

Configuring and Deploying the Dell PS Series Multi-Path I/O Device Specific Module with Microsoft Windows

Deploy and configure the Dell PS Series MPIO DSM for Windows Server to enable multi-path I/O for high availability access to a PS Series array.

Dell Storage Engineering
January 2017

Revisions

Date	Description
February 2008	Initial Release
June 2009	Host Integration ToolsV3.2 updates
February 2011	Windows® Server 2008 R2, 10 Gb support, and Host Integration Tools v3.5.1 updates
June 2013	Added support for HIT v4.5 and Windows® 2012
January 2017	New Format for Dell EMC

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2008 - 2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [1/25/2017] [Best Practices] [TR1036]

Dell EMC believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of contents

Revisions.....	2
Executive summary.....	4
Audience.....	4
1 Introduction.....	5
1.1 MPIO.....	5
1.2 Microsoft Multi-Path I/O.....	6
1.3 PS Series MPIO DSM.....	6
2 Getting started.....	8
2.1 10Gb networking with PS Series arrays.....	8
3 Installing, configuring and managing the PS Series MPIO DSM.....	10
3.1 Configuring the MPIO DSM.....	10
3.1.1 Default Load Balancing policy.....	12
3.1.2 Volume session connections.....	12
3.1.3 Microsoft iSCSI initiator.....	14
4 Connecting to volumes.....	17
4.1 Discovery of target portal.....	17
4.2 Connect to PS Series target.....	18
4.3 Make the PS Series Volume Available to the Server.....	20
4.4 Verify MPIO.....	21
4.5 Log off a volume with MPIO.....	22
5 Summary.....	24
A Using hardware initiators.....	25
B Tips, troubleshooting and hotfixes.....	26
B.1 Managing connections in large environments.....	26
C Technical support and resources.....	27

Executive summary

Dell™ Storage PS Series arrays optimize resources by automating performance and network load balancing. Additionally, PS Series arrays offer all-inclusive array management software, host software, and free firmware updates.

Audience


The information in this guide is intended for Microsoft® Windows Server® administrators.

1 Introduction

In our technology-driven world, data is one of the most valuable assets in any business organization. Increasingly, organizations are requiring that their data be available 24 x 7. Continuous availability of data is critical to all businesses. A data outage of any duration and size can result in a tremendous loss of revenue - perhaps thousands or even millions of dollars. High availability is commonly achieved by using fault tolerant and redundant components. For example, multiple servers may be used for redundancy, RAID or even failover clustering on the servers, but this type of protection alone is not enough because it relies on a storage area network (SAN) for its data. So, to ensure high availability, access to shared storage must be readily available and protected against failure. As SANs become populated with more and more data, the potential loss of access to data storage components becomes problematic. This is why many businesses today consider MPIO a must.

1.1 MPIO

Microsoft Multi-Path I/O (MPIO) framework uses redundant physical (or virtual) connections to deliver high availability to shared storage. Having multiple host connections, switches and SAN interfaces all help to eliminate single points of failure. Using MPIO, servers can send multiple I/O streams to SAN volumes concurrently. MPIO routes I/O over redundant paths that connect servers to storage and also manages these paths so that requests can be re-routed through another path in the event of a failure in one of the components along the way. MPIO also provides increased redundancy and can improve performance of application data hosted on the SAN. See Figure 1.



- 1: The eqlsm.sys kernel mode driver performs I/O path selection and error handling¹
- 2: The EHCMSservice.exe user mode service that manages iSCSI sessions¹
- 3: Two or more 1 Gb or 10 Gb Ethernet NICs or HBAs
- 4: Two or more redundant physical paths to redundant network switches
- 5: Two or more redundant physical paths provide a network switch interlink (ISL)
- 6: Two or more redundant physical paths from the network switches are connected to the separate redundant controllers in each Member in a PS Series Group.
- 7: Each Member in the Group has at least two physical paths to the network switches that are connected to separate redundant controllers.

¹This component is available when the Dell PS Series MPIO DSM is installed.

Figure 1 Basic MPIO configuration and components

1.2 Microsoft Multi-Path I/O

Microsoft MPIO enables storage vendors to develop multi-path modules that contain specific information about their storage hardware. This ultimately results in more efficient and intelligent use and connectivity to the SAN hardware. These vendor-developed plug-ins are referred to as Device Specific Modules (DSM). Dell has developed a Device Specific Module that plugs in to the Microsoft MPIO framework and is used to optimize Windows Server usage of PS Series storage arrays, Figure 2.

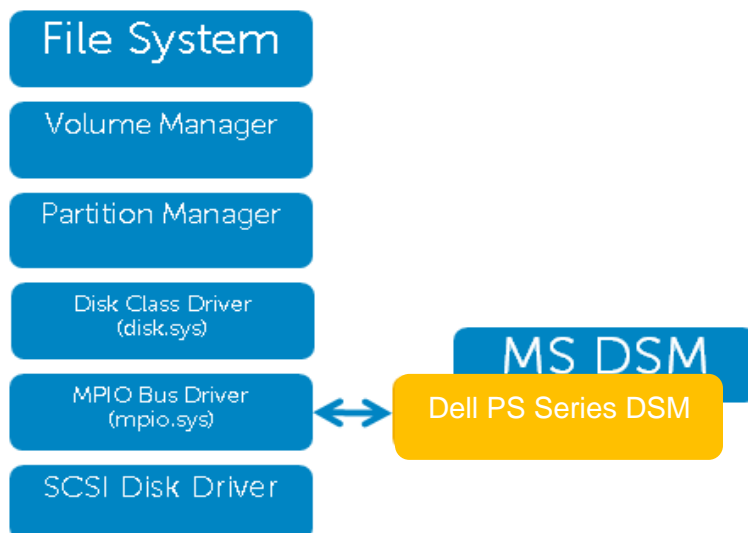


Figure 2 PS Series MPIO DSM is a driver component in the I/O stack and leverages the Microsoft MPIO framework

PS Series MPIO DSM is easy to install and is included as part of the Host Integration Tools for Microsoft (HIT/Microsoft). This HIT kit is included at no extra cost with the purchase of a PS Series array and can be downloaded from eqsupport.dell.com (login required).

1.3 PS Series MPIO DSM

The PS Series MPIO DSM is designed to deliver:

- Automatic connection management
- Automatic path failure detection and path failover
- Automatic load balancing across multiple paths
- Support for multiple connections to a single iSCSI target (volume)
- Increased bandwidth
- Reduced network latency
- Easy installation and management
- Automatically throttle down connections when a storage pool reaches 90% of maximum allowable connections

The PS Series DSM consists of two components:

1. A driver component in the I/O stack that works in conjunction with the Microsoft MPIO driver to route I/O to the desired path. This is the PS Series DSM for Microsoft Windows. It selects the best path for each I/O by using knowledge about how the volume is laid out on the PS Series Group. See Figure 3.
2. A Microsoft Windows Server service (PS Series Host Connection Manager – EHCM) that manages connections. The service periodically retrieves the current volume layout from the PS Series Group. It also monitors the iSCSI sessions and SAN topology and will reconfigure the iSCSI session mesh when necessary. See Figure 3.

Together, these components allow administrators to easily install and configure MPIO for iSCSI networks. The DSM components, along with a redundant iSCSI hardware configuration create the infrastructure needed for a completely fault tolerant and efficient MPIO solution.




Figure 3 PS Series MPIO DSM automatically manages iSCSI sessions and routes I/O directly to the member servicing it

Multiple I/O streams are concurrently routed through the host iSCSI sessions directly to the PS Series member servicing them. Using the PS Series MPIO DSM, the best path for each I/O is selected by using knowledge about the volume layout on the PS Series group. The PS Series DSM automatically manages these iSCSI sessions.

2 Getting started

Before installing PS Series HIT/Microsoft (that allows configuration of the PS Series MPIO DSM), plan the networking configuration.

- More than one network interface should be dedicated to iSCSI SAN traffic from each server.
- Separate the iSCSI SAN traffic from LAN network. This is especially important when implementing 10Gb and 1Gb infrastructures.
- The following two documents provide additional information and should be consulted before getting started.
 - [Dell PS Series Configuration Guide](#)
 - [Dell PS Series Storage iSCSI Initiator and Operating System Considerations](#)

2.1 10Gb networking with PS Series arrays

10 Gigabit Ethernet is designed to handle ten times the throughput as Gigabit Ethernet. If planned and configured properly, great throughput and performance increases can be expected on the iSCSI SAN. Ideally, all networking components from the server, through the network switch and to the SAN should support a 10Gb infrastructure and have Jumbo Frames enabled. This helps to prevent potential bottlenecks along the network path. When planning a 10Gb SAN infrastructure, restrict management access to 1Gb interfaces, so that the 10Gb NICs can be dedicated to SAN traffic. These settings can be confirmed through the PS Series Group Manager GUI. See Figure 4.

Lastly, as a best practice when mixing 1Gb and 10Gb PS Series arrays, dedicate a Storage Pool to 10Gb PS Series members, so that volumes are not potentially cross-striped across 1Gb and 10Gb arrays - which would not take advantage of a truly optimized 10Gb solution. If a Storage Pool crosses 1Gb and 10Gb members, then the performance will be potentially limited by the 1Gb member(s). The storage group dedicated 10Gb storage pool and the firmware network load balancer, in conjunction with the server PS Series MPIO DSM and EHCM service will ensure that the 10Gb iSCSI connections are always optimally routed to the proper 10Gb member of the storage group - even in a multi-member group comprised of both 10Gb and 1Gb arrays.

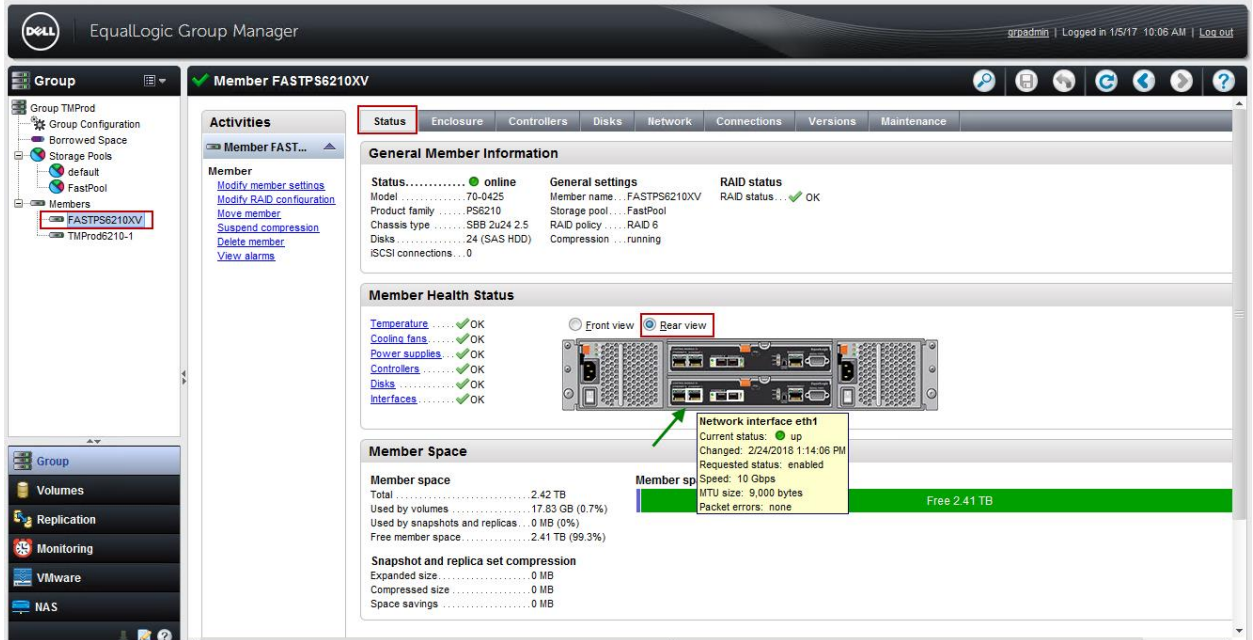


Figure 4 PS Series Group Manager

From the PS Series Group Manager, highlight the member on the left. From the **Status** tab, click on **Rear View** and mouse over the corresponding network interfaces to confirm Jumbo Frames. Lastly, modify IP settings to confirm restriction of management traffic to 1Gb management network interface.

3 Installing, configuring and managing the PS Series MPIO DSM

The PS Series MPIO DSM is included with HIT/Microsoft. For installation guidelines, please refer to the *Host Integration Tools for Microsoft User Guide* and Release Notes at eqlsupport.dell.com.

The following section assumes that access control to PS Series volumes has been configured with CHAP authentication, multiple redundant NIC IP addresses, or a unique iSCSI initiator name. Volume access control can be added to volumes through the PS Series Group Manager, either at the time of volume creation or afterwards.

A single ACL record for each host is typically sufficient for enabling volume connections from the MPIO DSM. When using IP addresses in the ACL, add an ACL record for each interface the host will use to connect to the volume.

For volumes connecting to cluster configurations, additional steps are needed to assure proper connection. Select the **Enable shared access to the iSCSI target from multiple initiators** box in the volume settings and add an ACL record for each cluster host adapter connecting to the volume. Additionally, there may be some specific configuration changes needed, including registry changes, to assure that each host in the cluster is using persistent reservations. Please refer to the latest version of *Release Notes for Host Integration Tools for Microsoft* for more information. This documentation is available at <https://eqlsupport.dell.com>.

3.1 Configuring the MPIO DSM

This section assumes PS Series HIT/Microsoft has been installed. This section covers:

- Using the PS Series Auto-Snapshot Manager to Configure MPIO
- Default Load Balancing policies

Launch the Dell EqualLogic Auto-Snapshot Manager application (EqualLogic is also referred to as PS Series). This is installed as part of the HIT/Microsoft.

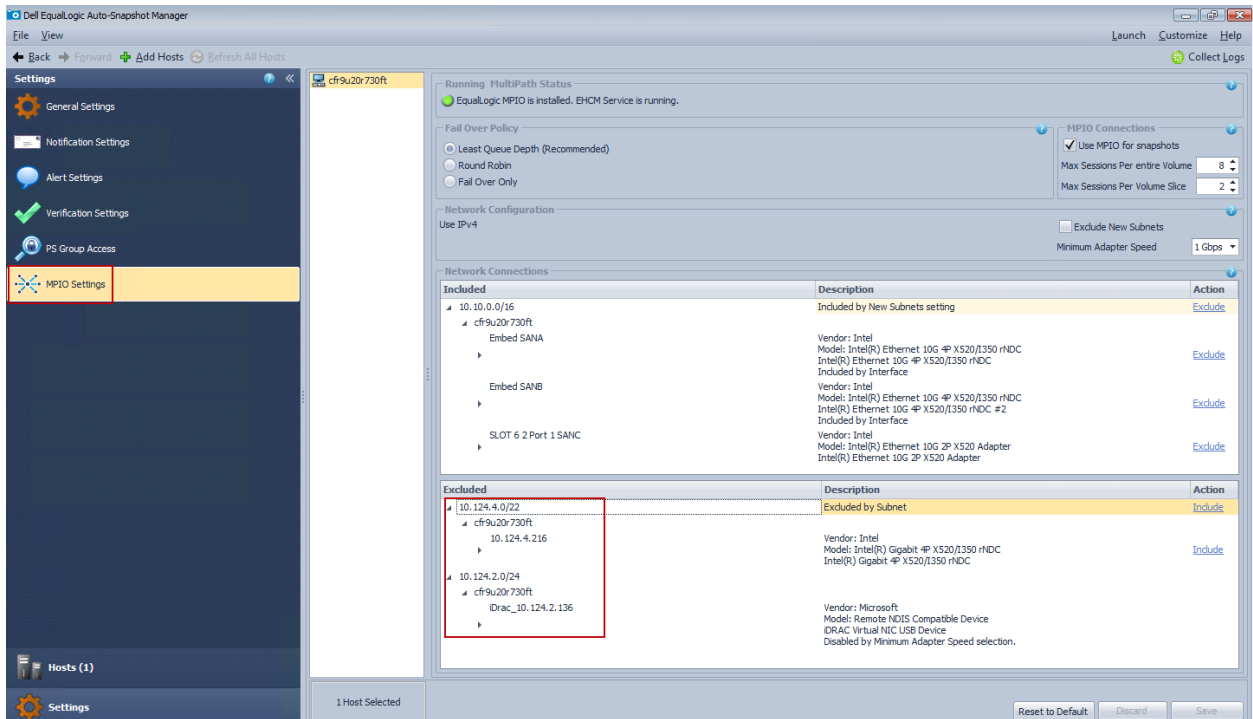


Figure 5 MPIO Settings showing included and excluded network connections

Select **Settings** at the bottom left of the application and click on **MPIO Settings** in the left hand object tree.

This will open the MPIO settings configuration view. See Table 1 for all the MPIO setting options.

Configure MPIO settings for a selected range of hosts. Select the subnets or individual adapters to Exclude MPIO and iSCSI subnets from SAN traffic. Optionally set the minimum adapter speed to match the network configuration. This will ensure that the MPIO DSM directs iSCSI traffic only through the desired network interfaces. See the next section for details on Default Load Balancing Policy and Max sessions for per volume. Click **Save** when finished.

Table 1 Auto-Snapshot Manager MPIO Settings Options

Setting	Description
Network Connections - Included	The IP range for subnets that are intended to be used for MPIO. (should be SAN only subnets)
Network Connections - Excluded	The IP range for subnets that are intended to be excluded from MPIO.
Fail Over policy	The default load balancing policy to use for MPIO sessions. (Recommend using Least Queue Depth)

Setting	Description
Max sessions per volume slice	The max number of session connections permitted per volume per member. This number represents the # of connections a volume will make to each member or array in the storage pool.
Max sessions per entire volume	The max number of session connections permitted for a volume. This value must be greater than or equal to the max sessions per volume slice.
Use MPIO for snapshots	Optionally include or exclude MPIO for connections to snapshots.
Minimum adapter speed	Slowest acceptable NIC speed to be used by MPIO

3.1.1 Default Load Balancing policy

When the PS Series 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. Generally, it is recommended to accept the default load balancing policy because it is optimized 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. These options are not available through the Remote Setup Wizard and will return an error message if they are selected using the iSCSI Initiator Properties for the target.

3.1.2 Volume session connections

The number of connections per target (volume) depends on the number of arrays in the storage pool and the number of NICs being used for MPIO. The two controlling settings for target connections are:

- Max sessions per entire volume
- Max sessions per volume slice (pieces of the volume)

By default, the max number of sessions per entire volume is set to 8 and the max number of sessions per volume slice is set to 2. Refer to Table 1 for descriptions of these settings.

The max number of sessions per slice are based on the number of arrays in the storage pool where the volume resides. Volume slice refers to the slice (or piece) of the volume on the member storage array. Each

volume will load balance across the arrays in the storage pool and each array will host a slice of the volume. The number of session connections per volume can be calculated based on the following equations.

2 sessions per volume slice * 1 members = 2 sessions per entire volume

2 sessions per volume slice * 2 members = 4 sessions per entire volume

2 sessions per volume slice * 3 members = 6 sessions per entire volume

Note: These calculations will change slightly with single port 10Gb PS Series array controllers. In this case the sessions per volume slice will default to 1.

The MPIO connections setting can be modified using Auto-Snapshot Manager in the **MPIO Settings** window, with the instructions on the previous page.

0 represents a volume distributed across three members with two iSCSI sessions per slice balanced across the NICs. Slices are pieces of the volume and are split across the members appropriately in accordance with the capacity of the members.

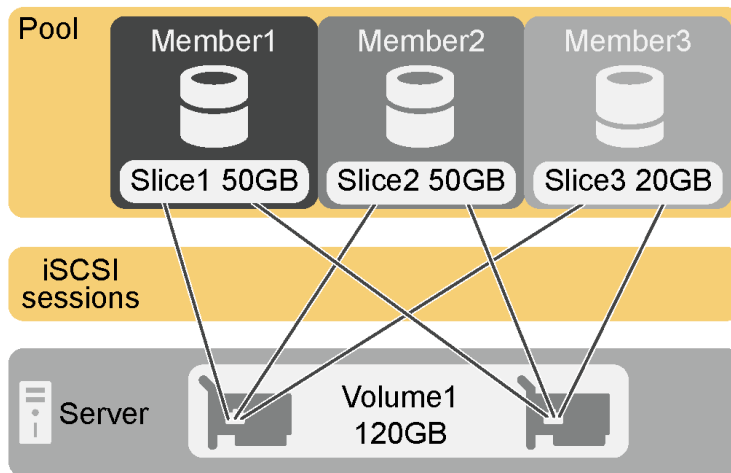


Figure 6 Slices across members presented to the volume

3.1.3 Microsoft iSCSI initiator

This section describes the Microsoft iSCSI initiator interface, including the Dell EqualLogic MPIO tab.

Launch the Microsoft iSCSI Initiator. See Table 2 on the next page for more information on each tab displayed.

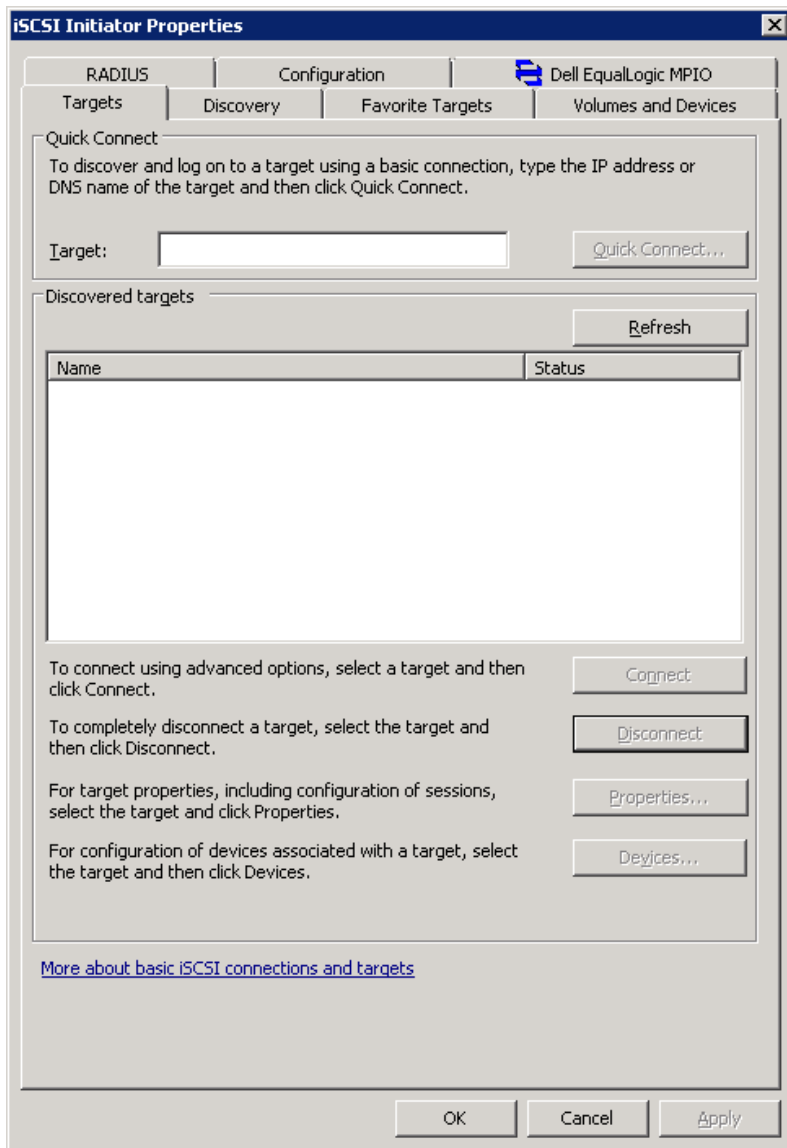


Figure 7 iSCSI initiator tool

Table 2 Microsoft iSCSI initiator tabs and descriptions

Tab Name	Available options...
Discovery	Add or remove and configure iSCSI portal settings
Targets	Connect or disconnect and view properties of iSCSI targets
Favorite Targets	View and configure persistent volume connections
Volumes and Devices	Configure drive letter bindings to a particular volume
RADIUS	Add and configure a RADIUS Server for RADIUS authentication
Configuration	<ul style="list-style-type: none"> • Change IQN Name • Set CHAP credentials • Configure IPsec settings • Generate reports of all connected targets/devices on system
Dell EqualLogic MPIO	<ul style="list-style-type: none"> • View subnets included/excluded for MPIO • To change MPIO Settings, launch the Auto-Snapshot Manager

The **Dell EqualLogic MPIO** Tab has been added to the Microsoft iSCSI Initiator properties window by the Host Integration Tools for Microsoft setup.

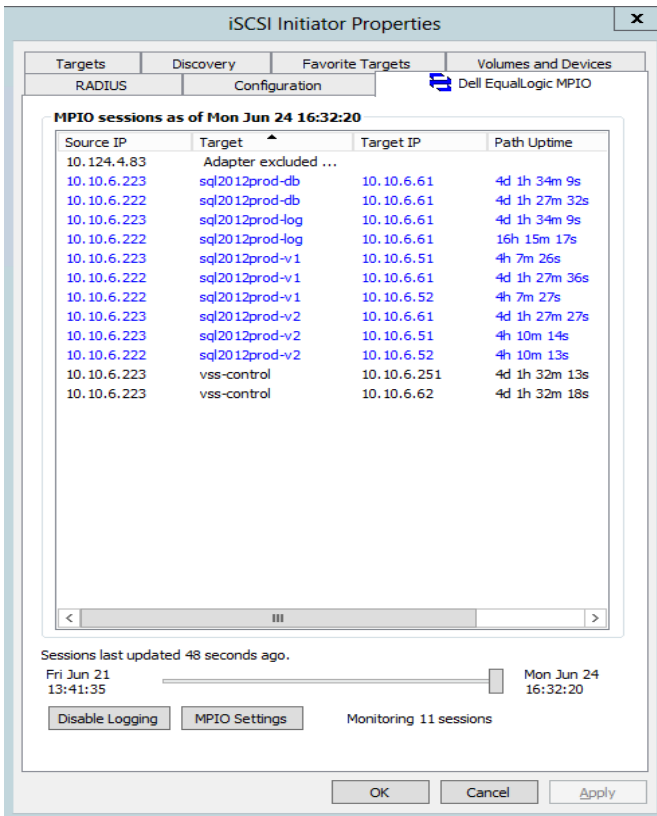


Figure 8 Dell EqualLogic MPIO tab

Both the included and excluded subnets are visible for MPIO as well as all iSCSI MPIO iSCSI Sessions. Note the even distribution of iSCSI sessions to each volume across each NIC. Finally, the scroll bar allows the ability to examine the MPIO status at a specific point in time.

Table 3 Dell EqualLogic MPIO Tab Information

Component	Description
Source IP	The IP address of the source device. This is the NIC or HBA installed in the computer that accesses an iSCSI target.
Target	The name of the iSCSI volume that is the target of the iSCSI session. When MPIO is configured, a listing will be shown for each active session to the same volume.
Target IP	The IP address of the PS Series array's Ethernet port
Path Uptime	The elapsed time during which this path is connected.
Managed	Indicates whether the sessions are actively managed. The text color also indicates this. Sessions colored in blue are managed while sessions colored black are not.

4 Connecting to volumes

If Auto-Snapshot Manager was used to configure PS Group access from the host then the Discover Target Portal Step can be skipped as Auto-Snapshot Manager will do this automatically. Please refer to the *Auto-Snapshot Manager User Guide* for more information about setting PS Series Group access. To configure the Target Portal manually, please follow the steps in the next section.

4.1 Discovery of target portal

If Auto-Snapshot Manager was not used, manually adding the discovery of the target portal will be needed.

First the MPIO must be installed and configured, then this section will step through connecting to a volume.

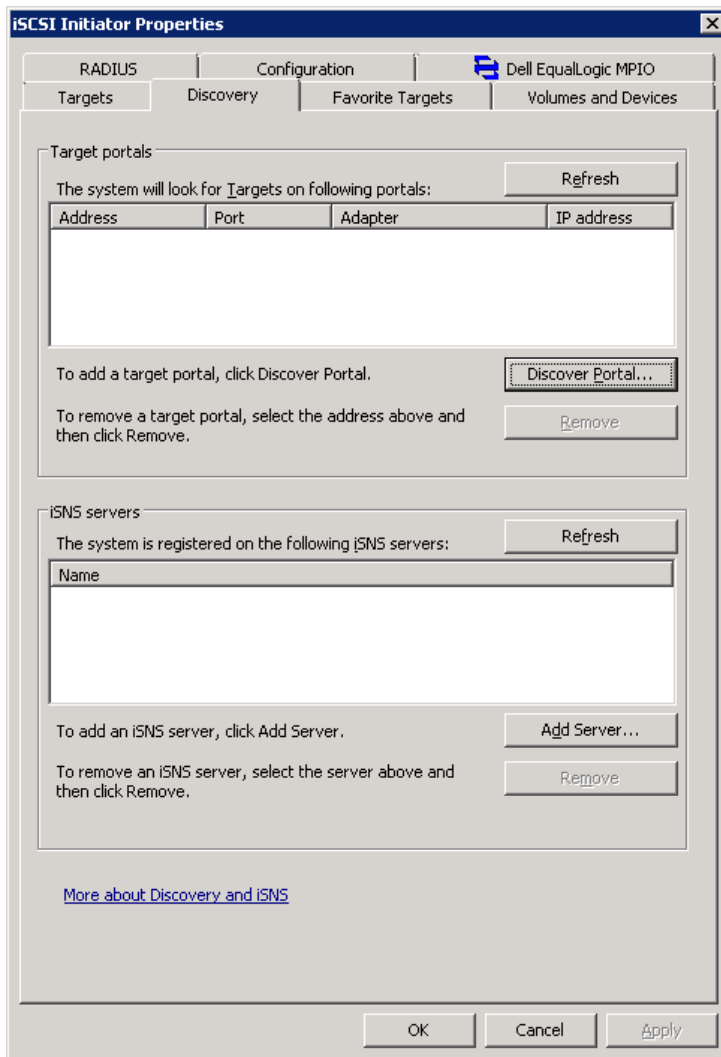


Figure 9 Discovery tab in the iSCSI Initiator tool

In the **Discovery** tab, click **Discover Portal**. Previous to this step, access control for PS Series volumes should have been already configured with CHAP authentication, multiple redundant NIC IP addresses, or a

unique iSCSI initiator name. Volume access controls are configured through the PS Group Manager, either at the time of volume creation or afterwards.

Add the PS Group Portal to the Initiator List

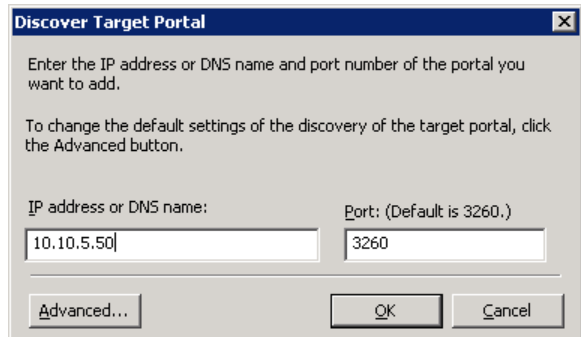


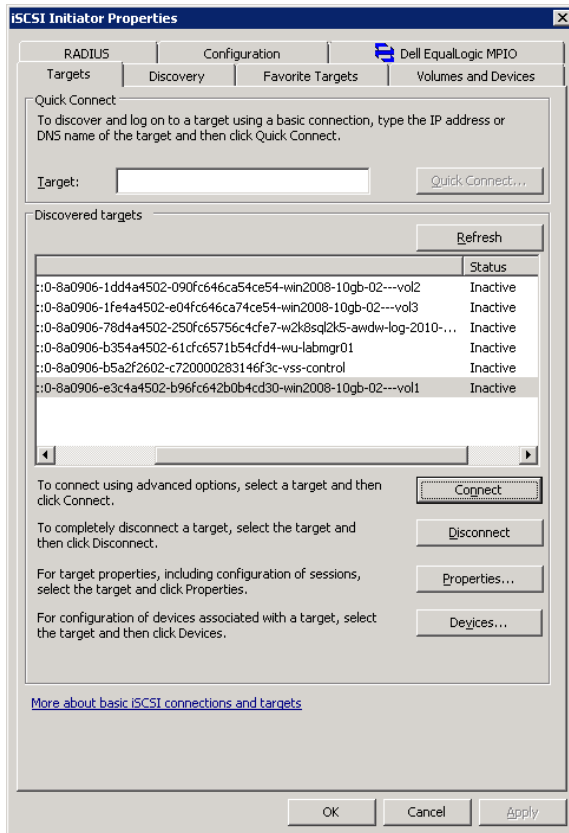
Figure 10 Discovery Target Portal

In the **Discover Target Portal** dialogue box, enter the PS Series group IP address and port. Consult the networking team for anything other than the default port number. Click **OK** to complete. The PS Group IP should now be listed in the Target Portals list.


Note: The discovery target portal IP address is not the PS Series group management IP. The PS Series group IP is found in the Group Manager GUI on the **General** tab of the Group Configuration.

4.2 Connect to PS Series target

Now that the iSCSI portal has been configured, click on the **Targets** tab in the initiator and Refresh the targets list.



Highlight the target of choice from the list and click **Connect**.



To make this iSCSI connection persistent after a reboot, select **Add this connection to the list of Favorite Targets**. The connection will then be added to the **Favorite Targets** tab of the iSCSI Initiator properties. To use CHAP credentials, click **Advanced**, enter the CHAP authentication information and click **OK**. Select **Enable multi-path** and click **OK**.

4.3 Make the PS Series Volume Available to the Server

Next, make the disk available to the server through Server Manager's Disk Management. Right-click on **Disk Management** and select **Re-Scan**. After Re-Scan, right-click the disk to set it online.

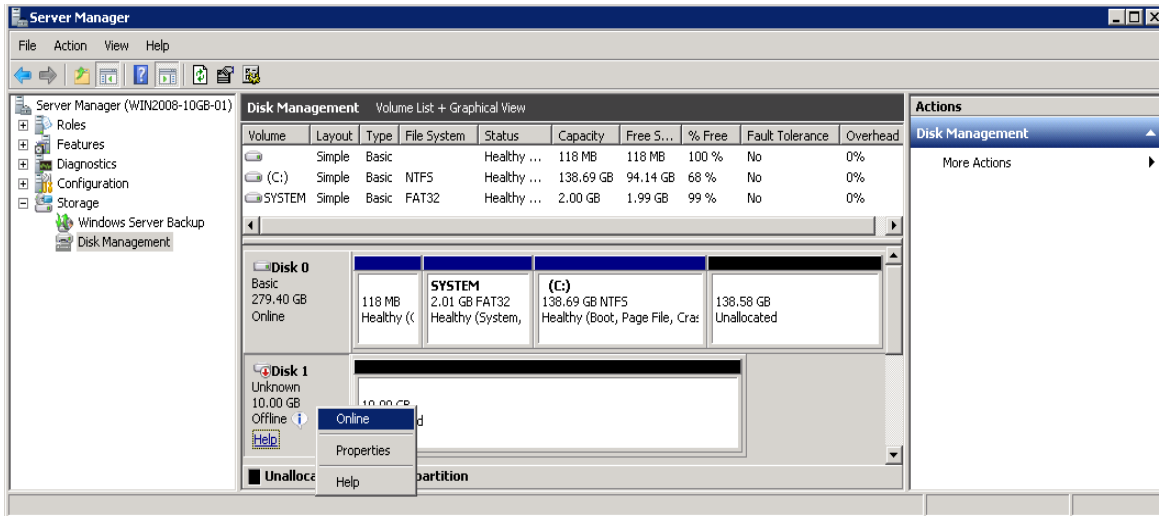


Figure 11 Disk Management

Now, use Server Manager/Disk Management, to format the volume as NTFS and assign a drive letter.

Optionally bind the newly available PS Series volume to a drive letter using the iSCSI Initiator properties however the volume must first be a Favorite Target. To configure a Favorite Target, see the [Connect to PS Series Target](#) section above. When connecting to the target, select the checkbox to **Add this connection to the list of Favorite Targets**.

From the **Favorite Targets** tab, either manually add the drive letter by clicking **Add** or choose **Auto Configure** to configure all available devices. Optionally, remove a device or clear the entire list by clicking **Remove** or **Clear**.

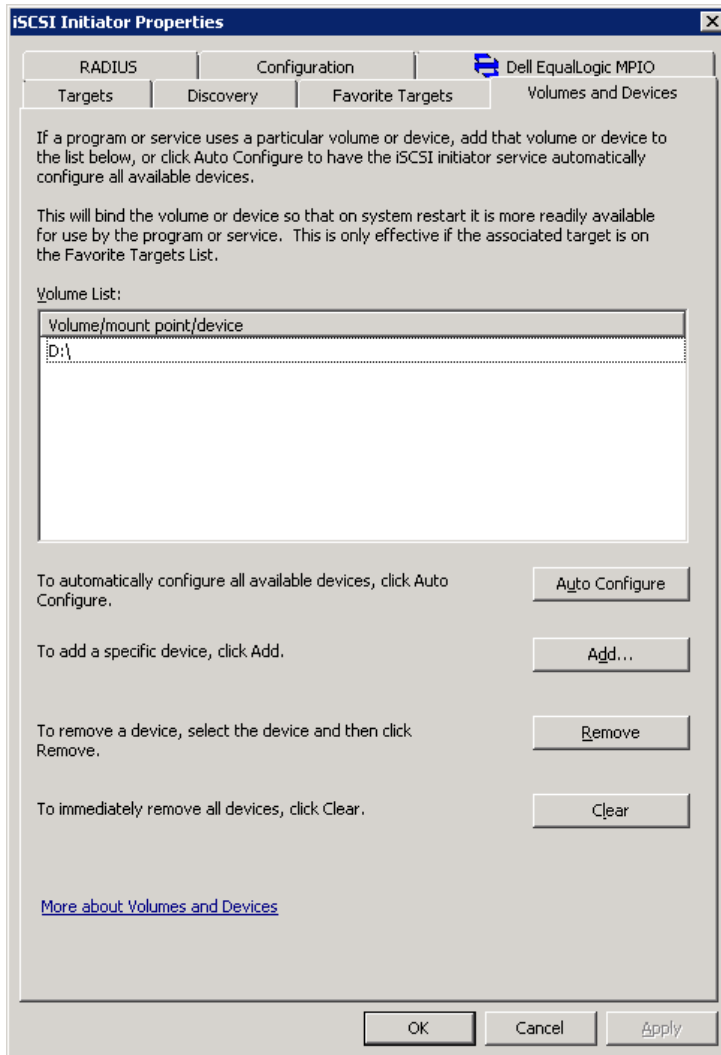


Figure 12 Adding the drive letter from the iSCSI Initiator tool's Volumes and Devices tab

4.4 Verify MPIO

There are two common places to look to verify volume session connections to the array.

1. The Microsoft iSCSI Initiator **Dell EqualLogic MPIO** Tab
2. From Server Manager/Disk Management:
Right-click the target disk and select **Properties**. The **MPIO** tab will display the MPIO Policy for the selected PS Series volume, as well as the multiple paths to the volume. The example below shows the Dell EqualLogic MPIO DSM Load Balancing Least Queue Depth Policy is enabled and that there are multiple (4) connection paths to the volume.

Note: The session connections in the initiator **Dell EqualLogic MPIO** tab will take a minute or two to show up after logging into the volume.

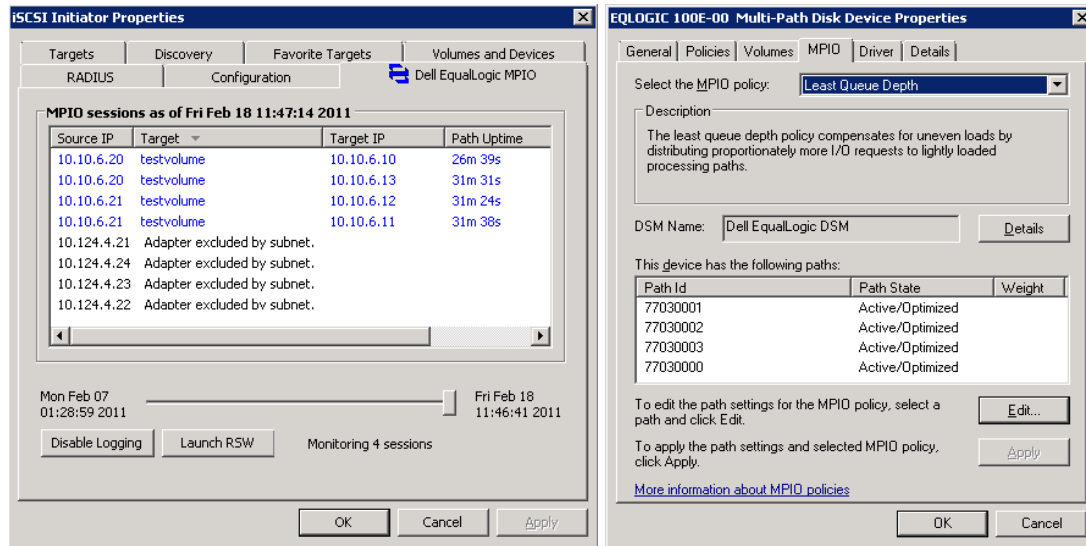


Figure 13 iSCSI Initiator and Disk Management screens

4.5 Log off a volume with MPIO

From the Microsoft iSCSI Initiator **Favorite Targets** tab, highlight each favorite target and Remove it. The Windows server will no longer try to re-connect to the target volume on startup.

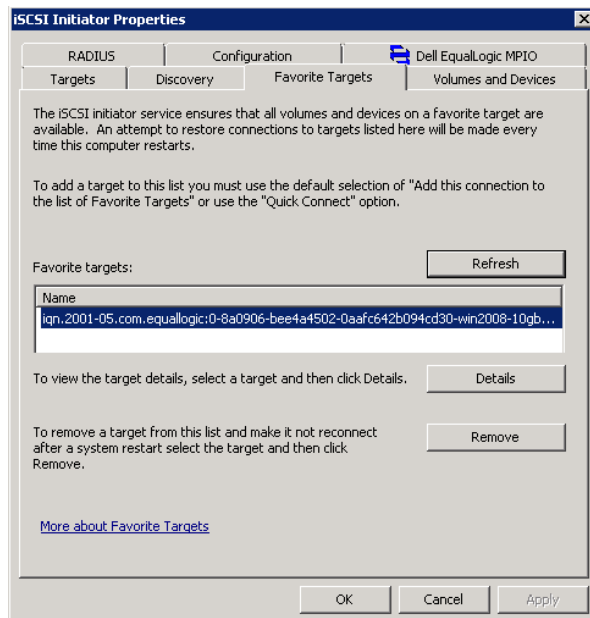


Figure 14 iSCSI Initiator **Favorite Targets** tab

From the **Targets** tab, highlight the target and click **Disconnect** to logoff each volume. This will disconnect the target volume from the Windows server.

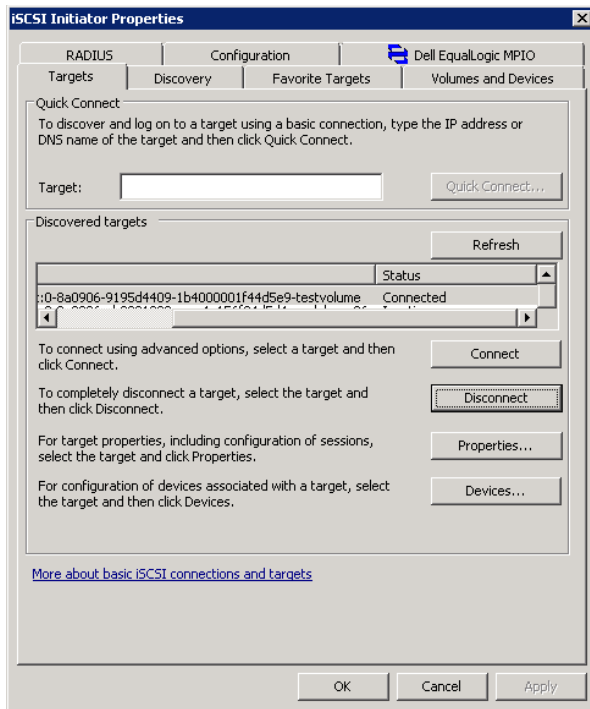


Figure 15 iSCSI Initiator's Targets tab

5 Summary

The Dell PS Series array supports multiple iSCSI SAN connections for performance and reliability. MPIO provides multiple paths from servers to storage, delivering fault tolerance, high availability, and improved performance. The PS Series MPIO DSM simplifies MPIO setup and configuration in a Microsoft Windows server environment by automatically managing multiple iSCSI sessions and also optimizes the operation of a PS Series array by using knowledge about volume layouts on a PS Series group and routing I/O directly to the PS Series arrays that will be servicing requests. Additionally, the PS Series MPIO DSM provides a Windows Server transparent redundancy in its iSCSI connections to a PS Series SAN.

The Host Integration Tools for Microsoft and documentation are available by logging onto the support website at <https://eqsupport.dell.com/support/download.aspx?id=3199>.

A Using hardware initiators

In addition to the Microsoft iSCSI Initiator, the PS Series DSM supports hardware initiators. At the time of this document, the Dell PS Series MPIO DSM supports the QLogic HBAs and iSCSI Offload for Broadcom CNA's. For information on these hardware initiators, reference the [Dell PS Series Configuration Guide](#). When using the Broadcom adapters with iSCSI Offload functionality, there are two methods to connect to iSCSI storage through the adapter:

- 1) Using the Microsoft iSCSI Software Initiator, as described in this document.
- 2) Using the iSCSI functionality implemented within the adapter hardware itself. The Dell PS Series EHCM service automatically prefers the hardware implementation if it is enabled and configured and iSCSI sessions are established using these interfaces instead of the iSCSI software interfaces. To use the Microsoft iSCSI software initiator instead, disable the iSCSI Offload functionality within the Broadcom Advanced Control Suite (BACS).

The Multi-Path I/O DSM connection manager uses temporary CHAP credentials to create additional iSCSI sessions to each target. As of this report, QLogic iSCSI HBAs do not support the use of CHAP credentials for these sessions. Therefore, one of the following must be done for each target to use MPIO in these configurations:

- 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 to be used for SAN.
- Create multiple access control records, one for each of the host computer's interfaces to be used for the SAN, where each record contains a specific IP addresses.

When using MPIO with a PS Series group running firmware prior to 5.0, CHAP authentication must be enabled. If the PS Series group is not configured to use local CHAP authentication, the server will receive login errors when attempting to create MPIO sessions. Use the Group Manager GUI to enable local CHAP authentication:

1. Click Group Configuration and open the iSCSI tab to display the options for iSCSI initiator authentication.
2. Check **Consult locally defined CHAP accounts** first.

B Tips, troubleshooting and hotfixes

Dell recommends that all hardware and software be kept up-to-date with patches and latest firmware and driver revisions. The latest version of the Dell Storage Compatibility Matrix can be found at:

<http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/20438558>

For details on hotfixes and recommendations when using Windows Servers and MPIO with Dell PS Series arrays, please consult the Dell PS Series Arrays iSCSI Initiator and Operating System Considerations at:

<http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/20371245>

B.1 Managing connections in large environments

Large cluster environments tend to push the maximum allowable connections to the PS Series group when using the PS Series MPIO modules. There are a number of ways to minimize and control connections including scaling down the Max sessions per entire volume and Max sessions per volume slice. For additional information and details see the PS Arrays: Scalability and Growth in Virtual Environments technical report at:

<http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19992296>

C Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

[Dell TechCenter](#) is an online technical community where IT professionals have access to numerous resources for Dell EMC software, hardware and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell EMC Storage platforms.