support." If it says anything other than that (e.g. "Disk Drive"), then it means that the Microsoft MPIO DSM has not been properly installed, or there is a problem recognizing an HBA as MPIO compliant.

Then hit the “Advanced” button, and look at the MPIO tab as in figure 12. You can change the MPIO policy accordingly with the drop down tab.

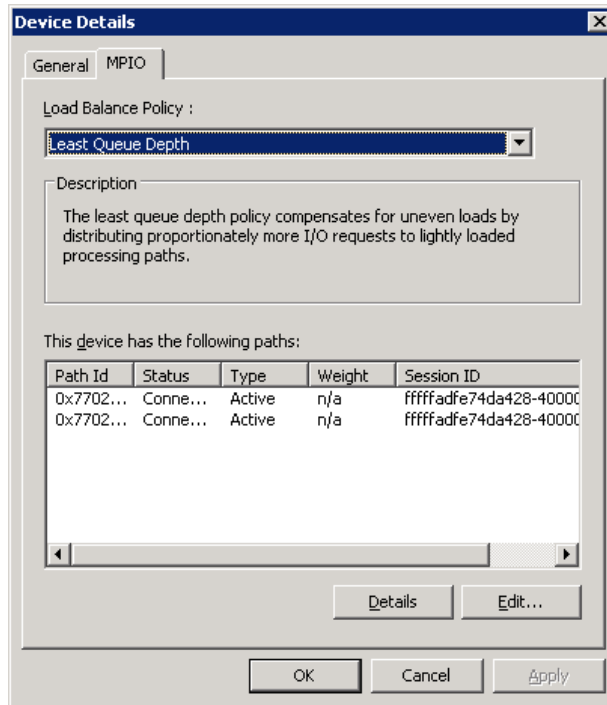


Figure 12: Target Device View - Advanced

SUMMARY

The Dell EqualLogic MPIO DSM allows for easy and time saving MPIO configurations by automatically making the MPIO connections for you. Additionally, you can designate IP subnets to use for MPIO by making changes in the Remote Setup Wizard.

The Host Integration Tools kit and documentation will be available on the support website (support.dell.com/EqualLogic) and on a CD-ROM that is shipped with the PS Series array.

APPENDIX A – DEFAULTS AND KNOWN ISSUES

MPIO with Microsoft Clusters

- Multipath I/O is only supported on clusters that are configured to use Persistent Reservations. Instructions on how to use Persistent Reservations are in the *Microsoft iSCSI Software Initiator 2.x Users Guide*. You must add the following registry key setting changes under the following:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\msiscdsm\PersistentReservation

- Microsoft suggests adding values to the key in the following manner:
 - **UsePersistentReservation** REG_DWORD, Set this hex value to 1 to enable persistent reservations.
 - **PersistentReservationkey** REG_BINARY, This is a 8-byte hex value that is unique to the cluster. All nodes in the cluster must have this value, see example:
 - Node 1: aabbccccbaa0001
 - Node 2: aabbccccbaa0002
 - Node 3: aabbccccbaa0003
 - Note the first 6 bytes must be identical and the last two bytes must be unique for each node in the cluster.

MPIO DSM Connection Defaults

- For each iSCSI target, the Dell EqualLogic Multipath I/O DSM connection manager will establish a maximum of:
 - Six connections to each group. You can configure this setting by modifying the following parameter in the registry:

HKEY_LOCAL_MACHINE\SOFTWARE\EqualLogic\EHCM\MaxDevicesPerMpioSession

- Two connections to each member. You can configure this setting by modifying the following parameter in the registry:

HKEY_LOCAL_MACHINE\SOFTWARE\EqualLogic\EHCM\MaxConnectionsPerMember

- One connection to each member from each adapter (NIC or HBA).

Known Issues

- When installing Microsoft MPIO support on a boot from SAN system, it is necessary to reboot several (3) times before MPIO is completely installed, as opposed to 1 reboot on a non boot from SAN system. The user will only be prompted to reboot once, though if you examine the system event log, events are logged from Plug-n-Play Manager requesting a reboot of the system. Once Microsoft MPIO support has been installed on a boot from SAN system, it should not be uninstalled, as doing so leaves the system in an instable state and may no longer boot successfully.
- Once the Microsoft MPIO support has been installed on a boot from SAN system, it is recommended not to uninstall it, as doing so may leave your system in a state where it is unable to boot.

- Before logging in to a recovery volume, be sure no initiator is connected to the original volume (A recovery volume is the result of promoting a replica set.). Otherwise, the Microsoft MPIO driver will treat the sessions as paths to the same target.
- The Microsoft MPIO driver handles a network path failure by failing over to another existing path to the same target. Because establishing new connections is a complex operation, the driver performs this operation as a background process, instead of during the failover itself. Therefore, if all the adapters with existing connections to a target fail simultaneously, you will lose access to the target even if there are other unused adapters available on the computer. The host connection management software reduces the chance of this happening by distributing connections among as many adapters as possible.
- The Multipath I/O DSM does not manage persistent connections to iSCSI targets. For maximum redundancy on a computer with multiple adapters, it is recommended that you log in persistently to the target from each adapter.
- Due to a known initiator problem, you may experience intermittent login failures when connecting to volumes or snapshots with a QLogic iSCSI HBA. The threshold at which the problem occurs depends on the group size:
 - For a single-member group, the computer must be connected to 18 targets.
 - For a two-member group, the computer must be connected to 9 targets.
 - For a group with three or more members, the computer must be connected to 6 targets.
- The Multipath I/O DSM connection manager uses temporary CHAP credentials to create additional iSCSI sessions to each target. This use of temporary CHAP credentials is not supported on a cluster. In addition, QLogic iSCSI HBAs do not support the use of CHAP credentials for these sessions on a non-clustered server. Therefore, in order to use MPIO with QLogic iSCSI HBAs, you must do one of the following for each target.
 - Create a single access control record with the iSCSI initiator name of the host computer.
 - Create a single access control record with an IP address range that includes all of the host computer's interfaces
 - Create multiple access control records, one for each of the host computer's interfaces, where each record contains a specific IP address.

Occasionally when swapping QLogic HBAs from one host to another the initiator name may change. Running the QLogic SANSurfer application will show the actual initiator name of the HBA port.

Excluding an IP Address from a Subnet Configured for Multipath

By default, when you configure the Dell EqualLogic Multipath I/O DSM to include a subnet, all the IP addresses on the subnet are configured for multipath. However, there might be cases where you want to exclude individual IP addresses on a subnet.

For example, your computer could have three network interface cards (NICs) that are all on the same subnet as the PS Series group, but you only want to use two of the NICs for iSCSI connections to the group volumes.

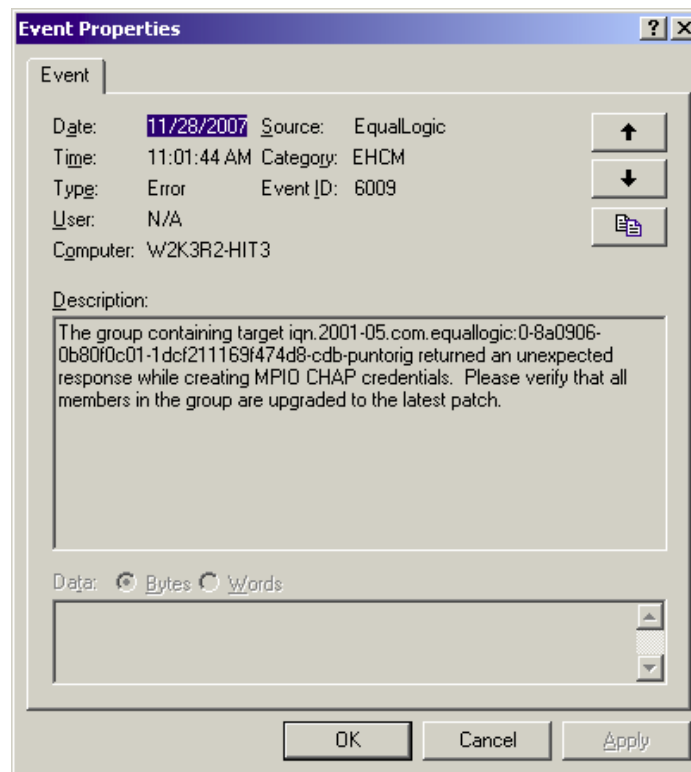
Note: The computer will automatically exclude any addresses that cannot reach the PS Series group. To exclude an individual IP address on an included subnet, you must manually edit the registry variables. Follow these steps:

1. Open the Registry Editor: `Start > Run > regedit`

2. Navigate to the HKEY_LOCAL_MACHINE\SOFTWARE\EqualLogic\EHCM\HostPortal directory.
3. Right-click the directory, select New > Key, and then enter the IP address you want to exclude.
4. Right-click the new subkey, select New > String Value, and enter Mask.
5. Double-click the new entry to open the Edit String dialog box and enter 255.255.255.255 in the Value data field.
6. Close the Registry Editor.

Troubleshooting the Dell EqualLogic MPIO DSM

The most common issue with the Dell EqualLogic MPIO DSM is that you are running the wrong version of firmware. HIT V3.0.0 and the MPIO DSM require that you be running a minimum PS Series firmware version 3.2.0 L2 on your PS Arrays. If you are not, you will see this message in your Windows System Event Logs:



"EHCM" is the Host Connection Manager, and this is telling you to version check the firmware on your PS Arrays. Update to 3.2.0 L2, or to any release after that.

Also, make sure your session had the "Enable Multi-path" option checked. If you are unsure, logout from the target, remove any persistent target information in the initiator, and log back in specifying that option.

Another common issue is that you have not actually installed Microsoft MPIO support when you installed your initiator. This can be seen if the Devices tab for your sessions shows "Disk Drive" instead of "Multi-Path Support," or if you have three separate disk devices in your Disk Manager instead of one. Re-run the MS Initiator installer, and tell it to include MPIO. This will require a restart of your server to activate. Be very careful not to use more than one instance of the same disk at once, if you have too many of them in Disk Manager. This is the same as mounting a single disk from two servers, and will destroy your data.

FOR MORE INFORMATION

For detailed information about PS Series arrays, groups, and volumes see the following documentation:

- *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.

The *QuickStart* and *Hardware Maintenance* manuals are printed and shipped with the PS Series array. They are also located on the documentation CD-ROM that is shipped with the array, along with the *Group Administration* and *CLI Reference* manuals and the Group Manager online help.

The Host Integration Tools kit and documentation will be available on the support website (support.dell.com/EqualLogic) and on a CD-ROM that is shipped with the PS Series array.

TECHNICAL SUPPORT AND CUSTOMER SERVICE

Dell's support service is available to answer your questions about PS Series arrays. If you have an Express Service Code, have it ready when you call. The code helps Dell's automated-support telephone system direct your call more efficiently.

Contacting Dell

Dell provides several online and telephone-based support and service options. Availability varies by country and product, and some services may not be available in your area.

For customers in the United States, call 800-945-3355.

Note: If you do not have an Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog.

To contact Dell for sales, technical support, or customer service issues:

1. Visit support.dell.com.
2. Verify your country or region in the Choose A Country/Region drop-down menu at the bottom of the window.
3. Click **Contact Us** on the left side of the window.
4. Select the appropriate service or support link based on your need.
5. Choose the method of contacting Dell that is convenient for you.

Online Services

You can learn about Dell products and services on the following websites:

- www.dell.com/
- www.dell.com/ap/ (Asian/Pacific countries only)

- www.dell.com/jp (Japan only)
- www.euro.dell.com (Europe only)
- www.dell.com/la (Latin American countries)
- www.dell.ca (Canada only)

You can access Dell Support through the following websites:

- support.dell.com
- support.dell.com/EqualLogic
- support.jp.dell.com (Japan only)
- support.euro.dell.com (Europe only)