

Best Practices for Sharing an iSCSI SAN Infrastructure with Dell PS Series and SC Series Storage using VMware vSphere Hosts

Dell Storage Engineering
January 2017

Revisions

Date	Description
March 2015	Initial release
April 2015	Added specific iSCSI NIC optimization settings for shared host
October 2015	Updated for VMware ESXi 6.0 and added dedicated host information
January 2017	Updated for VMware ESXi 6.5

Acknowledgements

This paper was produced by the following members of the Dell Storage team:

Engineering: Chuck Armstrong

Editing: Camille Daily

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 © 2015 - 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 [1/30/2017] [Best Practices] [2015-A-BP-INF]

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
Acknowledgements.....	2
1 Introduction.....	5
1.1 Scope.....	5
1.2 Audience.....	5
1.3 Terminology.....	6
2 Storage product overview.....	8
2.1 PS Series storage.....	8
2.2 SC Series storage.....	8
3 PS Series and SC Series iSCSI SAN coexistence.....	9
3.1 Topology of a shared iSCSI SAN infrastructure with shared hosts.....	9
3.2 Topology of a shared iSCSI SAN infrastructure with dedicated hosts.....	10
3.3 PS Series specific settings.....	10
3.4 SC Series specific settings.....	11
3.4.1 SC Series host physical connectivity and IP assignment.....	11
4 Enabling vSphere 6.5 host access to PS Series and SC Series iSCSI storage – shared.....	13
4.1 Configure access to PS Series storage.....	13
4.1.1 PS Series Multipathing Extension Module (MEM) for VMware.....	13
4.2 Configure access to SC Series storage.....	13
4.2.1 Configuring the VMware iSCSI software initiator.....	13
4.2.2 Configuring the VMware iSCSI software initiator to access SC Series volumes.....	14
4.2.3 VMware native multipathing.....	16
4.2.4 Setting Path Selection Policy and storage performance settings - PS Series.....	16
4.2.5 Setting Path Selection Policy and storage performance settings - SC Series.....	16
5 Enabling vSphere 6.5 host access to PS Series and SC Series iSCSI storage – dedicated.....	17
5.1 Configure access to PS Series storage.....	17
5.1.1 PS Series Multipathing Extension Module (MEM) for VMware.....	17
5.2 Configure access to SC Series storage.....	17
5.2.1 Configuring the VMware iSCSI software initiator to access SC Series volumes.....	17
5.2.2 VMware native multipathing.....	24
5.2.3 Setting Path Selection Policy and storage performance settings - SC Series.....	24
6 Test methodology.....	25

6.1	Test environment.....	26
6.2	I/O performance testing.....	26
6.2.1	I/O performance results and analysis: shared hosts.....	26
6.2.2	I/O performance results and analysis: dedicated hosts.....	28
6.3	High availability testing.....	28
7	Best practice recommendations.....	30
7.1	Switch fabric.....	30
7.2	Host connectivity.....	30
7.3	Storage.....	30
8	Conclusion.....	31
A	Technical support and resources.....	32
A.1	Related documentation.....	32

1 Introduction

Dell™ PS Series and Dell EMC™ SC Series storage systems both support storage area networks (SANs) over the iSCSI protocol. This document provides best practices for deploying:

- VMware® vSphere™ host servers connected to an existing PS Series storage target to simultaneously connect to an SC Series storage target over a shared iSCSI SAN infrastructure (shared)
- VMware vSphere host servers with both PS Series and SC Series storage targets, where only the iSCSI SAN infrastructure is shared: each host connects to either PS Series or SC Series storage targets (dedicated)

This paper also provides analysis of performance and high availability of the shared iSCSI SAN infrastructure consisting of PS Series and SC Series arrays.

1.1 Scope

The scope of this paper focuses on the following:

- Dedicated switches for iSCSI storage traffic
- Non-DCB (Data Center Bridging) enabled iSCSI SAN
- Standard TCP/IP implementations utilizing standard network interface cards (NICs)
- VMware vSphere ESXi operating-system-provided software iSCSI initiator
- Virtual LAN (VLAN) untagged solution
- IPv4 only for PS Series and SC Series

The scope of this paper does **not** include the following:

- 1GbE or mixed-speed iSCSI SAN (combination of 1GbE and 10GbE)
- DCB or sharing the same SAN infrastructure for multiple traffic types
- iSCSI offload engine (iSOE)
- NIC partitioning (NPAR)
- VLAN tagging at the switch, initiator, or target
- SC Series storage systems using Fibre Channel over Ethernet (FCoE) SAN connectivity
- Non-MPIO (Multipath Input/Output) implementation

1.2 Audience

This paper is for storage administrators, network administrators, SAN system designers, storage consultants, or anyone tasked with configuring a SAN infrastructure for PS Series and SC Series storage. It is assumed that readers have experience in designing and/or administering a shared storage solution. There are assumptions made in terms of familiarity with all current Ethernet standards as defined by the Institute of Electrical and Electronic Engineers (IEEE) as well as TCP/IP standards defined by the Internet Engineering Task Force (IETF) and FC standards defined by the T11 committee and the International Committee for Information Technology Standards (INCITS).

1.3 Terminology

The following terms are used throughout this document:

Converged network adapter (CNA): A network adapter that supports convergence of simultaneous communication of both traditional Ethernet and TCP/IP protocols as well as storage networking protocols such as internet SCSI (iSCSI) or Fibre Channel over Ethernet (FCoE) using the same physical network interface port.

Data Center Bridging (DCB): A set of enhancements made to the IEEE 802.1 bridge specifications for supporting multiple protocols and applications in the same data center switching fabric. It is made up of several IEEE standards including Enhanced Transmission Selection (ETS), Priority-based Flow Control (PFC), Data Center Bridging Exchange (DCBX), and application Type-Length-Value (TLV). For more information, see the document, [Data Center Bridging: Standards, Behavioral Requirements, and Configuration Guidelines with Dell EqualLogic iSCSI SANs](#).

EqualLogic Multipathing Extension Module (MEM) for VMware vSphere: The PS Series multipath I/O (MPIO) module for vSphere.

Fault domain (FD): A set of hardware components that share a single point of failure. For controller-level redundancy, fault domains are created for SC Series storage to maintain connectivity in the event of a controller failure. In a dual-switch topology, each switch acts as a fault domain with a separate subnet and VLAN. Failure of any component in an FD will not impact the other FD.

iSCSI offload engine (iSOE): Technology that can free processor cores and memory resources to increase I/Os per second (IOPS) and reduce processor utilization.

NIC partitioning (NPAR): A technology used by Broadcom and QLogic which enables traffic on a network interface card (NIC) to be split into multiple partitions. NPAR is similar to QoS on the network layer and is usually implemented with 10GbE.

Link aggregation group (LAG): A group of Ethernet switch ports configured to act as a single high-bandwidth connection to another switch. Unlike a stack, each individual switch must still be administered separately and function independently.

Local area network (LAN): A network carrying traditional IP-based client communications.

Logical unit (LUN): A number identifying a logical device, usually a volume that is presented by an iSCSI or Fibre Channel storage controller.

Multipath I/O (MPIO): A host-based software layer that manages multiple paths for load balancing and redundancy in a storage environment.

Native VLAN and default VLAN: The default VLAN for a packet that is not tagged with a specific VLAN or has a VLAN ID of 0 or 1. When a VLAN is not specifically configured, the switch default VLAN will be utilized as the native VLAN.

Network interface card (NIC): A network interface card or network interface controller is an expansion board inserted into the computer/server so that the computer/server can connect to a network. Most NICs are designed for a particular type of network (typically Ethernet) protocol (typically TCP/IP) and media.

Storage area network (SAN): A Fibre Channel, Ethernet, or other specialized network infrastructure specifically designed to carry block-based traffic between one or more servers to one or more storage and storage inter-process communications systems.

Virtual LAN (VLAN): A method of virtualizing a LAN to make it appear as an isolated physical network. VLANs can reduce the size of and isolate broadcast domains. VLANs still share resources from the same physical switch and do not provide any additional Quality of Service (QoS) services such as minimum bandwidth, quality of a transmission, or guaranteed delivery.

VLAN tag: IEEE 802.1Q: The networking standard that supports VLANs on an Ethernet network. This standard defines a system of tagging for Ethernet frames and the accompanying procedures to be used by bridges and switches in handling such frames. Portions of the network which are VLAN-aware (IEEE 802.1Q conformant) can include VLAN tags. When a frame enters the VLAN-aware portion of the network, a tag is added to represent the VLAN membership of the frame's port or the port/protocol combination. Each frame must be distinguishable as being within exactly one VLAN. A frame in the VLAN-aware portion of the network that does not contain a VLAN tag is assumed to be flowing on the native (or default) VLAN.

2 Storage product overview

The following sections provide an overview of the Dell storage products and technologies presented in this paper.

2.1 PS Series storage

PS Series arrays deliver the benefits of consolidated networked storage in a self-managing iSCSI SAN that is affordable and easy to use, regardless of scale. Built on an advanced, peer storage architecture, PS Series storage simplifies the deployment and administration of consolidated storage environments, enabling perpetual self-optimization with automated load balancing across PS Series members in a pool. This provides efficient scalability for both performance and capacity without forklift upgrades. PS Series storage provides a powerful, intelligent and simplified management interface.

2.2 SC Series storage

SC Series storage is the Dell EMC enterprise storage solution featuring multi-protocol support and self-optimizing, tiering capabilities. SC Series storage can be configured with all flash, as a hybrid system, or with only traditional spinning disks and features automatic migration of data to the most cost-effective storage tier. Efficient thin provisioning and storage virtualization enable disk capacity usage only when data is actually written, enabling a pay-as-you-grow architecture. This self-optimizing system can reduce overhead cost and free up the administrator for other important tasks.

3 PS Series and SC Series iSCSI SAN coexistence

PS Series and SC Series arrays can coexist in a shared iSCSI SAN, either with shared hosts or dedicated hosts. Shared hosts not only share the iSCSI SAN infrastructure, but also connect to storage targets on both the PS Series and SC Series arrays. Shared-host coexistence (see Figure 1) shares the iSCSI SAN infrastructure and has all hosts connected to both array platforms. When hosts are dedicated (see Figure 2), each host in the iSCSI infrastructure connects to targets from either the PS Series array or SC Series array, but not both. Dedicated host coexistence utilizes a shared iSCSI SAN infrastructure only.

3.1 Topology of a shared iSCSI SAN infrastructure with shared hosts




Figure 1 Shared iSCSI SAN with shared hosts reference topology

3.2 Topology of a shared iSCSI SAN infrastructure with dedicated hosts




Figure 2 Shared iSCSI SAN with dedicated hosts reference topology

3.3 PS Series specific settings

The use cases defined in this paper consist of SC Series arrays and connected VMware vSphere hosts sharing only the Ethernet iSCSI SAN infrastructure with existing PS Series storage and its connected hosts (dedicated), as well as SC Series arrays sharing not only the iSCSI SAN infrastructure, but also the VMware vSphere hosts (shared). It is assumed that the Ethernet network supporting the iSCSI SAN, as well as the VMware vSphere hosts accessing PS Series storage, are configured using best practice recommendations as defined in the documents, [Best Practices for Implementing VMware vSphere in a Dell PS Series Storage Environment](#).

The PS Series Virtual Storage Manager (VSM), Multipath Extension Module (MEM), and iSCSI port binding best practice settings for PS Series storage are applicable when the Ethernet iSCSI SAN network is shared with SC Series arrays.

Note: Additional PS Series-specific information can be found in [ESXi Versions 5.1, 5.5, or 6.0 Host Configuration Guide](#) and [Dell PS Series Configuration Guide](#).

3.4 SC Series specific settings

A typical SC Series iSCSI implementation involves two separate, dedicated Ethernet fabrics as two fault domains with an independent IP subnet and unique, non-default VLANs in each switch fabric. However, to enable SC Series to coexist with PS Series and share the Ethernet SAN infrastructure using the iSCSI storage protocol, use a single subnet for all host and storage ports.

To implement this correctly, a basic understanding of PS Series and SC Series storage is needed. This paper provides an overview of both storage types.

Each PS Series volume is presented as a unique target with LUN 0. The PS Series volumes that are accessible to the host are listed in the iSCSI initiator properties. When a volume is connected, the iSCSI initiator establishes the initial iSCSI session and then the PS Series MPIO plugin determines if additional sessions are necessary for redundancy.

Each SC Series array has both front-end and back-end ports. The front-end ports are presented with unique target LUN IDs. Every initiator IP has a connection to each port that it can access. Redundant connections are made by creating multiple sessions with each of the virtual iSCSI ports of the SC Series storage system. For example, one initiator port and two target ports in each fault domain means there will be four connections (two for each fault domain).

Note that the host port on one fault domain can access the target port on the other fault domain through the switch interconnection. Ensure that the host ports are connected to the appropriate fault domain and target port, physically and in iSCSI sessions. This minimizes inter-switch link (ISL) traffic and ensures that at least some iSCSI sessions will persist in the event that a component fails in a fault domain. The following sections discuss ways to ensure that the physical connectivity and iSCSI sessions are established correctly.

Note: Additional SC Series-specific information can be found in [Dell EMC SC Series Best Practices with VMware vSphere 5.x-6.x](#).

3.4.1 SC Series host physical connectivity and IP assignment

Depending on the OS-specific implementation, different methods are used to connect the arrays and assign IP addresses. Since SC Series fault domains are connected by an ISL and are in a single IP subnet, it is important to ensure that iSCSI sessions are properly established within their fault domains. Host ports connected to Fault Domain 1 should connect to switch fabric and storage ports on Fault Domain 1 physically. The same rule applies for Fault Domain 2. This step is important because, with a single subnet, it is possible for the hosts to access SC Series storage ports on both fault domains. The correct connectivity minimizes ISL traffic and ensures that at least some iSCSI sessions will persist in the event of a component failure.

Figure 3 depicts proper connection from each host port to the SC Series storage ports within the same fault domain without traversing the switch interconnection.




Figure 3 Connecting the host to SC Series ports

Note: With the approach discussed in this paper, misconfiguration of the SC Series connectivity (for example, host ports not connected to the correct fault domain) can lead to loss of volume access in the event of a switch failure.

4 Enabling vSphere 6.5 host access to PS Series and SC Series iSCSI storage – shared

This section assumes the environment is historically a PS Series storage environment and SC Series storage is being implemented into the existing environment, as shown in section 3.1

4.1 Configure access to PS Series storage

This section covers configuring access to PS Series storage through the installation and assumes VMware licensing is Enterprise or Enterprise Plus and supports the use of MEM.

4.1.1 PS Series Multipathing Extension Module (MEM) for VMware

The VMware vSphere host servers were configured using the best practices defined in the [Dell PS Series Configuration Guide](#), including the installation of MEM on the vSphere host servers.

For more information on MEM, see [Configuring and Installing the PS Series Multipathing Extension Module for VMware vSphere and PS Series SANs](#).

Note: MEM is only supported when using VMware vSphere Standard or above licensing. If MEM is not supported, manual configuration steps can be found in [Configuring iSCSI Connectivity with VMware vSphere 6 and Dell PS Series Storage](#).

After preparing the vSphere host, map volumes using Dell Storage manager or Group Manager.

Note: See the Dell Storage PS Series [Group Administration Guide](#) for information on creating volumes and mapping volumes to hosts (for existing customers; requires a valid portal account).

4.2 Configure access to SC Series storage

Configuring vSphere hosts to access SC Series storage when those hosts are already configured to access PS Series storage requires fewer steps as compared to when the host is not configured to access PS Series storage.

4.2.1 Configuring the VMware iSCSI software initiator

The first step to configure the vSphere hosts for SC Series arrays using a single subnet and iSCSI port binding is to create VMkernel (vmk) ports. For port binding to work correctly, the initiator must be able to reach the target. With the release of vSphere 6.5, routing is supported when using port binding. However, prior to vSphere 6.5, iSCSI port binding did not support routing: requiring the initiator and target be in the same subnet. VMware recommends associating each VMkernel port to a single vmnic uplink.

In this case, the NICs, vSwitches, and iSCSI software initiator have already been configured to access the PS Series storage and do not need additional configuration.

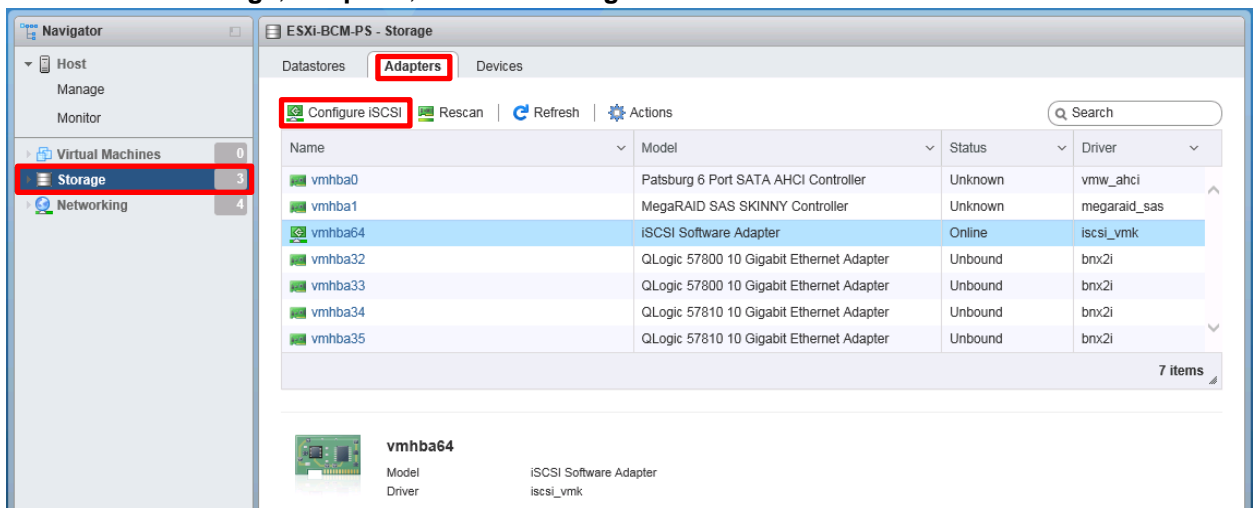
4.2.2 Configuring the VMware iSCSI software initiator to access SC Series volumes

For SC Series storage, the discovery portal addresses must be added manually to the dynamic discovery tab in the VMware iSCSI software initiator.

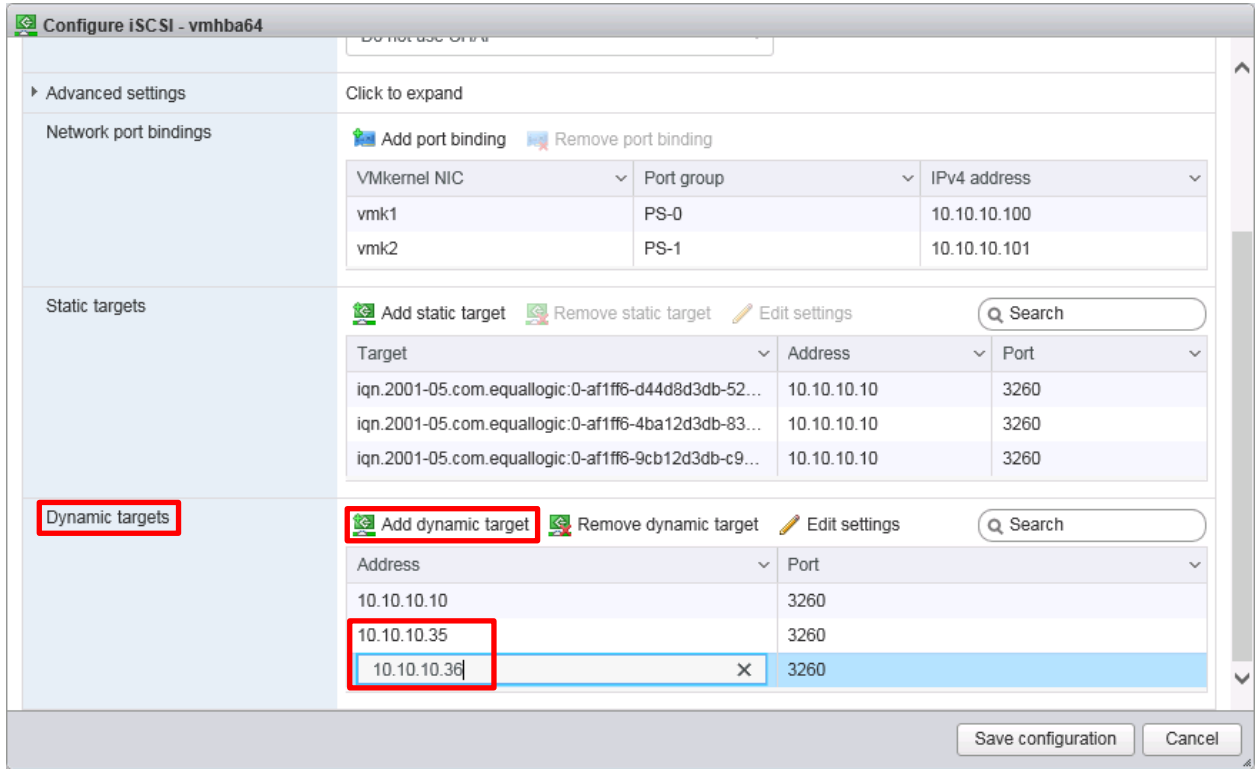
SC Series storage presents front-end target ports and each volume is presented as a unique LUN. Redundant connections are made by creating multiple sessions with each of the virtual iSCSI target ports the array.

For the purpose of this discussion, 10.10.10.35 and 10.10.10.36 are used as the SC Series iSCSI SAN discovery addresses.

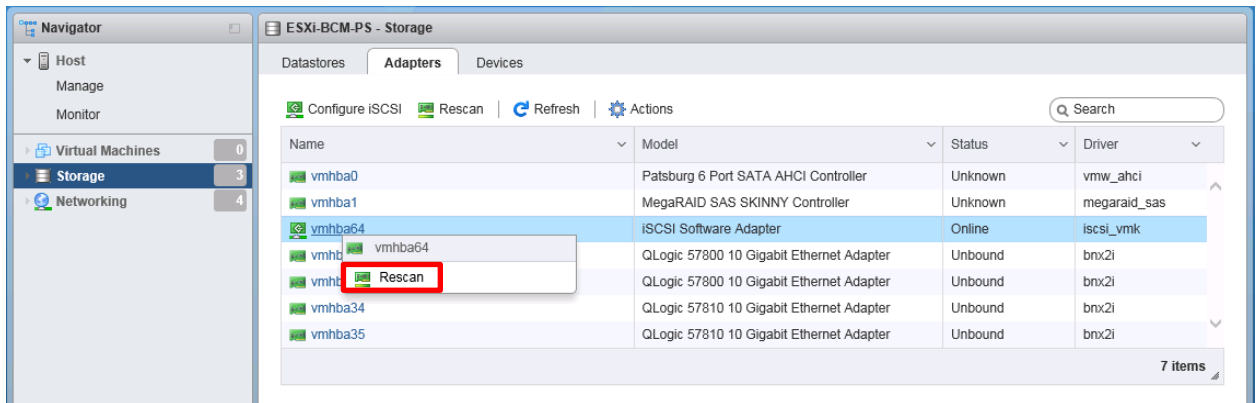
1. From the VMware ESXi Web GUI (local host management utility that has replaced the C# client), select **Storage, Adapters**, and click **Configure iSCSI**.



2. In the **Dynamic targets** section, click **Add dynamic target** twice and enter the two SC Series iSCSI target addresses and click **Save configuration**.



3. Right-click the **iSCSI Software Adapter** and click **Rescan** to connect to the targets.



Note: This approach with VMware does not provide a way to restrict initiators to connect only to targets on the same fault domain.

After preparing the vSphere host, create a server or cluster object in Dell Storage Manager and map volumes.

Note: See the [Create a Cluster Object in Enterprise Manager](#) and [Creating and mapping a volume in Enterprise Manager](#) videos for additional information.

4.2.3 VMware native multipathing

VMware provides native multipathing that can be used with any storage platform, if a vendor-provided multipath solution is not available.

4.2.4 Setting Path Selection Policy and storage performance settings - PS Series

In a shared environment, where VMware hosts connect to both PS Series and SC Series storage targets, the multipath and iSCSI settings optimized by the MEM (described previously) for all PS Series volumes. Therefore, no modifications are required

4.2.5 Setting Path Selection Policy and storage performance settings - SC Series

In a shared environment, where VMware hosts connect to both PS Series and SC Series storage targets, use VMware's native multipathing for SC Series volumes. Setting the Path Selection Policy (PSP) can be performed using either the vCenter GUI or a command line utility. Setting the PSP Default, and performance aspects of the PSP can only be performed using a command line utility at this time. Command line utility options are: ESXi console, SSH session, vSphere CLI, or VMware PowerCLI. The SSH session commands are shown here:

Note: At the time of this writing, the local vSphere host GUI, the new web-based utility that has replaced the C# client, does not have the ability to set PSP for volumes.

To set the Path Selection policy as default for new SC Series (SCOS v6.6 and up) volumes:

```
[root@ESXi-BCM-PS:~] esxcli storage nmp satp set -P VMW_PSP_RR -s VMW_SATP_ALUA
```

To set or change the PSP for an existing SC Series volume:

```
[root@ESXi-BCM-PS:~] esxcli storage nmp device set --device=naa.xxx --psp=VMW_PSP_RR
```

To set the Round Robin IOPS policy for an existing SC Series volume:

```
[root@ESXi-BCM-PS:~] esxcli storage nmp psp roundrobin deviceconfig set --device=naa.xxx/  
--type=iops --iops=3
```

Note: These settings must be applied manually for each new volume presented to the host. Or, as described in [Dell Storage SC Series Best Practices with VMware vSphere 5.x-6.x](#), default settings can be modified.

5 Enabling vSphere 6.5 host access to PS Series and SC Series iSCSI storage – dedicated

This section assumes the environment is historically a PS Series storage environment and SC Series storage with new hosts dedicated to the SC Series storage is being introduced into the environment, where only the iSCSI SAN infrastructure is being shared, as shown in section 3.2

5.1 Configure access to PS Series storage

This section covers configuring access to PS Series storage through the installation and assumes VMware licensing is Enterprise or Enterprise Plus and supports the use of MEM.

5.1.1 PS Series Multipathing Extension Module (MEM) for VMware

The VMware vSphere host servers were configured using the best practices defined in the [Dell PS Series Configuration Guide](#), including the installation of MEM on the vSphere host servers.

For more information on MEM, see [Configuring and Installing the PS Series Multipathing Extension Module for VMware vSphere and PS Series SANs](#).

Note: MEM is only supported when using VMware vSphere Standard or above licensing.

After preparing the vSphere host, map volumes using Dell Storage Manager or Group Manager.

Note: See the PS Series [Group Administration Guide](#) for information on creating volumes and mapping volumes to hosts (requires a login account).

5.2 Configure access to SC Series storage

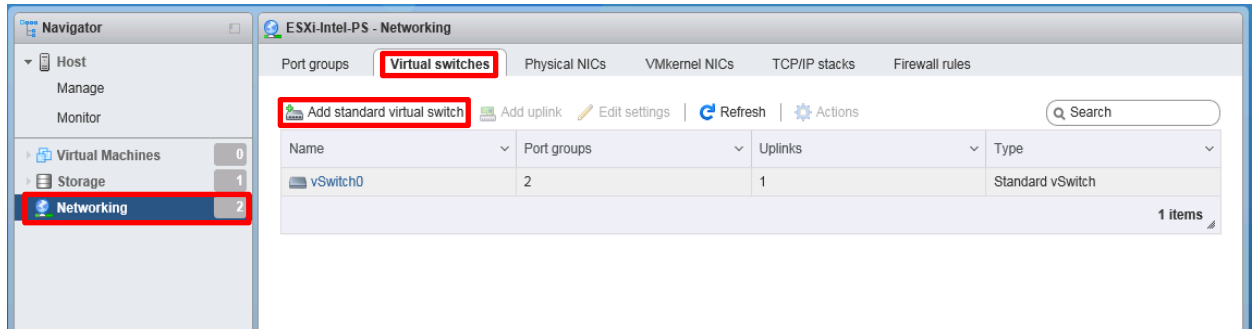
This section covers configuring vSphere hosts to access SC Series storage when those hosts have not been configured to access any other storage platform.

Note: If the hosts being configured have already been configured to access PS Series storage, go back to section 4.

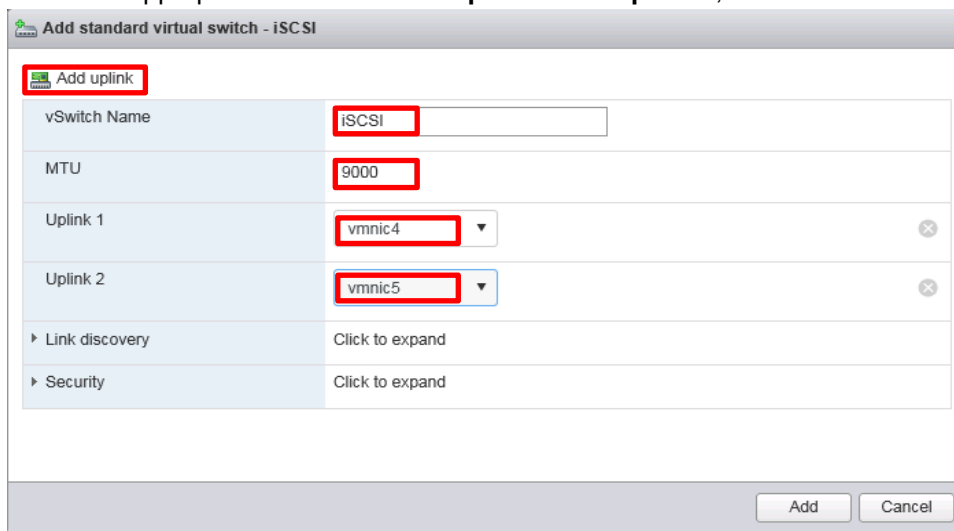
5.2.1 Configuring the VMware iSCSI software initiator to access SC Series volumes

The first step to configure the vSphere hosts for SC Series arrays using a single subnet and iSCSI port binding is to create VMkernel (vmk) ports. For port binding to work correctly, the initiator must be able to reach the target. With the release of vSphere 6.5, routing is supported when using port binding. However, prior to vSphere 6.5, iSCSI port binding did not support routing: requiring the initiator and target to be in the same subnet. VMware requires each VMkernel port used for iSCSI port binding to be associated with a single vmnic uplink.

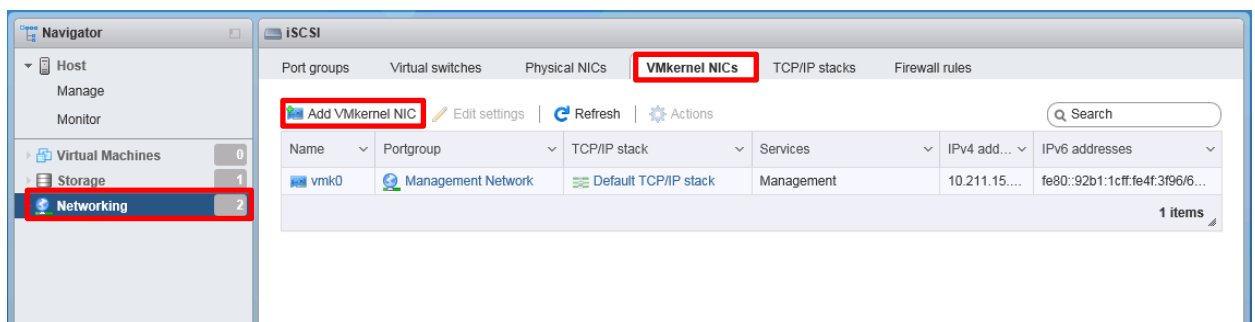
1. From the VMware ESXi Web GUI (local host management utility that has replaced the C-Sharp client), select **Networking**, **Virtual switches**, and click **Add standard virtual switch**.



2. Click on **Add uplink** to add the second uplink. Input the **vSwitch Name** (iSCSI), the **MTU** (9000), select the appropriate vmnics in both **Uplink 1** and **Uplink 2**, and click **Add**.



3. From **Networking**, select **VMkernel NICs** and click **Add VMkernel NIC**.



4. Input the name for the **New Port Group**, select the **Virtual Switch** from the drop-down box, enter the **MTU** (9000), select the **IP Version** from the drop-down box, expand the **IP settings** to enter **Static** information (or use DHCP), and click **Create**.

Add VMkernel NIC

Port group: New port group

New port group: iSCSI-01

Virtual switch: iSCSI

VLAN ID: 0

MTU: 9000

IP version: IPv4 only

IPv4 settings: DHCP Static

TCP/IP stack: Default TCP/IP stack

Services: vMotion Provisioning Fault tolerance logging
 Management Replication NFC replication

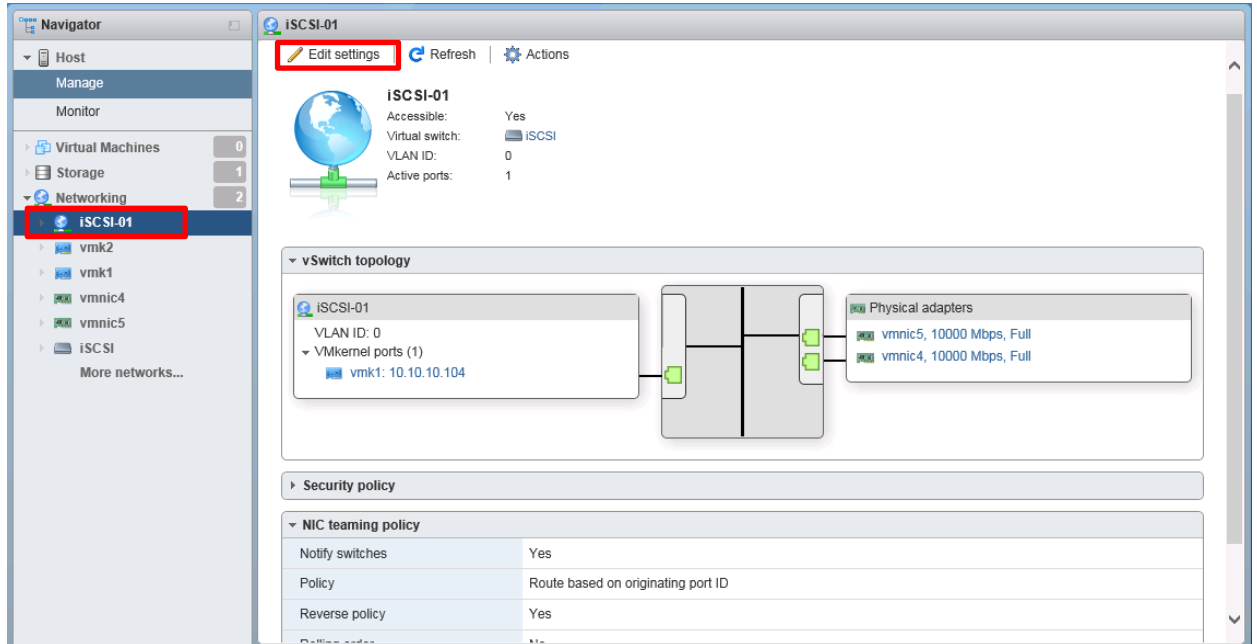
Create Cancel

5. Repeat steps 3-4 for the second **VMKernel NIC** (Shown completed).

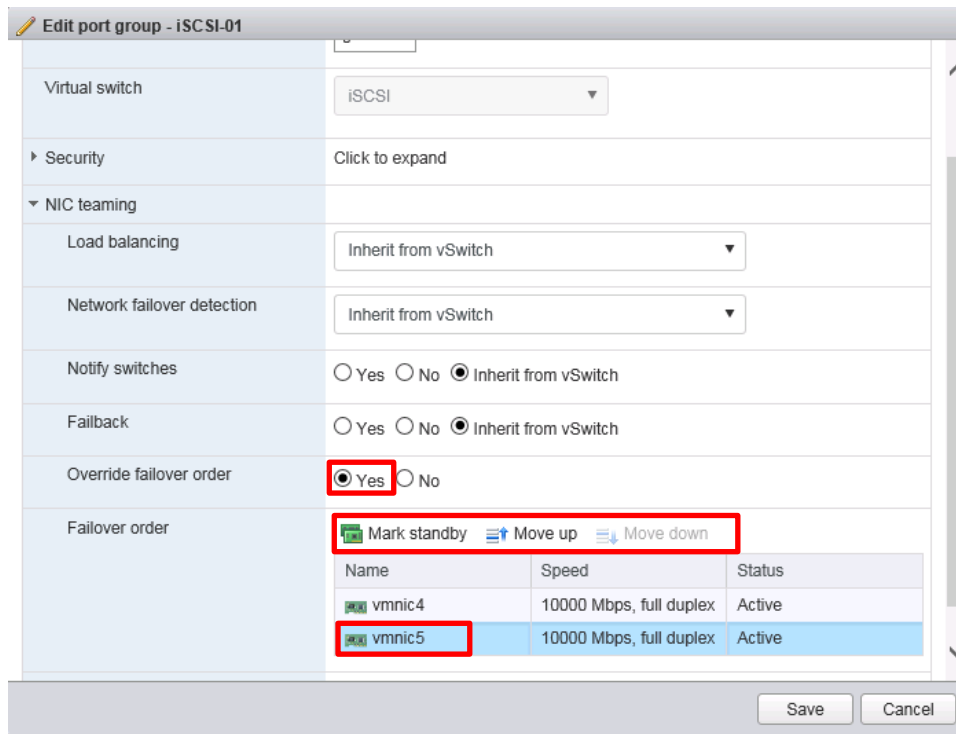
Name	Portgroup	TCP/IP stack	Services	IPv4 address	IPv6 addresses
vmk0	Management Network	Default TCP/IP stack	Management	10.211.15.199	fe80::92b1:1cff:fe4f:3f96/6...
vmk1	iSCSI-01	Default TCP/IP stack		10.10.10.104	fe80::250:56ff:fe6f:5f90/64
vmk2	iSCSI-02	Default TCP/IP stack		10.10.10.105	fe80::250:56ff:fe60:6bf4/64

Note: The use of port binding when multiple vmnic uplinks exist on the vSwitch, all but one uplink must be placed in an unused state for each VMkernel port group.

6. Expand **Networking**, select the VMkernel port group (**iSCSI-01**), and click **Edit settings**.



- When the ability to mark an adapter as unused becomes available, select **Yes** for **Override failover order**, select the **vmnic** to mark as unused, and click the **Unused (currently not an available option)**.



Note: At the time of this writing, the local vSphere host GUI, the new web-based utility that has replaced the C-Sharp client, does not have the ability to set the secondary vmnic to unused for the kernel port group. This can be accomplished either using the vCenter Web GUI, or a command line utility.

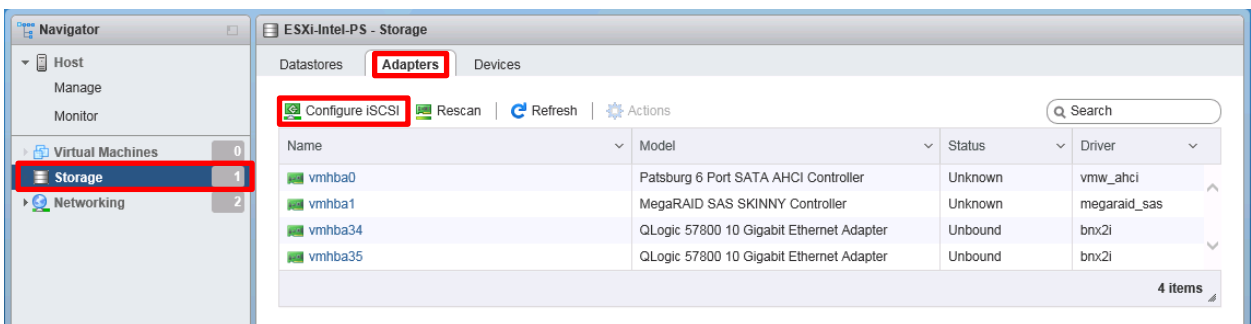
- Repeat steps 6-7 for the second **VMKernel port group (iSCSI-02)**.
Or, the command to execute at the command line to complete this task is:

```
[root@ESXi-Intel-PS:~] esxcli network vswitch standard portgroup policy failover/ set -p portgroupname -a vmnic-to-be-active
```

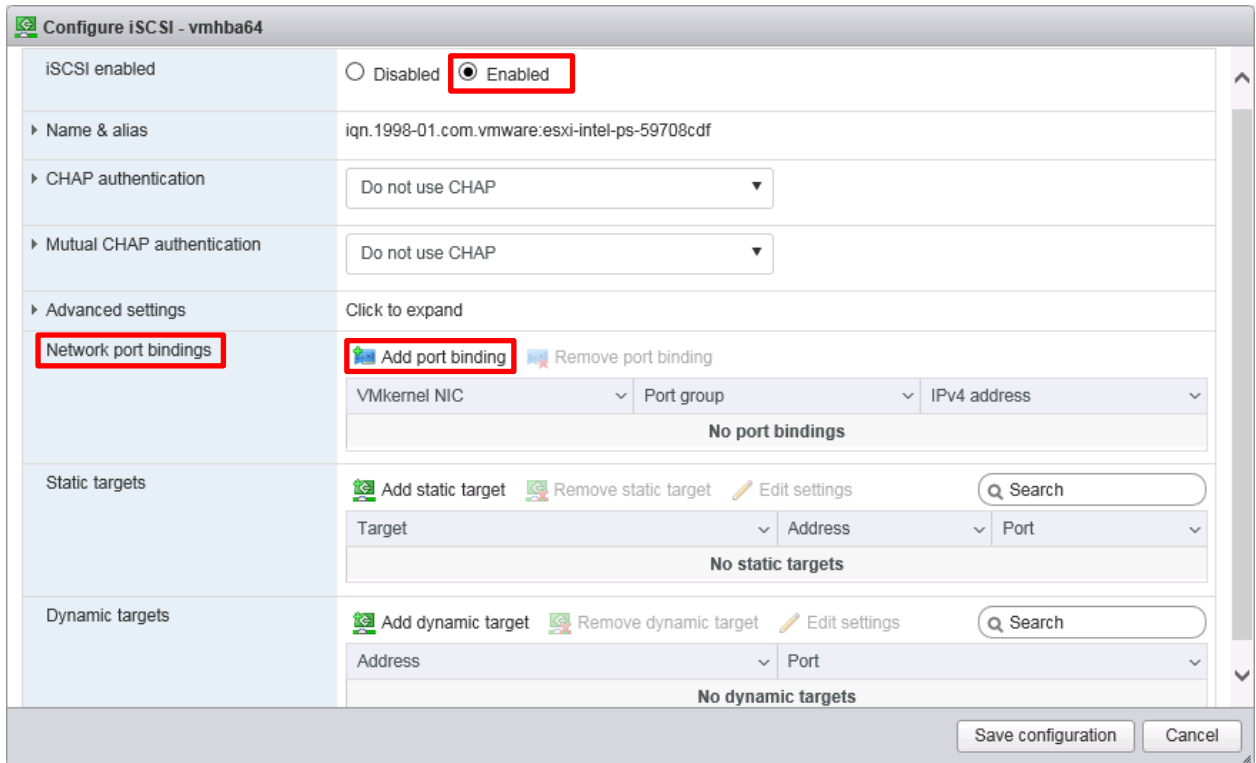
Note: Be sure to execute this command for each VMKernel port group being used, as shown below.

```
[root@ESXi-Intel-PS:~] esxcli network vswitch standard portgroup policy failover set -p iSCSI-01 -a vmnic4
[root@ESXi-Intel-PS:~] esxcli network vswitch standard portgroup policy failover get -p iSCSI-01
Load Balancing: srcport
Network Failure Detection: link
Notify Switches: true
Failback: true
Active Adapters: vmnic4
Standby Adapters:
Unused Adapters: vmnic5
Override Vswitch Load Balancing: false
Override Vswitch Network Failure Detection: false
Override Vswitch Notify Switches: false
Override Vswitch Failback: false
Override Vswitch Uplinks: true
[root@ESXi-Intel-PS:~] esxcli network vswitch standard portgroup policy failover set -p iSCSI-02 -a vmnic5
[root@ESXi-Intel-PS:~] esxcli network vswitch standard portgroup policy failover get -p iSCSI-02
Load Balancing: srcport
Network Failure Detection: link
Notify Switches: true
Failback: true
Active Adapters: vmnic5
Standby Adapters:
Unused Adapters: vmnic4
Override Vswitch Load Balancing: false
Override Vswitch Network Failure Detection: false
Override Vswitch Notify Switches: false
Override Vswitch Failback: false
Override Vswitch Uplinks: true
[root@ESXi-Intel-PS:~] █
```

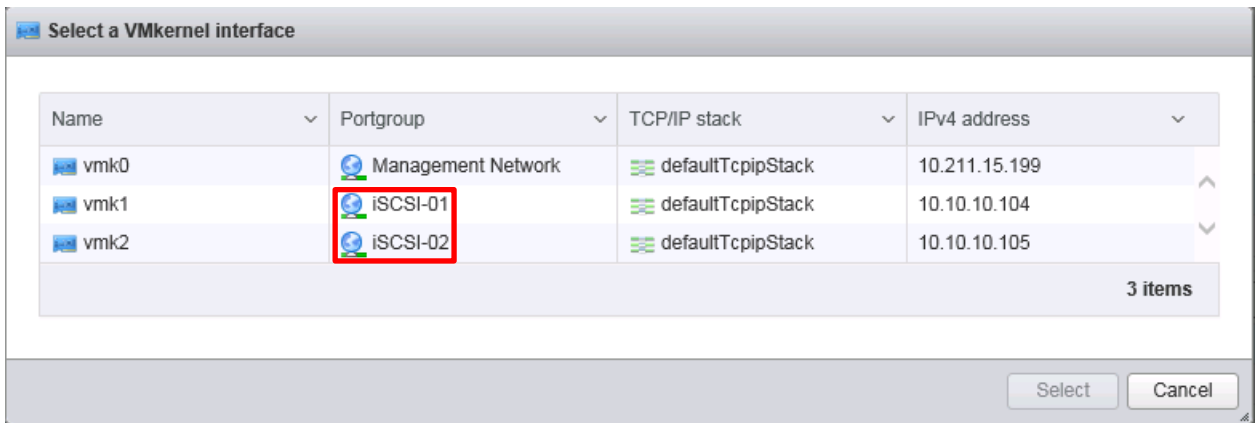
- Select Storage, Adapters, and click **Configure iSCSI**.



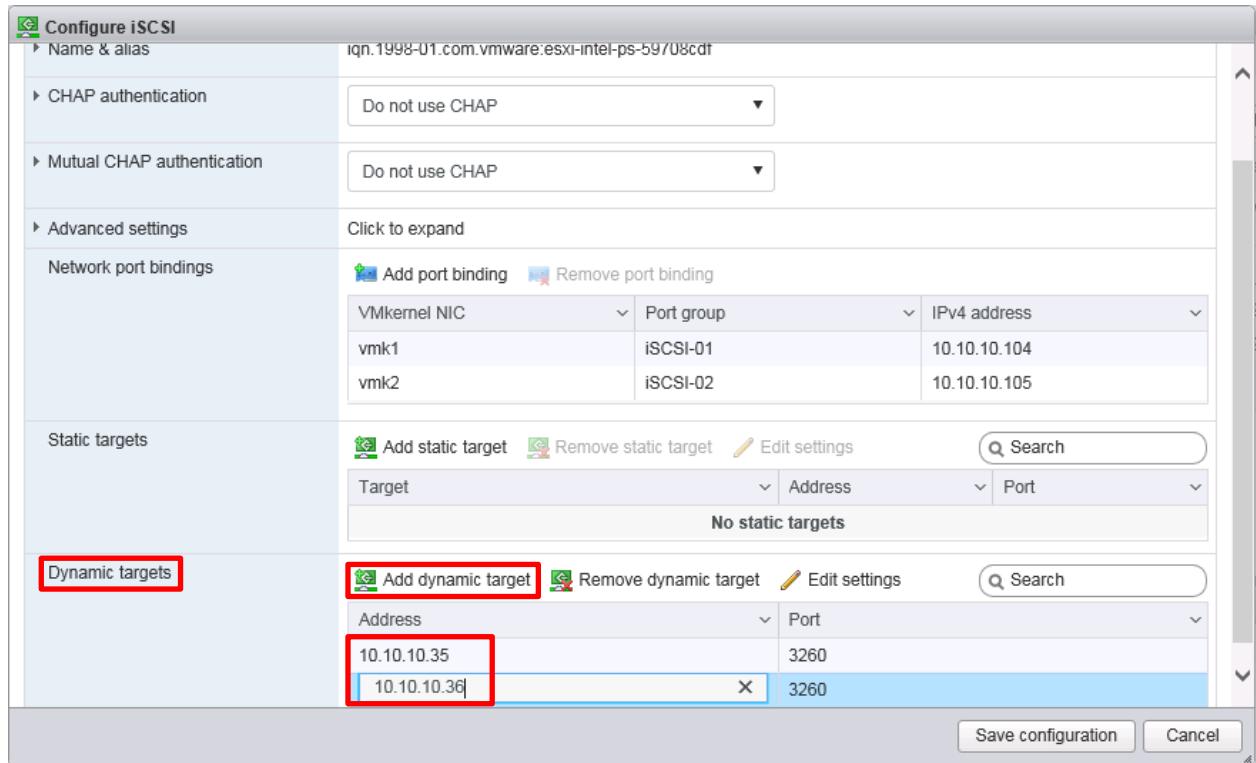
10. In the iSCSI enabled, click **Enabled** and click **Add port binding** in the **Network port bindings** section.



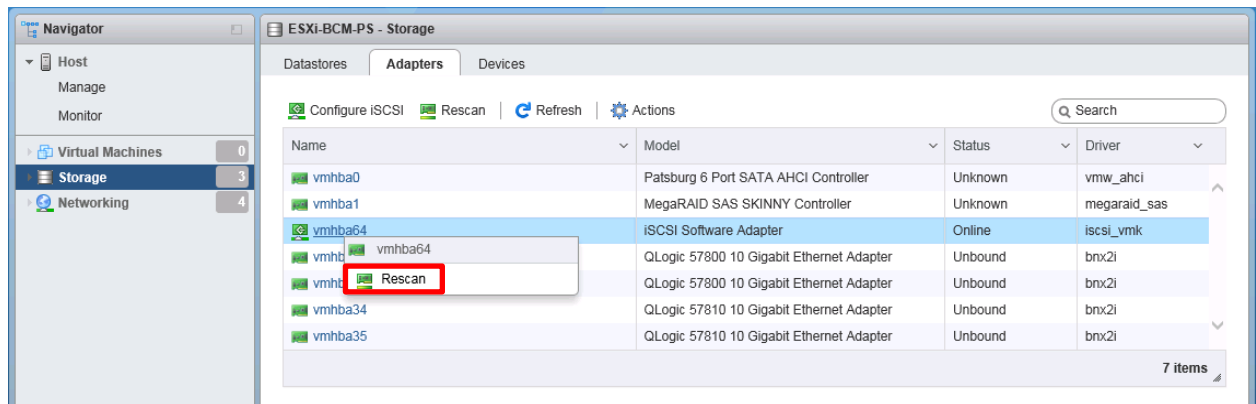
11. Select the first iSCSI VMKernel **Portgroup** and click **Select**. Repeat steps 10 and 11 for the second VMKernel Portgroup.



12. In the **Dynamic targets** section, click **Add dynamic target** twice, enter the two SC Series iSCSI target addresses, and click **Save configuration**.



13. Right-click the **iSCSI Software Adapter** and click **Rescan** to connect to the targets.



Note: This approach with VMware does not provide a way to restrict initiators to connect only to targets on the same fault domain.

After preparing the vSphere host, create a server or cluster object in Dell Storage manager and map volumes.

Note: See the [Create a Cluster Object in Enterprise Manager](#) and [Creating and mapping a volume in Enterprise Manager](#) videos for additional information.

5.2.2 VMware native multipathing

VMware provides native multipathing that can be used with any storage platform, if a vendor-provided multipath solution is not available.

5.2.3 Setting Path Selection Policy and storage performance settings - SC Series

VMware hosts connected to the SC Series storage platform must use VMware's native multipathing for SC Series volumes. Setting the Path Selection Policy (PSP) can be performed using either the vCenter GUI or a command line utility. Setting the PSP Default, and performance aspects of the PSP can only be performed using a command line utility at this time. Command line utility options are: ESXi console, SSH session, vSphere CLI, or VMware PowerCLI. The SSH session commands are shown here:

Note: At the time of this writing, the local vSphere host GUI, the new web-based utility that has replaced the C-Sharp client, does not have the ability to set PSP for volumes.

To set Path Selection Policy default for new SC Series (SCOS v6.6 and up) volumes:

```
[root@ESXi-BCM-PS:~] esxcli storage nmp satp set -P VMW_PSP_RR -s VMW_SATP_ALUA
```

To set or change the PSP for an existing SC Series volume:

```
[root@ESXi-BCM-PS:~] esxcli storage nmp device set --device=naa.xxx --psp=VMW_PSP_RR
```

To set the Round Robin IOPS policy for an existing SC Series volume:

```
[root@ESXi-BCM-PS:~] esxcli storage nmp psp roundrobin deviceconfig set --device=naa.xxx/  
--type=iops --iops=3
```

Note: Additional details, scripting options, and instructions for previous versions of SCOS are provided in [Dell Storage SC Series Best Practices with VMware vSphere 5.x-6.x](#).

6 Test methodology

Note: The test methodology used in this paper is only valid in the presented topology, components, and test cases. The results may not be directly applicable to environments with other components or variables. Additionally, the updated versions of vSphere, PS Series firmware, and SC Series SCOS show no indications of presenting different performance than that of previous versions. As such, the performance testing reflects the previous version of this document.

Key testing factors discussed in this section include:

- The test cases assess the impact from the operational and performance perspective of the shared host, host NICs, and shared Ethernet SAN infrastructure using the iSCSI protocol.
- The main objective was to ensure that the VMware vSphere hosts with both shared and dedicated NICs and shared Ethernet switches would sustain I/O workloads on Dell PS Series storage when Dell SC Series storage is added.
- SAN component failures were induced to confirm that the solution is highly available, and that a failure on one type of storage did not impact the other.
- The tests focused on validating the capability of the hosts and switches to handle both PS Series and SC Series storage with respect to network adapters on the host, switch hardware, buffer settings on the switch, Ethernet flow control, and resiliency of the storage network.

A baseline test running a specific I/O workload was performed with PS Series storage in a standalone environment. The standalone environment consisted of the VMware vSphere hosts accessing only PS Series arrays. The test was then repeated in both dedicated and shared coexistence environments (including PS Series and SC Series storage) with both storage types sharing the Ethernet iSCSI SAN to ensure that the data could be compared directly without introducing variables beyond those specified in the test plan. A high availability test for link failures, switch failure, and other factors was carried out to prove solution robustness.

In order to determine the relative performance of each SAN design, the performance tool [vdbench](#) was used to capture throughput values at four distinct I/O workloads:

- 8 KB transfer size, random I/O, 70% read
- 64 KB transfer size, random I/O, 70% read
- 256 KB transfer size, sequential I/O, 100% read
- 256 KB transfer size, sequential I/O, 100% write

Note: Information about vdbench is available at <http://sourceforge.net/projects/vdbench/>.

The above I/O workloads were first run for one hour each in a series with vSphere host servers accessing PS Series volumes only in the iSCSI SAN network to establish a baseline. Next, the same I/O workloads were run for one hour each with each of the vSphere hosts accessing either PS Series or SC Series volumes sharing the same iSCSI SAN network. Finally, the same I/O workloads were run for one hour each with both of the vSphere hosts accessing both PS Series and SC Series volumes using the same host NICs with both storage platforms sharing the same iSCSI SAN network. The results for both shared and dedicated tests were compared against the baseline for any change in performance.

IOPS, latency measured in milliseconds (ms), throughput measured in megabytes per second (MBPS), retransmits (an indicator of network congestion), and other statistics (such as switch packet drops, discards, and pause frames) were collected from host, switch, and storage ports. These parameters were carefully analyzed to ensure that the host and iSCSI SAN network performance were not adversely impacted with both PS Series and SC Series storage being accessed simultaneously.

6.1 Test environment

The test environment consisted of the following:

- Two Dell PowerEdge™ R620 servers running vSphere 6.0
- Three Dell PS6210 arrays (72 x 15K disks)
- Two Dell SC8000 storage controllers with three Dell SC220 enclosures (72 x 15K disks)

One Windows Server® 2008 R2 guest VM was created on each vSphere host. PS Series MEM was installed on each vSphere host and native VMware iSCSI Software Initiators were used to connect to PS Series and SC Series volumes. Each vSphere host had 8 LUNs or volumes from PS Series arrays and 8 LUNs from SC Series storage mapped for use. On each mapped LUN or volume, a VMware datastore was created. On each of the datastores, as required by each testing scenario, a virtual disk was created and presented to the Windows guest VM.

6.2 I/O performance testing

Testing scenarios for this paper are as follows:

- Baseline test: The vSphere hosts were connected to volumes on PS Series arrays only.
- Shared coexistence environment test: Both vSphere hosts were connected to both PS Series and SC Series arrays to show the effects of adding SC Series storage to an existing PS Series SAN environment, which includes hosts accessing both storage platforms.
- Dedicated coexistence environment test: One vSphere host was connected to volumes on PS Series storage only, and the other vSphere host was connected to volumes on the SC Series storage only. This test shows the effects of adding SC Series storage traffic to the existing PS Series iSCSI SAN infrastructure only.

6.2.1 I/O performance results and analysis: shared hosts

Performance analysis in the shared-host environment is designed to show the impact to the PS Series storage environment when SC Series storage is added to the existing environment, and existing hosts access storage from both the PS Series and SC Series storage platforms over the same iSCSI adapters and iSCSI infrastructure.

The following figures show the baseline IOPS and latency information gathered on the PS Series environment and how it compares after the introduction of the SC Series storage to the environment. Figure 4 and Figure 5 show the baseline information on the left and the shared environment data on the right, as a percentage of the baseline. As shown in the charts, there was no significant decrease in IOPS nor increase in latency with the introduction of the SC Series storage platform to the environment.



Figure 4 Shared host performance data: 8K and 64K

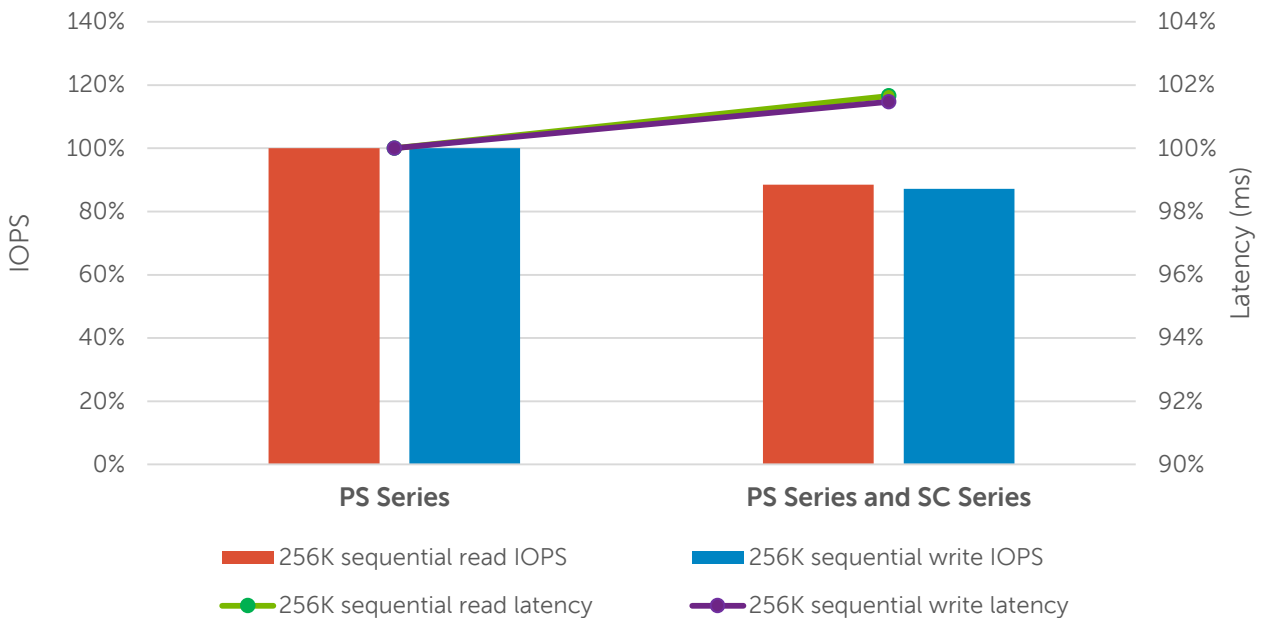


Figure 5 Shared host performance data: 256K read and write

6.2.2 I/O performance results and analysis: dedicated hosts

Performance analysis in the dedicated host environment is designed to show the impact to the PS Series storage environment when a dedicated SC Series storage environment is added to the same iSCSI infrastructure.

Figure 6 below shows the baseline IOPS and latency information gathered on the PS Series environment on the left side. The right side shows the PS Series environment IOPS and latency, as a percentage of the baseline, with the same workload running on the SC Series environment that shares the iSCSI infrastructure. This chart shows no significant decrease in IOPS nor increase in latency when SC Series storage environment shares the iSCSI infrastructure.

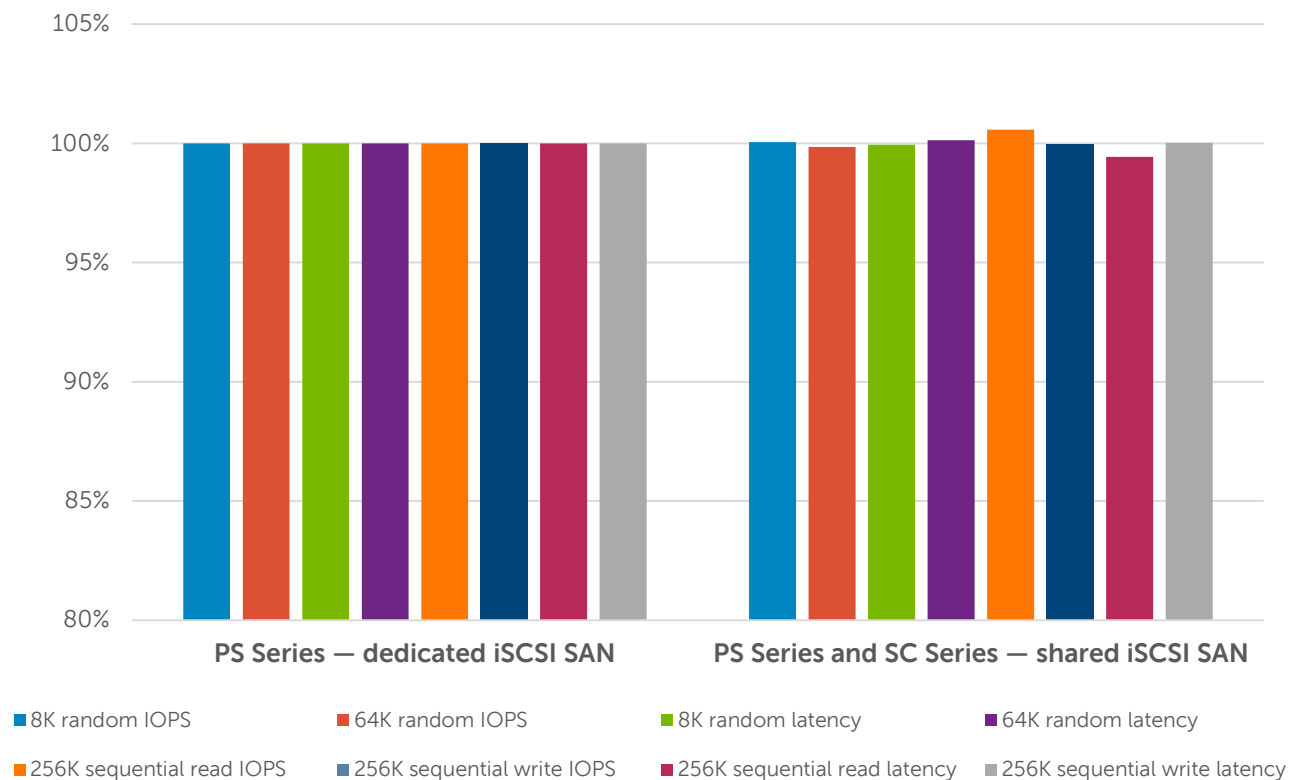


Figure 6 Dedicated host performance data

6.3 High availability testing

High availability was demonstrated by showing performance during an iSCSI switch failure. In this test scenario, both VMware vSphere servers were connected to both PS Series and SC Series storage platforms. The workload VMs, one on each vSphere server, each had eight virtual disks on each storage platform.

Both VMs started running a vdbench 256K Write workload. Roughly midpoint in the hour-long test, one of the two iSCSI switches were powered off, simulating a switch failure. The output chart in Figure 7 shows a drop in IOPS and a spike in latency at the time of the failure, as would be expected. Also, as expected, the latency quickly recovers to roughly double the latency observed prior to the failure. Additionally, at the time of the failure, IOPS encounter a momentary drop, but quickly recover to roughly half of the pre-failure IOPS. Overall, the performance observed during the switch failure is consistent with reducing the paths over which storage traffic is passed, by half.

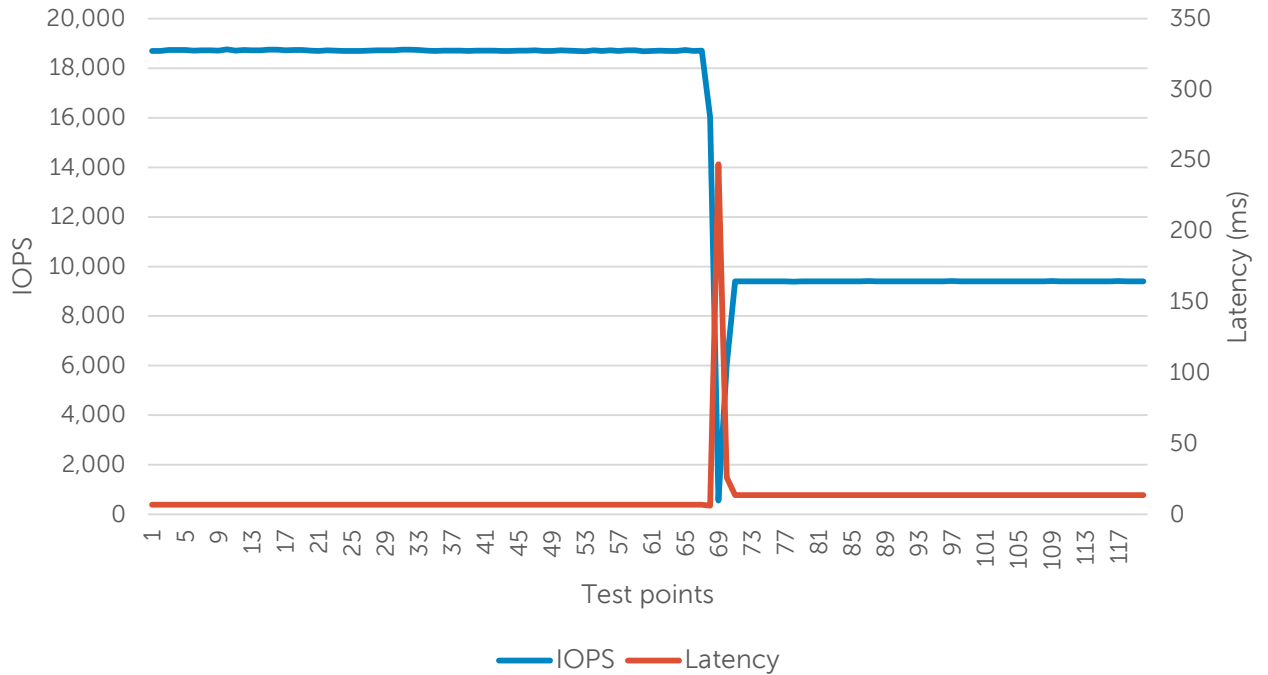


Figure 7 High availability test

7 Best practice recommendations

Always check the [Dell Storage Compatibility Matrix](#) to ensure that SAN components (such as switches and host NICs) are supported for both PS and SC Series deployments. For full support, the switches used for array connectivity must be validated for use with each solution individually.

7.1 Switch fabric

Place all storage interfaces (including storage controller ports and host ports) in the default (native) VLAN when the iSCSI SAN infrastructure is shared. If a non-default VLAN is assigned, set the ports as **untagged** or put them in **access** mode.

Use the Dell [switch configuration guides](#) for PS Series SANs to configure the switch fabric. These guides contain optimal settings for jumbo frames, flow control, and other common settings recommended for Dell storage.

A switch interconnect (such as LAG, stack, VLT, or VPC) must be configured to maintain a single Layer 2 fabric topology so that PS Series arrays can share the same infrastructure.

7.2 Host connectivity

Hosts must be configured for IPv4 connectivity. IPv6 is not supported at this time for shared connectivity.

Enable Jumbo frames (MTU of 9000 or greater) on host NICs used for storage connectivity to allow for greater throughput with large I/O request (block) sizes.

7.3 Storage

PS Series storage will auto-negotiate jumbo frame size. For SC Series storage, enable jumbo frames on each iSCSI SAN port through Dell Storage Manager. Refer to the SC Series administration guide on the Knowledge Center at the [Customer Portal](#) (login required) for the configuration steps.

8 Conclusion

As a result of testing and analysis shown in this document, introducing SC Series storage into an existing PS Series storage iSCSI environment, whether in a dedicated or shared host configuration, will not significantly affect the performance of the PS Series storage. Understanding the effects of a coexisting storage environment is key to enabling future growth and expansion of existing PS Series environments.

A Technical support and resources

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.

We encourage readers of this publication to provide feedback on the quality and usefulness of this information by sending an email to StorageSolutionsFeedback@dell.com.

A.1 Related documentation

See the following referenced or recommended Dell publications:

- Best Practices for Implementing VMware vSphere in a Dell PS Series Storage Environment:
http://en.community.dell.com/techcenter/extras/m/white_papers/20434601
- Configuring and Installing the PS Series Multipathing Extension Module for VMware vSphere and PS Series SANs:
<http://en.community.dell.com/dell-groups/dtcmmedia/m/mediagallery/19991633/>
- Configuring iSCSI Connectivity with VMware vSphere 5 and Dell PS Series Storage:
<http://en.community.dell.com/dell-groups/dtcmmedia/m/mediagallery/19997606>
- Dell PS Series Configuration Guide:
<http://en.community.dell.com/dell-groups/dtcmmedia/m/mediagallery/19852516>
- Dell PS Series Storage *Group Administration Guide* (requires a valid portal account):
https://eqsupport.dell.com/support/download_file.aspx?id=2926
- Dell Storage Compatibility Matrix:
<http://en.community.dell.com/dell-groups/dtcmmedia/m/mediagallery/19856862.aspx>
- Dell EMC SC Series Best Practices with VMware vSphere 5.x-6.x:
http://en.community.dell.com/techcenter/extras/m/white_papers/20441056
- Switch Configuration Guides for PS Series or SC Series SANs:

<http://en.community.dell.com/techcenter/storage/w/wiki/4250.switch-configuration-guides-for-EqualLogic-or-compellent-sans.aspx>

- VMware ESXi 5.1, 5.5, or 6.0 Host Configuration:

<http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/20094619/>