

Configuring and Installing the PS Series Multipathing Extension Module for VMware vSphere and PS Series

Abstract

This document describes the benefits of the Dell™ PS Series Multipathing Extension Module (MEM) for VMware® vSphere® that provides MPIO for highly available access to the PS Series SAN.

November 2019

Revisions

Date	Description
November 2011	Initial release
February 2012	General availability updates
September 2012	Updated Storage Heartbeat recommendation for vSphere 5.1
September 2013	Updated to reflect firmware 6.0 and vSphere 5.5
June 2015	Updated to include Virtual Volume support
July 2017	Updated to include web client UI changes to vSphere Update Manager, and MEM 1.5
November 2019	vVols branding update

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.

© 2011–2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Dell 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
Acknowledgements.....	Error! Bookmark not defined.
Table of contents	3
Executive summary.....	4
1 Introduction.....	5
1.1 Audience.....	5
2 Deploying the MEM	6
2.1 Deployment considerations and requirements	6
2.2 Deployment considerations for iSCSI connection count	6
2.3 Deployment considerations for iSCSI login timeout	8
2.4 Deployment considerations and best practice parameters	8
3 Configuring an iSCSI vSwitch for multipathing.....	9
3.1 Interactive mode configuration	9
3.2 Unattended mode configuration	11
4 Installing the MEM	12
5 Verification of MEM iSCSI session creation	13
6 Advanced iSCSI connection configuration parameters.....	15
6.1 Increasing the default values.....	15
6.2 Decreasing the default values	16
6.3 Setting the EHCM configuration values under vSphere.....	16
7 Using MEM with Virtual Volumes	18
7.1 MEM configuration and installation.....	18
7.2 Access controls.....	18
7.3 Virtual Volumes impact on iSCSI connection count	18
8 Summary	19
A Example iSCSI vSwitch configurations	20
B Installing the MEM with VMware Update Manager	21
C Using the setup.pl script to configure Round Robin MPIO.....	24
D Software and firmware versions	25
E Technical support and resources	26

Executive summary

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

High availability is an important requirement of any system in the data center. This availability is especially critical if that system is a component in the virtual infrastructure where a virtualized datacenter is built. Redundant hardware and RAID technologies form a critical foundation. When using shared storage, the paths from the servers to the storage need to also be redundant and highly available.

This paper details the benefits of the Dell Multipathing Extension Module (MEM) for VMware® vSphere®, as well as the installation and configuration process to provide Multi-Path I/O (MPIO) for high available access to PS Series storage. Also covered are a number of overall virtual environment iSCSI design considerations and best practices.

1 Introduction

VMware vSphere offers many enhancements to the software iSCSI initiator beyond basic iSCSI SAN connectivity. The most significant of these enhancements is the API support for third-party multipathing plugins. This provides a framework that enables the MEM to intelligently route and efficiently load balance iSCSI traffic across multiple NICs.

The MEM provides the following benefits:

- Ease of installation
- Increased bandwidth
- Reduced network latency
- Automatic failure detection and failover
- Automatic load balancing across multiple active paths
- Automatic connection management
- Multiple connections to a single iSCSI target

The MEM utilizes the same multipathing iSCSI vSwitch as VMware Round Robin multipathing. As part of the continuous efforts at Dell to help customers simplify the configuration of their IT environments, the iSCSI MPIO vSwitch configuration process for the MEM has been reduced to either a single command with a few parameters or a guided question-and-answer process. The MEM installation process is equally straightforward, requiring that only a single command be executed. The MEM can also be installed using the VMware vSphere Update Manager.

Once installed, the MEM automatically creates iSCSI sessions to each member that a volume spans. As the storage environment changes, the MEM responds by automatically adding or removing needed iSCSI sessions.

As storage I/O requests are generated on the VMware ESXi™ hosts, the MEM intelligently routes each requests to the appropriate array member. This results in efficient load balancing of the iSCSI storage traffic, reduced network latency and increased bandwidth.

1.1 Audience

The information in this guide is intended for administrators of a VMware vSphere environment utilizing PS Series iSCSI arrays.

2 Deploying the MEM

Deploying the MEM consists of two steps:

1. Configuring a vSwitch for iSCSI multipathing
2. Installing the MEM

The entire process of creating the multipathing-compatible vSwitch and installing the MEM is quick and efficient. Once deployed, it provides both new and existing ESXi hosts with increased performance to PS Series storage resources.

2.1 Deployment considerations and requirements

The MEM has the following requirements:

- VMware vSphere ESXi with Standard licensing (see Table 1)
- vSphere Command-Line Interface (CLI) or vSphere Management Assistant (vMA) compatible with above version of ESXi
- Compatible PS Series array firmware (refer to the release notes on eqlsupport.dell.com — login required)

Prior to deploying the MEM, complete the following steps:

1. Place the ESXi host(s) in maintenance mode.
2. Download and install either the VMware vMA or vSphere CLI. Refer to [VMware Documentation](#) to install and configure these tools.
3. Download the EqualLogic Multipathing Extension Module for VMware from eqlsupport.dell.com (login required). The download is a zip archive named EqualLogic-ESX-Multipathing-Module.zip.
4. Unpack the zip archive files:
 - MEM-Release_Notes.pdf
 - MEM-User_Guide.pdf
 - README.txt
 - setup.pl
 - dell-eql-mem-esx<vSphere_version>-<release version>.zip

Note: Do not extract dell-eql-mem-esx<vSphere_version>-<release version>.zip. It is a vSphere Installation Bundle (VIB) offline bundle.

5. If using the vMA, upload setup.pl and the VIB offline bundle to the vMA.
6. Retain the MEM-Release_Notes.pdf, MEM-User_Guide.pdf (*EqualLogic Multipathing Extension Module Installation and User Guide*), and README.txt files for reference.

2.2 Deployment considerations for iSCSI connection count

The MEM provides increased bandwidth and lower latency access to PS Series storage. This performance benefit is achieved through leveraging several behaviors that are unique to the PS Series peer storage architecture. It is important to understand how MEM achieves performance gains, and the potential impact on storage if it is not deployed.

The MEM achieves its performance gains by creating multiple iSCSI connections to each PS Series group member that has a datastore volume. Assuming the MEM default settings are in use, and depending on the configuration of the environment, there will be up to six iSCSI connections to a datastore volume. This is a significant increase, when compared with the VMware Fixed path policy, which utilizes one iSCSI connection per volume. This increase in connections must be planned at the time of deployment, and again as the vSphere environment and its supporting PS Series storage environment scales.

To calculate the number of iSCSI connections that the MEM will consume in a particular environment, utilize the following formula:

$$\text{Number of iSCSI connections} = H \times V \times M \times 2$$

H = number of ESXi hosts

V = number of volumes

M = number of PS Series members in the pool (maximum value is 3)

2 = default number of iSCSI sessions per volume portion

For example, in a vSphere environment with eight ESXi hosts, and using a two-member PS Series pool for hosting the ten volumes needed for virtual machines, the number of iSCSI connections would be:

$$8 \times 10 \times 2 \times 2 = 320 \text{ iSCSI connections}$$

While this is within the 1,024 iSCSI connection limit per pool of the firmware, it does not take into account other iSCSI connections to the PS Series storage pool from other servers in the data center. Consider how this virtual environment may grow in the future with respect to additional ESXi hosts, datastore volumes, and additional PS Series members.

Using the SCSI commands, MEM and the array can communicate certain information. This enables the array to notify the MEM of the number of total iSCSI connections to the pool from all sources. Should the number of iSCSI connections to the pool begin to approach the limit, the MEM will reduce the number of iSCSI sessions it is creating to the array, while still maintaining redundancy.

Depending on the overall environment and future requirements, changes may need to be made to one or more of the following variables if there are concerns about the iSCSI connection count growing too large:

- The MEM parameters `membersessions`, `volumessessions` or `totalsessions` (see the *EqualLogic Multipathing Extension Module Installation and User Guide* for details).
- Increasing the size of a given volume to allow more virtual machines to reside on it and reduce the total number of volumes needed.
- Reconfigure larger clusters into multiple smaller clusters to reduce the number of hosts accessing particular volumes and the number of iSCSI connections consumed.
- Reduce the number of PS Series members per pool and increase the number of pools to increase the number of available iSCSI connections.

For more information on this topic, read the document, [Dell Storage PS Series Arrays: Scalability and Growth in Virtual Environments](#).

The MEM iSCSI connection algorithm works to avoid creating unnecessary sessions while still maintaining redundancy. The Dell PS6110, PS4110, and PS-M4110 arrays have only one 10Gb Ethernet port and therefore fewer iSCSI connections are created compared to arrays with multiple Ethernet ports per controller. With the PS6110, PS4110, and PS-M4110 arrays in a single-member pool, MEM creates two iSCSI connections per volume, for redundancy. In a two- or three-member pool, MEM creates a single iSCSI

connection per member, resulting in creating two or three iSCSI connections per volume. This balances the bandwidth and throughput from an individual ESXi host with that of an individual PS Series member.

Note: The vSphere 6.0 feature Virtual Volumes changes the way the iSCSI connections are made between the host and the volumes. See section 7, Using MEM with Virtual Volumes, for details.

2.3 Deployment considerations for iSCSI login timeout

The default value of 5 seconds for iSCSI logins on vSphere ESXi is too short in some circumstances, such as a large configuration where the number of iSCSI sessions to the array is close to the limit of 1024 per pool. If a severe network disruption were to occur, such as the loss of a network switch, a large number of iSCSI sessions need to be reestablished. With such a large number of logins occurring, completely processing the logins takes longer than the five-second default timeout period. The MEM installer automatically sets the iSCSI Login Timeout to the value of 60 seconds as recommended by Dell EMC.

The iSCSI Login Timeout value can be set using `esxcli` with the following syntax:

```
esxcli iscsi adapter param set --adapter=vmhba --key=LoginTimeout --value=60
```

Note: By default, the MEM `setup.pl` will attempt to update this value at configuration time.

2.4 Deployment considerations and best practice parameters

The MEM `setup.pl` script for configuring the ESXi host iSCSI vSwitch can be pass the `--bestpractice` parameter, which enables the following best practices:

Disabling Delayed ACK: Delayed ACK is a TCP/IP method of reducing I/O overhead by allowing segment acknowledgements to piggyback on each other or on other data passed over a connection. One side effect of delayed ACK is that if the pipeline is not filled, acknowledgement of data is delayed. In SAN Headquarters (SANHQ), this can be seen as higher latency during lower I/O periods. Latency is measured from the time the data is sent to when the acknowledgement is received. Since we are talking about disk I/O, any increase in latency can result in poorer performance. Additional information can be found in [VMware Knowledge Base article 1002598](#).

Disabling Large Receive Offload: Similar to Delayed ACK, Large Receive Offload (LRO) works by aggregating packets into a buffer before the received data is sent to the TCP stack. With iSCSI storage, this additional latency inserted into the process could potentially reduce performance for some workloads. Additional information can be found in [VMware Knowledge Base article 2055140](#).

Note: While iSCSI login timeout is considered a best practice, it is also consider a requirement, and therefore is always set to 60 seconds during configuration.

3 Configuring an iSCSI vSwitch for multipathing

The MEM utilizes the same multipathing vSwitch as the VMware native Round Robin multipathing. Therefore, if the ESXi hosts are already configured for Round Robin, this step can be omitted. However, if the ESXi hosts are configured for Fixed Path, remove the iSCSI initiator and associated vSwitch before continuing. The MEM setup.pl script provides powerful functionality and is easy to use compared to the manual configuration process.

Note: The instructions provided here refer to using the vMA, however, the syntax is the same whether the vMA or CLI is used. The optional parameters `--username` and `--password` have been excluded from the examples below for clarity. If not included and `vi-fastpass` is not configured, the user will be prompted to provide the username and password.

1. Connect to the vMA and change the directory to where the setup.pl was uploaded.
2. Create a multipathing iSCSI vSwitch by invoking the setup.pl in interactive mode.

```
setup.pl --configure --server=hostname
```

3.1 Interactive mode configuration

In this configuration mode, a series of questions are posed to the user, and default options presented when available. To accept a default, press **[Enter]**.

Note: The following example was generated using MEM 1.3 and vSphere 6.0, the output may differ slightly with other combinations.

```
setup.pl --configure --server=10.124.6.223
Do you wish to use a standard vSwitch or a vNetwork Distributed Switch
(vSwitch/vDS) [vSwitch]:
Found existing switches vSwitch0.
vSwitch Name [vSwitchISCSI]:
Which nics do you wish to use for iSCSI traffic? [vmnic1 vmnic2 vmnic3]: vmnic2
vmnic3
IP address for vmknic using nic vmnic2: 192.168.0.215
IP address for vmknic using nic vmnic3: 192.168.0.216
Netmask for all vmknics [255.255.255.0]:
What MTU do you wish to use for iSCSI vSwitches and vmknics? Before increasing
the MTU, verify the setting is supported by your NICs and network switches.
[1500]: 9000
What prefix should be used when creating VMKernel Portgroups? [iSCSI]:
The SW iSCSI initiator is not enabled, do you wish to enable it? [yes]:
What PS Group IP address would you like to add as a Send Target discovery
address (optional)?: 192.168.0.200
What CHAP user would you like to use to connect to volumes on this group
(optional)?: BLUEcluster
What CHAP secret would you like to use to connect to volumes on this group
(optional)?:
Configuring iSCSI networking with following settings:

    Using a standard vSwitch 'vSwitchISCSI'
```

Configuring an iSCSI vSwitch for multipathing

```
Using NICs 'vmnic2,vmnic3'  
Using IP addresses '192.168.0.215,192.168.0.216'  
Using netmask '255.255.255.0'  
Using MTU '9000'  
Using prefix 'iSCSI' for VMkernel Portgroups  
Using SW iSCSI initiator  
Enabling SW iSCSI initiator  
Adding PS Series Group IP '192.168.0.200' with CHAP user 'CHAPuser' to  
Send Targets discovery list.
```

The following command line can be used to perform this configuration:

```
setup.pl --configure --server=10.124.6.223 --vswitch=vSwitchISCSI --mtu=9000 --  
nics=vmnic2,vmnic3 --ips=192.168.0.215,192.168.0.216 --netmask=255.255.255.0 --  
vmkernel=iSCSI --nohwiscsi --enableswiscsi --groupip=192.168.0.200 --  
chapuser=BLUEcluster  
Do you wish to proceed with configuration? [yes]:
```

Note: The CHAP secret is not displayed as part of the command line that Interactive Mode generates.

As the script executes, it provides status notifications through several steps while creating the multipathing iSCSI vSwitch. Once complete, a multipathing iSCSI vSwitch will be created on the ESXi host, as shown in Figure 1.

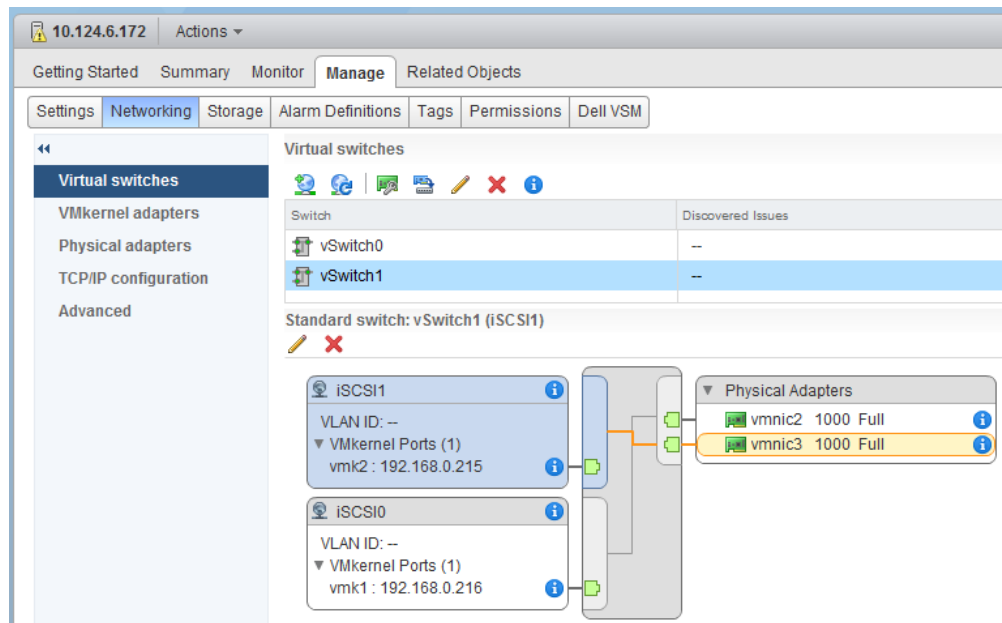


Figure 1 Multipathing iSCSI vSwitch created using MEM setup.pl script

A detailed list of all the parameters, options, and usage is provided in the *EqualLogic Multipathing Extension Module Installation and User Guide*. Additional configuration examples are provided in appendix A.

3.2 Unattended mode configuration

While setup.pl scripts interactive mode is extremely helpful, using it to configure all the hosts in a cluster is not efficient. Instead, passing the values of the various parameters to the setup.pl script provides a more efficient means of configuring several hosts.

At the end of the interactive mode configuration, the resulting command-line questions and answers are presented as shown in the example below.

```
setup.pl --configure --server=10.124.6.223 --vswitch=vSwitchISCSI
--mtu=9000 --nics=vmnic2,vmnic3 --ips=192.168.0.215,192.168.0.216
--netmask=255.255.255.0 --vmkernel=iSCSI
--nohwiscsi --enableswiscsi --groupip=192.168.0.200
--chapuser=BLUEcluster
```

Using this as an example, modify the hostname and IP information to provision additional ESXi hosts with the same vSwitch configuration.

```
setup.pl --configure --server=10.124.6.224 --vswitch=vSwitchISCSI
--mtu=9000 --nics=vmnic2,vmnic3 --ips=192.168.0.225,192.168.0.226
--netmask=255.255.255.0 --vmkernel=iSCSI
--nohwiscsi --enableswiscsi --groupip=192.168.0.200
--chapuser=BLUEcluster --chapsecret=BLUEsecret
```

Note: The CHAP secret is not displayed as part of the generated interactive mode command line, but was added to the example above. If it is not included, setup.pl will prompt for the CHAP secret at run time, if a CHAP user is specified.

Depending on the required configuration of the multipathing iSCSI vSwitch, additional parameters may be required or removed. For a full list of parameters and their defaults, see “Configuring Your Network for the MEM Plugin” in the *EqualLogic Multipathing Extension Module Installation and User Guide*.

4 Installing the MEM

The MEM can be installed using the command-line tools (CLI or vMA) or with vSphere Update Manager (VUM). The command-line tools are used in the example in this section. See appendix B for installation instructions using VUM.

Setup.pl is a Perl script wrapper around a number of VMware CLI commands. It provides administrators with a similar installation process, regardless of the vSphere version.

Note: Prior to installation, setup.pl copies the VIB to the first datastore it finds if a particular datastore has not been specified using the `--datastore` parameter. Due to API limitations, the VIB is not deleted by setup.pl.

Note: When performing an upgrade or uninstall of the MEM, a reboot is required on all versions of vSphere.

1. Using the vMA or vSphere CLI, install the MEM on each host with the command:

```
setup.pl --install --server=hostname
```

It may take a few minutes to install the module to the ESXi host.

2. Reboot the ESXi host to enable the plugin multipathing functionality. This may be done manually during the host configuration process or automatically after the MEM module is installed by appending the `--reboot` parameter to the above command line.
3. To verify the availability of the MEM, execute the command:

```
setup.pl --query --server=hostname
```

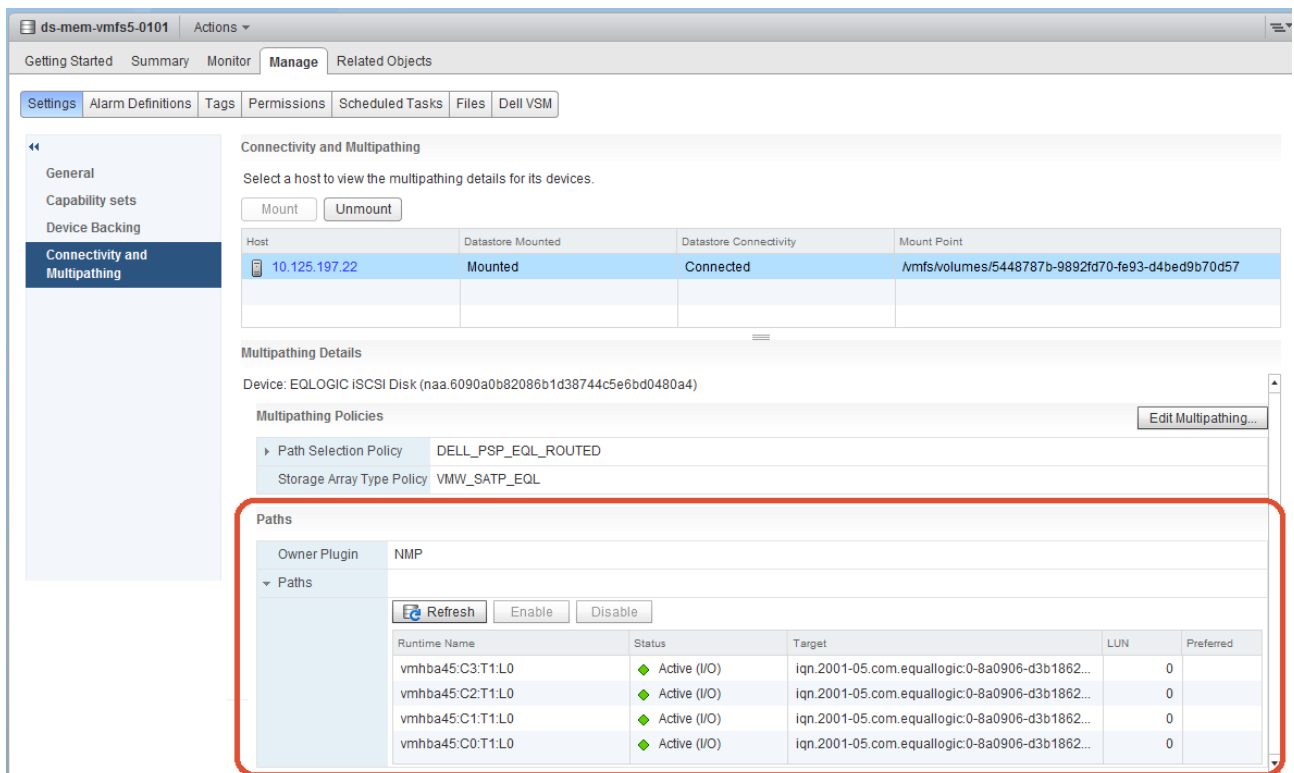
As shown in the following, running this command displays the presence and status of the MEM on the host and any volumes it is the active PSP for.

```
Found Dell EqualLogic Multipathing Extension Module installed: 1.5.0
Default PSP for EqualLogic devices is DELL_PSP_EQL_ROUTED.
Active PSP for naa.6090a088f03eebc07ee924f465696412 is
DELL_PSP_EQL_ROUTED.
```

5 Verification of MEM iSCSI session creation

By default, the MEM claims existing and new PS Series volumes and then creates the additional iSCSI session used by the MEM to route the iSCSI data. To verify these additional sessions, perform the following steps:

1. From the vSphere Web Client, click the storage icon, and on the vCenter object pane on the left, select a datastore.
2. For that datastore, select the **Manage** tab, and click **Settings**.
3. Select an individual host from the cluster where the datastore is mounted, and expand the paths.



The screenshot displays the vSphere Web Client interface for a datastore named 'ds-mem-vmfs5-0101'. The 'Manage' tab is selected, and the 'Connectivity and Multipathing' settings are visible. The 'Paths' section is highlighted with a red box, showing four active paths to the target volume.

Host	Datastore Mounted	Datastore Connectivity	Mount Point
10.125.197.22	Mounted	Connected	/vmfs/Volumes/5448787b-9892fd70-fe93-d4bed9b70d57

Multipathing Policies

- Path Selection Policy: DELL_PSP_EQL_ROUTED
- Storage Array Type Policy: VMW_SATP_EQL

Paths

Runtime Name	Status	Target	LUN	Preferred
vmhba45:C3:T1:L0	Active (I/O)	iqn.2001-05.com.equallogic:0-8a0906-d3b1862...	0	
vmhba45:C2:T1:L0	Active (I/O)	iqn.2001-05.com.equallogic:0-8a0906-d3b1862...	0	
vmhba45:C1:T1:L0	Active (I/O)	iqn.2001-05.com.equallogic:0-8a0906-d3b1862...	0	
vmhba45:C0:T1:L0	Active (I/O)	iqn.2001-05.com.equallogic:0-8a0906-d3b1862...	0	

Notice that the **Path Selection Policy** is set to `DELL_PSP_EQL_ROUTED`, and that there are multiple paths with active I/O to the target volume backing the datastore.

This can also be verified from the array web GUI on the **Connections** tab of the volume. The following example shows two iSCSI sessions have been established from each of the VMkernel ports on the ESXi host to the volume.

The screenshot displays the web interface for Volume skynet01. The top navigation bar includes tabs for Status, Access, Snapshots, Replication, Collections, Schedules, and Connections. The main content area is divided into three sections:

- Activities:** A sidebar menu for Volume skynet01 with options like Modify settings, Clone, and Move volume.
- Volume iSCSI Settings:** Shows the iSCSI target as `iqn.2001-05.com.equallogic:0-1cb196-f7432d738-f61bc3894c158ab6-skynet01` and the public alias as `skynet01`.
- iSCSI Connections:** A table showing 4 total connections. The table lists the Initiator address and Connection time for each session.

Initiator address	Connection time
10.10.12.238	38 day 16 hour 51 min
10.10.12.238	38 day 16 hour 51 min
10.10.12.233	38 day 16 hour 51 min
10.10.12.233	38 day 16 hour 51 min

6 Advanced iSCSI connection configuration parameters

The few configuration parameters that are needed with the MEM are used to alter the runtime behavior of the EqualLogic Host Connection Manager (EHCM) that manages the iSCSI sessions to the PS Series volumes. For information on all of the parameters, refer to the *EqualLogic Multipathing Extension Module Installation and User Guide*.

Three configuration parameters affect the number of iSCSI sessions created:

- *totalsessions*: The maximum total number of sessions that may be created on all PS Series volumes accessible to the host. The default value is 512.
- *volumesessions*: The maximum number of sessions that may be created on an individual PS Series volume. The default value is 6.
- *membersessions*: The maximum number of session that may be created on the portion of a volume that resides on a single member. The default value is 2.

When determining how many sessions to create, EHCM chooses a value that meets all these constraints. The default values provide good performance without unnecessarily consuming excess resources. In the majority of installations, these settings do not need to be changed from the defaults.

Prior to altering these values, it is strongly advised that an administrator fully understand the impact of such a change on the total number of iSCSI connections. For information on the number of iSCSI connections that particular firmware versions can handle at a per pool and group level, refer to the release notes for that firmware. Refer to the applicable VMware documentation for details on vSphere limits that may apply.

Section 2.2, Deployment considerations for iSCSI connection count, includes a formula for calculating the iSCSI connection count based on the number of hosts, sessions per volume portion, PS Series members in the pool, and the number of volumes the hosts are connected to. Take into account connections from other systems in the total iSCSI connection count such as:

- Other physical servers
- Virtual machines with iSCSI within the guest connections
- Backup servers
- Short-lived connections which are used by the vSphere VAAI Full Copy primitive and replication

6.1 Increasing the default values

In some lab settings, it may be necessary to increase the number of iSCSI connections for benchmarking I/O to an individual volume or to test the limits.

In such a configuration, it is possible that the default membersession value is limiting the number of iSCSI connections to fewer than the number of VMkernel ports bound to the iSCSI initiator. To utilize all the VMkernel ports bound to the iSCSI initiator, increase membersessions sufficiently to enable EHCM to create iSCSI connections through all bound VMkernel ports.

Generally, this increase would not be beneficial in a production environment. In a production environment, multiple volumes are used, and EHCM would, in aggregate, create and balance an optimal number of iSCSI sessions to effectively utilize all of the network ports.

6.2 Decreasing the default values

In large configurations, reduce the number of iSCSI connections if the limits of the array firmware are exceeded. This can be achieved by:

- Reducing the number of volumes: By increasing the size of a given volume, more virtual machines can reside on it, thereby reducing the total number of volumes needed.
- Reconfiguring larger clusters into multiple smaller clusters: By splitting a large cluster into smaller clusters, the number of hosts accessing particular volumes is reduced, thereby reducing the number of iSCSI connections consumed.
- Altering the configuration of the PS Series group: Reducing the number of PS Series members per pool, and increasing the number of pools, will increase the number of available iSCSI connections.

These changes typically take some time to implement. Altering the EHCM parameters provides a means of quickly reducing the iSCSI connection to be within limits. However, depending on the overall environment configuration and I/O workload, a drop in overall performance may be observed.

- *totalsessions*: This parameter is defaulted to a maximum of 512, and can be lowered to 64. It is the maximum number of iSCSI connections that EHCM will create from this individual ESXi host to all PS Series arrays. If an additional volume is added to the host, EHCM automatically rebalances the sessions to include the new and existing volumes, without consuming more than the totalsessions limit.
- *volumesessions*: This parameter is defaulted to a maximum of 6, and can be lowered to 1. It is the maximum number of iSCSI connections that EHCM will make from this individual ESXi host to each individual PS Series volume, regardless of the number of array members the volume may span as the group expands. This value applies to all volumes that the host accesses, and cannot be specified to affect only particular volumes. As additional members are added to the pool, and volumes are distributed to include this new member, EHCM creates additional sessions to the volume within this limit.
- *membersessions*: This parameter is defaulted to a maximum of 2, and can be lowered to 1. It is the number of iSCSI connections that EHCM makes from this individual ESXi host to the portion of a volume that resides on an individual PS Series member. Setting membersessions to 1 will halve the iSCSI connection count from the ESXi host to the array. However, this effectively disables the Least Queue Depth part of the MPIO algorithm, as there is one queue to each portion of the volume. The Intelligent Routing part of the algorithm remains in use, so I/O requests will be efficiently directed to the appropriate member.

Note: Similar default values for VMware® vSphere® Virtual Volumes™ (vVols) are defined in ehcmd.conf. See the product documentation for details.

6.3 Setting the EHCM configuration values under vSphere

PS Series arrays have added a number of commands that can be used to control and report on the status of the MEM. It is through these commands that the EHCM parameters are changed.

The `esxcli` command can be executed remotely from the vMA, CLI, or directly on the ESXi host. The syntax of this command is:

```
esxcli equallogic param set --name=parameter_name
--value=parameter_value
```


For example:

```
esxcli equallogic param set --name=membersession --value=1
```

To query the current EHCM values, use the following syntax:

```
esxcli equallogic param list
```

7 Using MEM with Virtual Volumes

VMware introduced the Virtual Volumes (vVols) feature with vSphere 6.0. It enables per-VM granularity for many storage-related tasks, as opposed to per-volume granularity. While vVols do not require MEM, the inclusion of MEM in a PS Series vVol environment provides more efficient routing of I/O, reduced latency, and increased bandwidth utilization. For additional information on Virtual Volumes, see [VMware vSphere Virtual Volumes on Dell PS Series](#).

7.1 MEM configuration and installation

While vVol changes many things when it comes to architecting a vSphere environment, it does not alter the creation of the iSCSI vSwitch for multipathing, or the installation process of the MEM.VIB.

7.2 Access controls

One aspect that has changed with vVols is access controls. With traditional volumes, access to a volume is managed through an access control list or access control policy, on a per-volume basis. However, with vVols volumes, this no longer applies. Access is restricted by an access control list on the protocol endpoint, and managed from the PS Series Virtual Storage Manager (VSM) plugin for vCenter. For detailed instructions on vVol access controls, see the document, [Dell Virtual Storage Manager: Installation Considerations and Local Data Protection](#).

7.3 Virtual Volumes impact on iSCSI connection count

With traditional VMFS datastores, the iSCSI connection terminated at the individual volume, but with vVols, this has changed. The iSCSI connection now terminates at the protocol endpoint, resulting in a lowering of the iSCSI connection count, regardless of the number of storage containers or the number vVol-based virtual machines. With MEM and vVols, the iSCSI connection count is now simply one iSCSI connection per VMkernel/PS Series member pair, with a minimum of at least two iSCSI connections from each host to the ESXi host to the protocol endpoint. The formula can be expressed as:

Number of VMkernel ports used for iSCSI x Number of members in group x two sessions per member

For example, if there are three PS6210 arrays, and each host is configured with two VMkernel ports for iSCSI, then six iSCSI connections will be created from each host.

As with all rules, there are exceptions. With the PS6110, PS4110, and PS-M4110 arrays, if they are in a single-member pool, MEM will create two iSCSI connections to the protocol endpoint for redundancy. With two or more such members in the pool, MEM will create a single iSCSI connection per member, resulting in one iSCSI connection being created per member to the protocol endpoint.

When mixing traditional volumes and virtual volumes, the iSCSI connection pool limit applies to the combined total of iSCSI connections from traditional volumes and the iSCSI connection to the protocol endpoint.

8 Summary

Through its intelligent routing and load balancing, the MEM can provide a reduction in network latency and an increase in bandwidth to PS Series storage arrays. Through automated setup and host connection management, the MEM also reduces the steps for deployment and ongoing management of advanced vSphere iSCSI configurations.

A Example iSCSI vSwitch configurations

This appendix provides some examples of the additional parameters used with the `setup.pl` configuration script. A detailed list of all the parameters and their usage can be found in the *EqualLogic Multipathing Extension Module Installation and User Guide*.

A.1 Overriding hardware iSCSI offload default utilization

ESXi can utilize the iSCSI offload capabilities of the Broadcom® NetXtreme II network adapters, resulting in significantly lower software iSCSI CPU utilization. By default, the `setup.pl` configuration script will utilize this iSCSI offload capability if it is present. If there is a preference not to use the iSCSI offload capability, this must be specified when configuring the iSCSI vSwitch. This is shown as follows:

```
setup.pl --configure --server=172.17.5.121 --nics=vmnic2,vmnic3  
--ips=192.168.0.215,192.168.0.216 --nohwiscsi
```

A.2 Utilizing a vNetwork distributed switch

With Enterprise Plus licensing, VMware provides a virtual switch which spans many ESXi hosts. This abstracts the configuration of individual vSwitches on the host level and enables centralized management through vSphere vCenter server. If utilizing a vNetwork Distributed Switch for iSCSI traffic, the `--vds` parameter must be specified. Should the name of the vNetwork Distributed Switch differ from the default utilized by the configuration script, it can be specified using the `--vswitch` parameter. This is shown as follows:

```
setup.pl --configure --server=172.17.5.121 --nics=vmnic2,vmnic3  
--ips=192.168.0.215,192.168.0.216 --vswitch vdsISCSI --vds
```

A.3 Enabling Jumbo Frames

With the vSphere release of ESXi support for Jumbo Frames has been extended to VMkernel traffic, which includes the iSCSI stack. To utilize Jumbo Frames, they must be enabled on all networking components used. By default, the `setup.pl` script uses an MTU of 1500 when creating the iSCSI vSwitch and VMkernel interfaces. To use a larger value, specify it as demonstrated in the following example:

```
setup.pl --configure --server=172.17.5.121 --nics=vmnic2,vmnic3  
--ips=192.168.0.215,192.168.0.216 --mtu=9000
```

A.4 Setting the iSCSI discover address

It is possible to specify the group IP of an array to be set as the send targets discovery address for the iSCSI initiator.

```
setup.pl --configure --server=172.17.5.121 --nics=vmnic2,vmnic3  
--ips=192.168.0.215,192.168.0.216 --groupip=192.168.0.200
```

B Installing the MEM with VMware Update Manager

Note: vCenter Update Manager is only capable of installing the MEM. vMA or vCLI will need to be utilized to configure the MPIO iSCSI vSwitch for the MEM.

VMware Update Manager (VUM) has the ability to install and upgrade third-party packages to the ESXi hosts. This enables administrators to not only manage the patching and updating of ESXi hosts but also to update third party packages installed on the ESXi hosts.

Installing the MEM consists of four major steps. Perform steps 1-3 once per version release of the MEM:

1. Import the MEM to the patch repository.
2. Create an extension baseline that includes the MEM.
3. Attach the baseline to the environment.

Perform step 4 for each host in the cluster or data center.

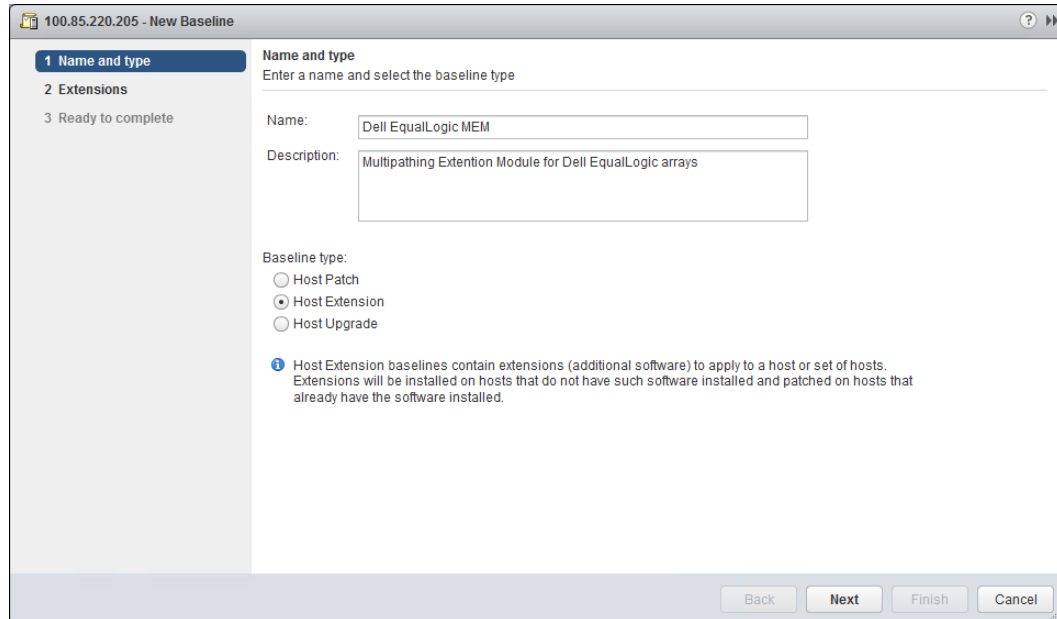
4. Install the MEM to each host.

B.1 Step 1: Importing the MEM to the patch repository

1. Download the MEM from eqsupport.dell.com (login required).
2. Unpack the zip archive. Do not unpack the embedded zip file, which is a VIB offline bundle and is expected to be in this format.
3. From the vSphere client **Home** section, under **Operations and Policies** select **Update Manager**. Select the vCenter to be administrated, and click the **Manage** tab.
4. Select **Patch Repository** and click **Import Patches** to start the Import Patches wizard.
5. Click **Browse...**, browse to where the zip archive was unpacked, select the embedded zip file named **dell-eql-mem-*<versionnumber>*.zip** and click **Open**.
6. The Import Patches wizard will upload and analyze the file. This may take a few minutes. When completed, click **Finish**.

B.2 Step 2: Creating an extension baseline

1. Continuing from the **Manage** tab, select **Host Baselines**, and click **New Baseline...** to start the New Baseline wizard.
2. Provide the baseline with a suitable name and optional description, select the baseline type **Host Extension**, and click **Next** shown in the following example:



3. From the list of extensions, select **Dell EqualLogic iSCSI MEM 1.x.x**, and click **Next**. If there are a large number of extensions in the repository, use **EqualLogic** as a keyword for filtering.
4. On the Ready to Complete page of the wizard, verify that the information is correct, and click **Finish**.

B.3 Step 3: Attaching the extension baseline

1. Continuing from the **Manage** tab, on the far right click the **Go to compliance view** button.
2. From the tree pane on the left, select the vCenter, Datacenter, Cluster, or Host vSphere object that the MEM Extension Baseline will attach to, and click **Attach Baseline...** to start the Attach Baseline wizard. Typically the baseline will be attached as at the top level, unless circumstances dictate otherwise. It may be necessary to re-select the **Update Manager** tab after selecting the vSphere object.
3. Select the MEM extension baseline created previously and click **OK** to attach the Extension Baseline.

B.4 Step 4: Installing the MEM

1. From the vSphere client, select the vSphere object to which MEM is to be installed to, and select the **Update Manager**.
2. In the **Baselines** pane, under **Independent baselines**, highlight the MEM extension baseline that was attached earlier and click the **Remediate...** button to start the Remediation wizard.
3. On the **Select baselines** page, verify that the correct baseline is listed. Click **Next** to continue.
4. On the **Select target** objects, unselect any hosts to which the MEM is not to be installed to at this time. Click **Next** to continue.

5. The **Patches and extensions** page lists the patches and extensions to be applied. Click **Next** to continue.
6. On the **Advanced options** page, there is the option to schedule the deployment for a future time and date. Click **Next** to continue.
7. On the **Host remediation options** page, there are options for altering the behaviors of the virtual machines during the install. Refer to the *Installing and Administering VMware vSphere Update Manager* from the [VMware Documentation](#) site for details. Click **Next** to continue.
8. On the **Ready to complete** page, verify that the information is correct, and click **Finish**. Unless the task was schedule to run at a later time, it will be immediately executed.
9. As the task is been executed, it will display status updates in the Recent Tasks pane of the vSphere client. If several host were selected, the task will execute against one host at a time, until the MEM has been installed on all the selected hosts.
10. Once the task has been completed, the MEM will have been installed or updated on all the hosts selected.

C Using the setup.pl script to configure Round Robin MPIO

MEM is dependent upon the vStorage APIs for Multipathing which are only available with the appropriate licensing. For vSphere environments with licensing which do not include these APIs, it is recommended that the vSphere native Round Robin Multipathing be used. For additional information about Round Robin Multipathing, see the document, [Best Practices for Implementing VMware vSphere in a Dell PS Series Storage Environment](#).

C.1 Configuring ESXi for Round Robin MPIO

Configuring an ESXi host for native Round Robin requires three steps:

1. Create an MPIO-capable vSwitch.
2. Set the PSP claim rules for PS Series volumes to Round Robin.
3. Optimize the number of IOs per path.

The first step is easily accomplished as both the MEM and native Round Robin utilize the same iSCSI vSwitch configuration. See section 3 of this document for detailed instructions on configuring an iSCSI vSwitch for multipathing. However, this will set the PSP (Path Selection Plugin) claim rules for PS Series volumes to utilize the MEM, which is not desired.

Steps two and three can be manually completed by following the steps outlined in the document, *Best Practices for Implementing VMware vSphere in a Dell PS Series Storage Environment*, or by utilizing a new feature in the setup.pl script. Running the setup.pl script with the parameters below will disable the MEM default routing of DELL_PSP_EQL_ROUTED and set the routing to VMWARE_PSP_RR (Round Robin). It will also change the IOs per path from 1000 to 3.

```
setup.pl --configure --disable --RR
```


D Software and firmware versions

Table 1 lists the software and firmware used for the preparation of this document.

Table 1 Software and firmware versions

Vendor	Model	Software revision
VMware	vSphere ESXi 6.0	6.0 and 6.5
Dell	Dell PS Series SAN	6.0 or above
Dell	EqualLogic Multipathing Extension Module for VMware vSphere 5.x	1.4
Dell	EqualLogic Multipathing Extension Module for VMware vSphere 6.x	1.5

E Technical support and resources

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

[Storage technical documents and videos](#) on provide expertise that helps to ensure customer success on Dell EMC storage platforms.

E.1 Related resources

Table 2 lists the referenced or recommended resources related to this document. The VMware documents listed in this table are available at the [VMware Documentation](#) site.

Table 2 Referenced or recommended resources

Vendor	Document title
Dell	Dell EqualLogic Group Manager Administrator's Manual at eglsupport.dell.com
Dell	EqualLogic Multipathing Extension Module Installation and User Guide
Dell	PS Series Firmware Compatibility with EqualLogic Products
VMware	vSphere Storage Guide
VMware	vSphere Installation and Setup Guide
VMware	Installing and Administering VMware vSphere Update Manager
VMware	vSphere Command-Line Interface Concepts and Examples
VMware	vSphere Management Assistant Guide