

# Best Practices for Sharing an iSCSI SAN Infrastructure with PS Series and SC Series using Windows Hosts

Dell EMC Engineering  
April 2017

## Revisions

Date	Description
March 2015	Initial release
April 2015	Added granular Windows setting and specific NIC advanced settings
December 2015	Added dedicated host information
January 2017	Template update and minor revisions
April 2017	Update to TCP Connection offload best practice

## Acknowledgements

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

Engineering: Mark Boeser and Chuck Armstrong

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. [4/21/2017] [Best Practices] [2014-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 PS Series specific settings .....	10
3.2 SC Series specific settings .....	11
2.3.1 SC Series host physical connectivity and IP assignment.....	11
4 Enabling Windows host access to PS Series and SC Series iSCSI storage .....	13
4.1 Host Integration Tools (HIT) kit for Windows.....	13
4.2 Windows DSM for SC Series storage multipathing .....	13
4.3 Configuring the Microsoft iSCSI software initiator to access SC Series volumes .....	15
4.4 Shared host connections to both PS Series and SC Series .....	17
4.5 Dedicated host connections to each PS Series or SC Series .....	17
5 Test methodology .....	18
5.1 Test environment.....	19
5.2 I/O performance testing.....	19
2.5.1 I/O performance results and analysis: shared hosts .....	20
2.5.2 I/O performance results and analysis: dedicated hosts .....	21
5.3 High availability testing .....	24
3.5.1 High availability results and analysis .....	24
6 Best practice recommendations .....	25
6.1 Switch fabric .....	25
6.2 Host connectivity.....	25
6.3 Storage .....	25
7 Conclusion.....	26
A IP address tables.....	27

B	iSCSI host networking best practice settings .....	28
B.1	Windows registry settings .....	28
B.2	Windows Server 2008 R2 NIC best practices .....	33
B.2.1	BCM57810 NDIS mode recommended configuration .....	34
B.2.2	Intel X520 recommended configuration .....	34
B.3	Windows 2012 R2 NIC best practices .....	35
B.3.1	Broadcom BCM57810 NDIS mode recommended configuration .....	35
B.3.2	Intel X520 recommended configuration .....	35
B.4	NIC-specific settings summary .....	36
B.5	iSCSI-specific settings .....	36
B.6	Windows 2012 Server, R2, and Hyper-V .....	37
C	Additional resources .....	38
C.1	Technical support and resources .....	38
C.2	Related resources .....	38

# 1 Introduction

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

- 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)
- 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 storage 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)
- Microsoft® Windows® 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
- Microsoft Hyper-V® configuration

## 1.2 Audience

This paper is intended 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 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](#).

**Device specific module (DSM):** Microsoft Multipath I/O (MPIO) provides a well-defined MPIO plug-in framework that allows storage providers to develop multipath solutions (DSMs) that contain the hardware-specific information needed to optimize connectivity with their storage arrays.

**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.

**Host Integration Tools (HIT):** PS Series Host Integration Tools include several easy-to-use, host-based software modules that tightly integrate PS Series SANs with hosts and applications, delivering advanced data protection, high availability and performance, and simplified management of application data and virtual machines for Microsoft and Linux® environments.

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

**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.

**Microsoft device-specific module (DSM):** Microsoft DSM is an inbox DSM and can work with many storage arrays that are active-active and SPC-3 compliant.

**Microsoft Multipath I/O (MPIO):** Windows Server includes integrated MPIO for high availability of Windows-based servers to SANs. Microsoft MPIO architecture supports iSCSI, Fibre Channel and serial-attached storage (SAS) SAN connectivity by establishing multiple sessions or connections to the storage array. Multipathing solutions use redundant physical path components — adapters, cables, and switches — to

create logical paths between the server and the storage device. In the event that one or more of these components fails, causing the path to fail, multipathing logic uses an alternate path for I/O so that applications can still access their data. MPIO can also provide significant I/O performance benefits in high workload environments by load balancing the parallel data paths in a round robin or least queue depth configuration.

**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 or server so that the computer or server can connect to a network. Most NICs are designed for a particular type of network (typically Ethernet) protocol (typically TCP/IP) and media.

**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.

**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

This section provides an overview of the Dell EMC™ storage products and technologies presented in this paper.

### 2.1 PS Series storage

The PS Series delivers 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 mid-range 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

The SC Series is a Dell EMC midrange storage solution featuring multi-protocol support and self-optimizing, tiering, compression, and deduplication capabilities. SC Series arrays can be configured with all flash, as a hybrid system, or with only traditional spinning disks, and feature 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-host solutions not only share the iSCSI SAN infrastructure, but also have hosts that simultaneously 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 or SC Series array, but not both. Dedicated-host coexistence utilizes a shared iSCSI SAN infrastructure only.




Figure 1 Shared iSCSI SAN with shared hosts reference topology




Figure 2 Shared iSCSI SAN with dedicated hosts reference topology

### 3.1 PS Series specific settings

The use case defined in this paper includes SC Series arrays sharing Windows host servers and the Ethernet iSCSI SAN with deployed and operational PS Series array(s). It is assumed that an Ethernet network using iSCSI and Windows host servers accessing PS Series storage is already configured using the best practice recommendations defined in the [Dell PS Series Configuration Guide](#). The best practice settings for PS Series storage apply when the Ethernet iSCSI SAN is shared with SC Series storage arrays.

The PS Series Host Integration Tools (HIT) for Microsoft Windows settings, including iSCSI port binding best practice settings for PS Series storage, are applicable when the Ethernet iSCSI SAN network is shared with SC Series arrays.

## 3.2 SC Series specific settings

A typical SC Series iSCSI implementation typically involves two separate, dedicated Ethernet fabrics as two fault domains, each 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, a single subnet for all host and storage ports is used.

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 SCSI target with a single LUN (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 in a single IP subnet SAN infrastructure, it is possible that a host port on one fault domain can access a target port on the other fault domain through the switch interconnection. Ensure that all host and array ports are connected to the same fault domain, physically and in iSCSI sessions. This eliminates 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 subsections discuss ways to ensure that the physical connectivity and iSCSI sessions are established correctly.

### 2.3.1 SC Series host physical connectivity and IP assignment

Since SC Series SAN switches are connected by an ISL and are in a single IP subnet, it is important to ensure that the fault domains are properly connected and that iSCSI sessions are properly established within their correct 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 and interconnected switches, it is possible for the hosts to access SC Series storage ports on both fault domains. The correct connectivity eliminates 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.

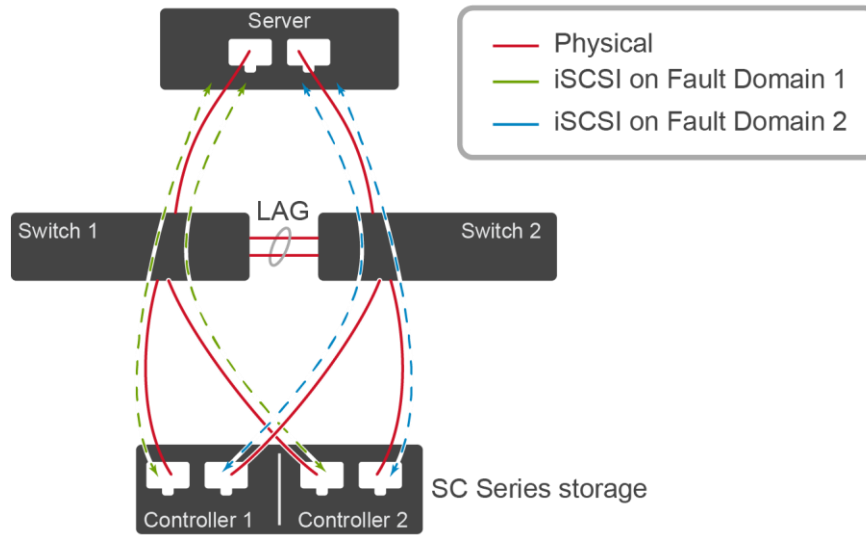


Figure 3 Connecting the host to SC Series ports

**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 Windows host access to PS Series and SC Series iSCSI storage

To enhance the coexistence of PS Series and SC Series array connectivity with shared or dedicated hosts, follow the best practices described in this section.

### 4.1 Host Integration Tools (HIT) kit for Windows

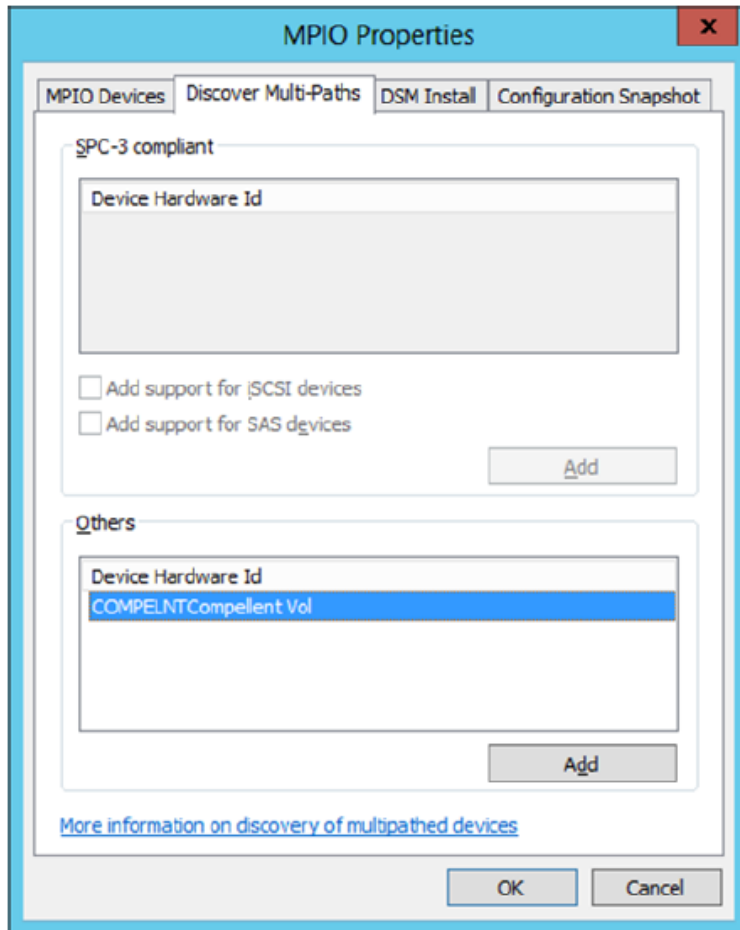
The Windows servers used in testing were configured using PS Series best practices defined in the [Dell PS Series Configuration Guide](#), including installing the HIT kit on the Windows host servers. Tests were performed with SC Series and PS Series iSCSI volumes served by the same Windows host servers with HIT installed. MPIO/iSCSI settings optimized by HIT either aligned or did not interfere with SC Series best practice settings. No performance impact was observed when Windows hosts serving PS Series storage began serving SC Series storage.

### 4.2 Windows DSM for SC Series storage multipathing

The Windows servers in the testing environment were configured with all recommended best practices for PS Series and HIT installed with the device-specific module (DSM) provided by PS Series. For the same host to access SC Series storage, it is required to implement the additional settings defined in this section.

SC Series iSCSI storage uses the native Windows MPIO that is built into Windows Server 2008 and later. The following steps associate the SC Series volume with a Windows native DSM. Configure the servers in the Enterprise Manager/Dell Storage Manager and map at least one volume to the server with MPIO enabled. For additional details, refer to the best practices document, [Dell Storage SC Series: Microsoft Multipath I/O Best Practices](#).

1. To open the MPIO control panel, click **Start > Administrative Tools > MPIO**.
2. Click **Discover Multi-Paths**. Under **Others**, **COMPELNTCompellent Vol** appears.
3. Click **Add** and click **OK**.



Another option for providing associations is to use the MPCLAIM CLI command. Open a command prompt with elevated (administrator) privileges and type:

```
mpclaim.exe -r -i -d "COMPELNTCompellent Vol"
```

This command provides the same result as the control panel option. It associates an SC Series volume and then restarts the server.

**Note:** Once the server reboots, verify that there is only one instance of each SAN volume listed in the Windows Disk Management administration tool.

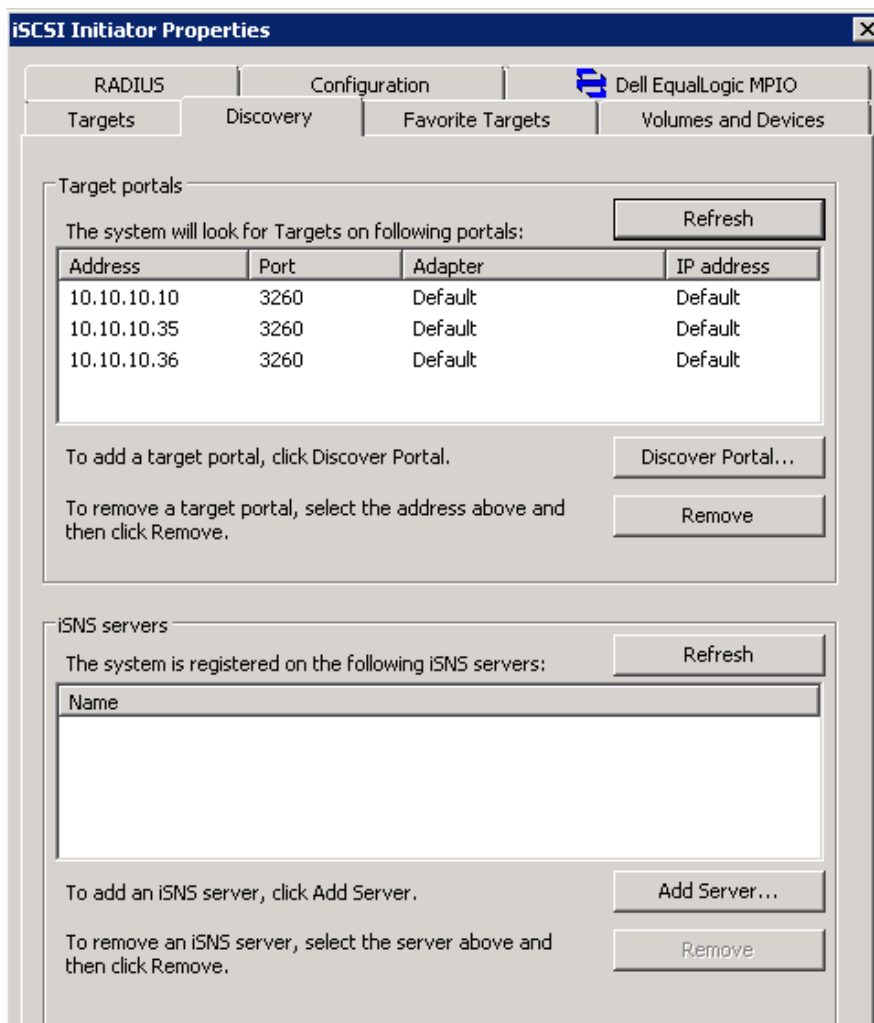
## 4.3 Configuring the Microsoft iSCSI software initiator to access SC Series volumes

Configure the Microsoft software iSCSI initiator accommodating both SC Series fault domains. For PS Series, refer to the [Dell PS Series Configuration Guide](#) for details.

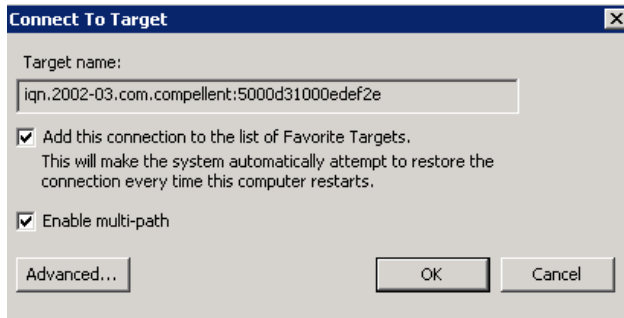
To configure the Microsoft software iSCSI initiator for SC Series storage using a single subnet:

4. In the Microsoft iSCSI initiator > **Discovery** tab, add the discovery portal addresses manually.

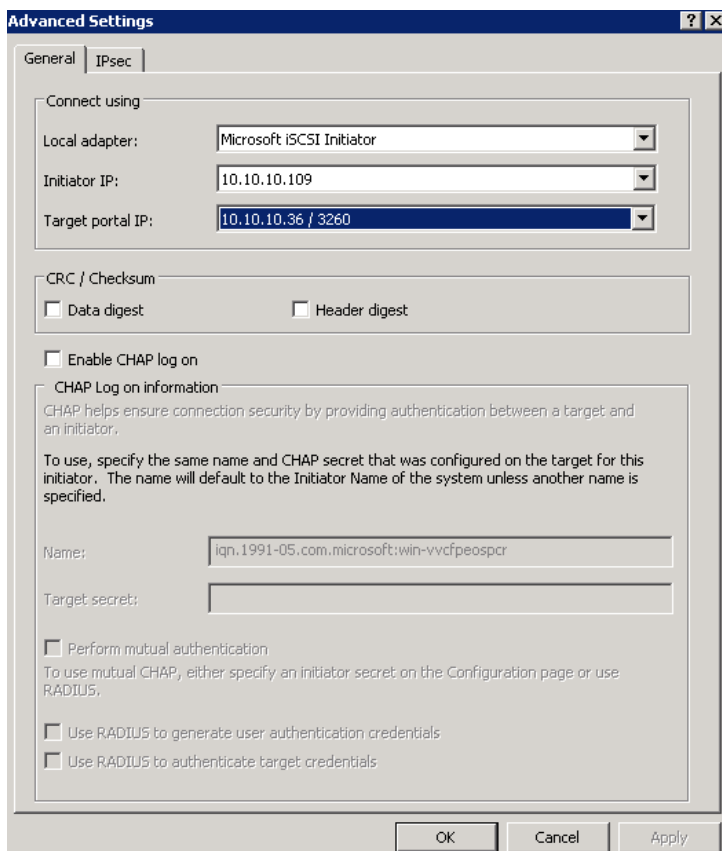
Refer to appendix A for an IP configuration used in testing. For the purpose of this discussion, 10.10.10.10 is the PS Series group IP, 10.10.10.35 is the control port for SC Series fault domain 1, and 10.10.10.36 is the control port for SC Series fault domain 2.



5. Make a connection to the targets. When creating the connections, select the **Enable multi-path** checkbox to enable MPIO for load balancing and path failover.



6. Click the **Advanced** button, and create a connection specifying the Microsoft iSCSI initiator, an initiator IP, and a corresponding target portal IP address. Click **OK**.
7. Make sure that the host port and target portal belong to the same fault domain. For the purpose of this discussion, 10.10.10.109 and 10.10.10.110 are used as the host NIC port IP addresses.



8. Repeat step 7 to connect every host NIC to its target based upon the IP address plans for each fault domain.
9. Once all connections are established with the targets, register the hosts with Storage Center or through Enterprise Manager/Dell Storage Manager and assign volumes (LUNs) to the host.



**Note:** For detailed Windows configuration settings such as registry value settings and timeout values, refer to appendix A.

## 4.4 Shared host connections to both PS Series and SC Series

In a shared environment, where Windows hosts connect to both PS Series and SC Series storage targets, the multipath and iSCSI settings optimized by HIT (described previously) either aligned or did not interfere with SC Series storage best practices.

## 4.5 Dedicated host connections to each PS Series or SC Series

When Windows hosts are dedicated to either the PS Series or SC Series storage platform, the native multipathing must be used when connecting to SC Series storage.

The following sample output shows the multipath selection policy changed from Fixed to Round Robin for SC Series volumes.

```
=====
MPIO Storage Snapshot on Wednesday, 09 January 2013, at 15:30:44.689
Registered DSMs: 1
=====
+-----+-----+-----+-----+-----+-----+
| DSM Name                | Version                | PRP | RC | RI | PVP | PVE |
+-----+-----+-----+-----+-----+-----+
| Microsoft DSM            | 006.0002.09200.16384 | 0020 | 0003 | 0001 | 030 | False |
+-----+-----+-----+-----+-----+-----+

Microsoft DSM
=====
MPIO Disk0: 02 Paths, Round Robin, ALUA Not Supported
SN: 600D310039C0000000892
Supported Load Balance Policies: FOO RR RRWS LQD WP LB

Path ID          State          SCSI Address      Weight
-----
0000000077050003 Active/Optimized 005|000|003|001    0
  Adapter: Microsoft iSCSI Initiator... (B|D|F: 000|000|000)
  Controller: 46616B65436F6E74726F6C6C6572 (State: Active)

0000000077050002 Active/Optimized 005|000|002|001    0
  Adapter: Microsoft iSCSI Initiator... (B|D|F: 000|000|000)
  Controller: 46616B65436F6E74726F6C6C6572 (State: Active)

MSDSM-wide default load balance policy: N\A
No target-level default load balance policies have been set.
=====
```

## 5 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.

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 Windows hosts with both shared and dedicated NICs and shared Ethernet switches would sustain I/O workloads on PS Series storage when 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 Windows host accessing only PS Series storage 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/>.

These I/O workloads were first run for one hour each in a series with Windows 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 Windows 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 ESXi 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.

## 5.1 Test environment

The test environment consisted of the following:

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

Two of the servers are configured to be Windows 2008 R2 hosts and two are configured as Windows 2012 R2 hosts. Each host has 8 LUNs or volumes from PS Series storage arrays and 8 LUNs from SC Series exposed to itself. The PowerEdge servers are connected to volumes of PS Series storage arrays for the baseline test in the standalone environment, and then are connected with the volumes on the SC Series along with the volumes of PS Series in a coexistence environment when SC Series storage is introduced in the environment.

IOPS, latency measured in ms, throughput measured in MB/s, retransmits (an indicator of network congestion), and other statistics (such as switch packet drops, discards, and pause frames) were observed from host, switch, and storage ports. These parameters were carefully analyzed to ensure that the host and iSCSI SAN network performance was not adversely impacted with both PS Series and SC Series storage in use simultaneously.

## 5.2 I/O performance testing

Testing scenarios for this paper are as follows:

**Baseline test:** The Windows hosts were connected to volumes on PS Series storage arrays only.

**Shared coexistence environment test:** Windows hosts were connected to both PS Series and SC Series storage 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:** Two Windows hosts were connected to volumes on PS Series storage only, and the other Windows hosts were 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.

## 2.5.1 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 when existing hosts access storage from both the PS Series and SC Series storage platforms over the same iSCSI adapters and iSCSI infrastructure.

Figure 4 outlines the performance impact on PS Series storage for a small block random I/O workload. It compares the performance of an 8K block size random I/O workload (8K random) and a 64K block size random I/O workload (64K random) on a PS Series storage array in a standalone and a coexistence environment. The IOPS for the 8K random workload has no effect and the IOPS for 64K random workload is reduced marginally. However, the reduction of IOPS for the 64K random workload is accompanied by a reduction in latency, suggesting that a bottleneck occurs on the server while accessing PS Series and SC Series LUNs concurrently, but there is not a limitation on the PS Series storage arrays. These observations show that there is minimal-to-no impact on the PS Series performance when SC Series storage is introduced in the same SAN environment with common Windows hosts accessing LUNs from the PS Series and SC Series.



Figure 4 Effect on PS Series small block random I/O performance by introducing SC Series

Throughput performance for large block sequential I/O on PS series in standalone and coexistence environments is shown in Figure 5. 256K block size sequential read and 256K sequential write workload performance on PS Series storage is minimally affected by the introduction of SC Series in the SAN environment.

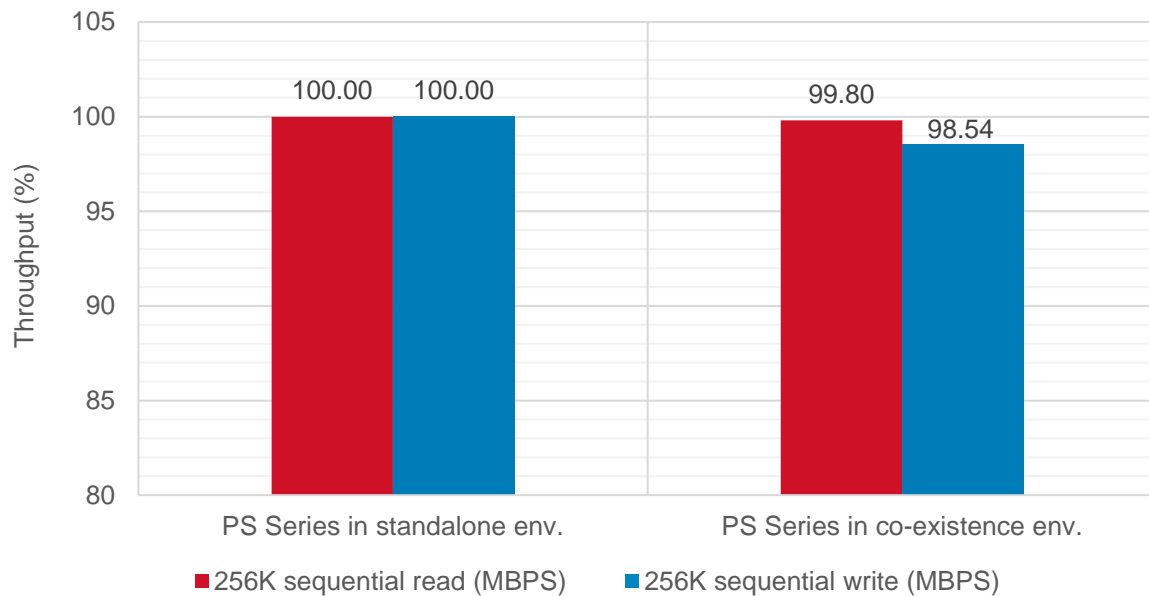


Figure 5 Effect on PS Series large block workload performance by introducing SC Series

The previous observations show there is minimal-to-no effect on the PS Series performance after introducing SC Series storage in the same environment. Latency almost stayed at the same levels as baseline, which is a clear indicator that the shared hosts and iSCSI SAN were able to sustain the shared storage efficiently.

## 2.5.2 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 5 shows the performance of standalone PS Series I/O with its own dedicated servers compared to I/O performance of PS Series with SC Series storage sharing the SAN infrastructure. Throughput for large block sequential I/O did not deteriorate substantially when SC Series I/O was sharing the same iSCSI SAN. The same results were seen for IOPS with random I/O. No major deterioration was observed. Latency of random I/O did not spike above acceptable limits. It almost stayed at the same level; a clear indicator that the shared iSCSI SAN was able to sustain the shared storage efficiently.

**Note:** There was not a significant impact to the IOPS, latency or throughput after adding a SC Series array to the existing PS Series iSCSI SAN infrastructure.

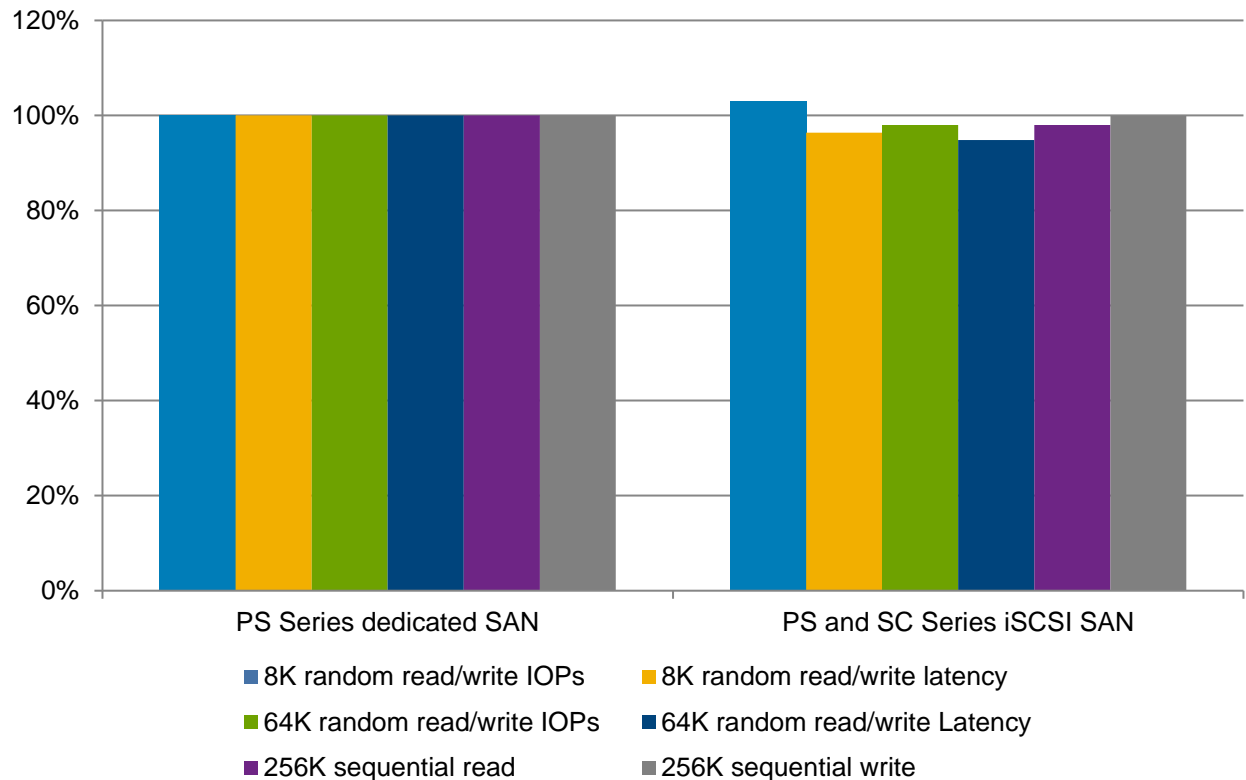


Figure 6 PS Series baseline versus performance of PS Series storage with SC Series storage sharing the iSCSI SAN infrastructure

Figure 7 shows the I/O performance of standalone SC Series storage with dedicated servers compared to I/O performance of combined SC Series and PS Series storage in a shared SAN infrastructure. Throughput for large block sequential I/O did not deteriorate substantially when SC Series I/O was sharing the same iSCSI SAN. The same applies for IOPS with random I/O. No major deterioration was observed. Latency of the random I/O did not spike above acceptable limits and was better or stayed at the same level. This is a clear indicator that the shared iSCSI SAN was able to sustain the shared storage efficiently without any adverse network impact.

**Note:** These test results show that the shared iSCSI SAN was able to sustain the shared storage efficiently without any adverse network impact when an SC Series array was added to the existing PS Series iSCSI SAN.

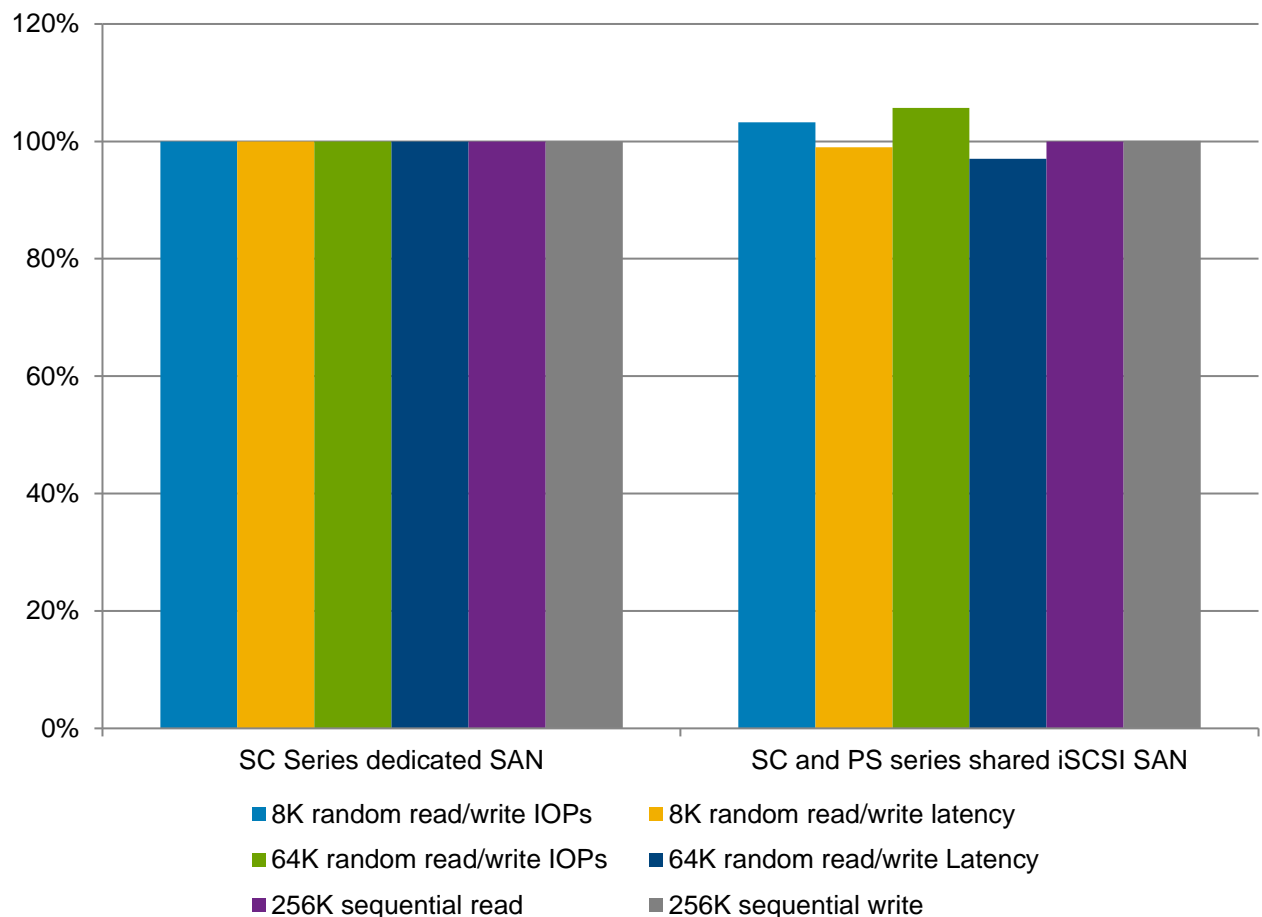


Figure 7 SC Series baseline compared to performance of SC Series storage with PS Series storage sharing the iSCSI SAN infrastructure

## 5.3 High availability testing

High availability tests were performed on the host, network, and storage to ensure that the hosts sharing the storage and iSCSI SAN infrastructure are able to sustain a single point of failure without causing any adverse impact on the SAN. The tests were performed with an I/O workload of 256K sequential write with read verification running on the hosts simultaneously accessing both PS Series and SC Series volumes. Monitoring parameters for performance and errors were carefully observed and gathered to analyze the impact of the failure on the SAN and its components. I/O was monitored for data corruption errors which could be an indicator that something went wrong when the failure was introduced.

### 3.5.1 High availability results and analysis

The following tests were performed on the Windows hosts, switches, and storage systems to verify that the SAN with shared Windows hosts, NICs, and network could sustain a single point of failure without any adverse effects.

- Single host port failure (with dual port adapter)
- Switch power cycle (one switch at a time)
- PS Series and SC Series storage array controller failover

There was no data corruption, and I/O either continued with no impact or resumed without any timeouts or halts in all HA tests. Figure 8 shows the I/O behavior for the PS Series and SC Series storage array controller failover HA test. This test involved array controller failover for PS Series and SC Series storage arrays. In the figure, I/O resumed after controller failover without any timeouts during the HA test, and latency spiked momentarily during the test and resumed at the normal levels once the I/O resumed. No iSCSI connection drops were observed during the tests. TCP retransmits, pause frames, CRC, and packet error counters were monitored during failure or recovery.

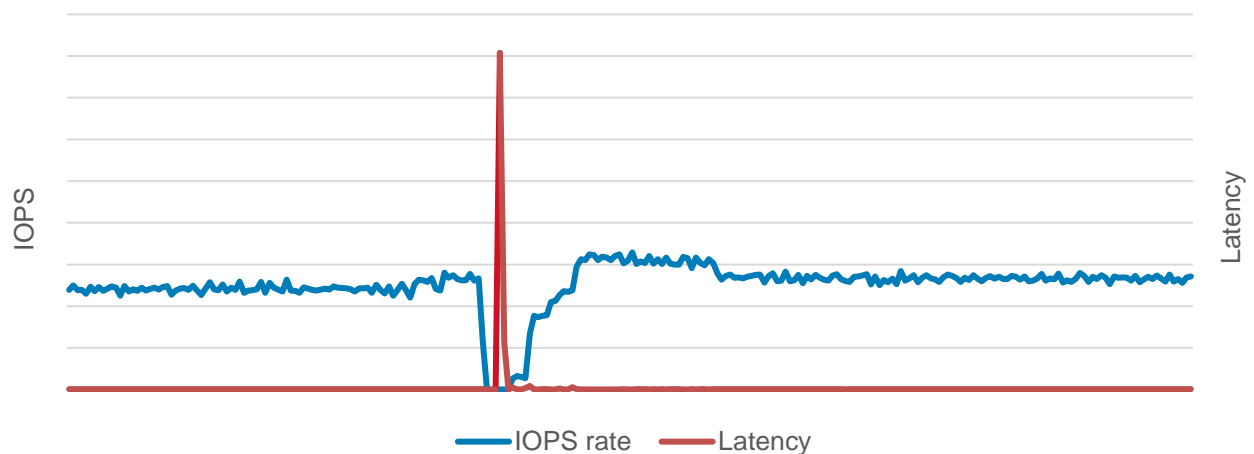


Figure 8 I/O behavior for PS Series and SC Series storage array controller failover high availability test

The results indicated that the shared Windows hosts and iSCSI infrastructure can sustain I/O to both PS Series and SC Series iSCSI storage despite the failures. No adverse impact to functionality or data availability was observed in the test cases.



## 6 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.

### 6.1 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.

### 6.2 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.

### 6.3 Storage

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

## 7

## 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 IP address tables

The MAC and IP addresses of the host and storage ports used for testing are provided in the following tables:

Table 1 MAC and IP addresses of the host and storage ports

Host server	NIC port 1 (MAC/IP)	NIC port 2 (MAC/IP)
Windows 2008 R2	00:10:18:EF:81:B0/ 10.10.10.110	00:10:18:EF:81:B2/ 10.10.10.111
Windows 2012 R2	00:10:18:ED:51:50/ 10.10.10.116	00:10:18:ED:51:52/ 10.10.10.117

SC Series	Port 1 (MAC/IP)	Port 2 (MAC/IP)
Controller 1	00:07:43:09:76:9A/ 10.10.10.31	00:07:43:09:76:9B/ 10.10.10.32
Controller 2	00:07:43:09:78:02/ 10.10.10.41	00:07:43:09:78:03/ 10.10.10.42
	<b>Control port fault domain 1 IP</b>	<b>Control port fault domain 2 IP</b>
	10.10.10.35	10.10.10.36

PS Series	Port 1 (MAC/IP)	Port 2 (MAC/IP)
Array 1	F0:1F:AF:DB:D3:D8/ 10.10.10.20	F0:1F:AF:DB:D3:D9/ 10.10.10.21
Array 2	F0:1F:AF:DB:D3:12/ 10.10.10.22	F0:1F:AF:DB:D3:13/ 10.10.10.23
Array 3	F0:1F:AF:E4:23:07/ 10.10.10.24	F0:1F:AF:E4:23:08/ 10.10.10.25
<b>Group IP: 10.10.10.10</b>		

## B iSCSI host networking best practice settings

In a shared Windows host environment with PS Series and SC Series volumes mapped to a common server, there are several granular networking settings that are recommended as best practices for reliability and performance. These include disk timeout, iSCSI initiator, MPIO, TCP interface, and specific NIC configuration settings. These settings apply to both Windows Server 2008 R2 and Windows Server 2012. Most apply to Windows Server 2012 R2 unless noted.

### B.1 Windows registry settings

Most of the settings provided in this section can be changed using the Windows Registry Editor (regedit.exe). Extreme caution must always be taken when using this tool because inadvertently changing the wrong setting can have serious implications to the Windows operating environment. The settings that are labeled REQUIRED must be set for acceptable connectivity reliability.

**Note:** If the PS Series HIT kit for Microsoft has been run on the Windows server, all of these settings will be set to their best practice setting automatically.

Table 2 Windows registry settings for disk

Setting	Scope/device affected	Default value per MS iSCSI	Shared host value	Setting details	Required/option
Disk Timeout	Global Disk	60 seconds	60	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\DiskTimeout	REQUIRED

Table 3 Windows registry settings for MPIO

Setting	Scope/device affected	Default value per MS iSCSI	Shared host value	Setting details	Required/option
PathRecoveryInterval	MPIO	2 x PDORemovePeriod	60	HKLM\System\CurrentControlSet\Services\mpio\Parameters\PathRecoveryInterval	REQUIRED
UseCustomPathRecoveryInterval	MPIO	0	enabled (1)	HKLM\System\CurrentControlSet\Services\mpio\Parameters\UseCustomPathRecoveryInterval	REQUIRED
PDORemovePeriod	MPIO	25 seconds	120	HKLM\System\CurrentControlSet\Services\mpio\Parameters\PDORemovePeriod	REQUIRED
PathVerifyEnabled	MPIO	disabled (0)	enabled (1)	HKLM\System\CurrentControlSet\Services\mpio\Parameters\PathVerifyEnabled	REQUIRED for MPIO,

Setting	Scope/ device affected	Default value per MS iSCSI	Shared host value	Setting details	Required/ option
					OPTIONAL otherwise
PathVerification Period	MPIO	30	30	HKLM\System\CurrentControlSet\ Services\mpio\Parameters\Path VerificationPeriod	REQUIRED for MPIO, OPTIONAL otherwise
RetryCount	MPIO	3	3	HKLM\System\CurrentControlSet\ Services\mpio\Parameters\Retry Count	REQUIRED
RetryInterval	MPIO	1 second	1	HKLM\System\CurrentControlSet\ Services\mpio\Parameters\Retry Interval	REQUIRED

Table 4 Windows MPCLAIM command changing MPIO default policy

Setting	Scope/devic e affected	Default value per MS MPIO	Shared host value	Setting details	Required/ option
Load balance policy	MPIO/volume	Round robin	Round robin	mpclaim.exe -L -M <0-7> -d "COMPLNTCompellent Vol"	RECOMMENDED

Table 5 Windows registry settings for iSCSI initiator

Setting	Scope/ device affected	Default value per MS iSCSI	Shared host value	Setting details	Required/ option
MaxRequestHold Time	iSCSI initiator	60	90	HKLM\SYSTEM\CurrentControlS et\Control\Class\{4D36E97B- E325-11CE-BFC1- 08002BE10318}\<Instance Number>\Parameters\MaxReque st HoldTime	REQUIRED for Non-MPIO
LinkDownTime	iSCSI initiator	15 seconds	35	HKLM\SYSTEM\CurrentControlS et\Control\Class\{4D36E97B- E325-11CE-BFC1- 08002BE10318}\<Instance Number>\Parameters\LinkDown Time	REQUIRED for Non-MPIO

Setting	Scope/ device affected	Default value per MS iSCSI	Shared host value	Setting details	Required/ option
TCPConnectTime	iSCSI initiator	15	15	HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters\TCPConnectTime	DEFAULT, REQUIRED to be SET
TCPDisconnectTime	iSCSI initiator	15	15	HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters\TCPDisconnectTime	DEFAULT, REQUIRED to be SET
WMIRequestTimeOut	iSCSI initiator	30	30	HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters\WMIRequestTimeout	DEFAULT, REQUIRED to be SET
DelayBetweenReconnect	iSCSI initiator	1 second	1s	HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters\DelayBetweenReconnect	DEFAULT, REQUIRED to be SET
MaxConnectionRetries	iSCSI initiator	-1 (indefinitely)	default	HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters\MaxConnectionRetries	DEFAULT, REQUIRED to be SET
MaxPendingRequests	iSCSI initiator	255	255	HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters\MaxPendingRequests	DEFAULT, REQUIRED to be SET
EnableNOPOut	iSCSI initiator	disabled (0)	enabled (1)	HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters\EnableNO	REQUIRED

Setting	Scope/ device affected	Default value per MS iSCSI	Shared host value	Setting details	Required/ option
				P Out	
MaxTransfer Length	iSCSI initiator	262144 (256KB)		HKLM\SYSTEM\CurrentControlS et\Control\Class\{4D36E97B- E325-11CE-BFC1- 08002BE10318}\<Instance Number>\Parameters\MaxTransf er Length	RECOMMENDED
MaxBurstLength	iSCSI initiator	262144 (256 KB)		HKLM\SYSTEM\CurrentControlS et\Control\Class\{4D36E97B- E325-11CE-BFC1- 08002BE10318}\<Instance Number>\Parameters\MaxBurst Length	RECOMMENDED
FirstBurstLength	iSCSI initiator	65536		HKLM\SYSTEM\CurrentControlS et\Control\Class\{4D36E97B- E325-11CE-BFC1- 08002BE10318}\<Instance Number>\Parameters\FirstBurst Length	RECOMMENDED
MaxRecvData SegmentLength	iSCSI initiator	65536		HKLM\SYSTEM\CurrentControlS et\Control\Class\{4D36E97B- E325-11CE-BFC1- 08002BE10318}\<Instance Number>\Parameters\MaxRecv DataSegmentLength	RECOMMENDED
IPSecConfig Timeout	iSCSI initiator	15		HKLM\SYSTEM\CurrentControlS et\Control\Class\{4D36E97B- E325-11CE-BFC1- 08002BE10318}\<Instance Number>\Parameters\IPSecConf ig Timeout	RECOMMENDED if supported
InitialR2T	iSCSI initiator	0		HKLM\SYSTEM\CurrentControlS et\Control\Class\{4D36E97B- E325-11CE-BFC1- 08002BE10318}\<Instance Number>\Parameters\InitialR2T	RECOMMENDED
ImmediateData	iSCSI Initiator	1 (yes)		HKLM\SYSTEM\CurrentControlS et\Control\Class\{4D36E97B- E325-11CE-BFC1- 08002BE10318}\<Instance	RECOMMENDED

Setting	Scope/ device affected	Default value per MS iSCSI	Shared host value	Setting details	Required/ option
				Number>\Parameters\Immediate Data	
PortalRetryCount	iSCSI initiator	1		HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters\PortalRetr y Count	DEFAULT, REQUIRED to be SET
NetworkReady RetryCount	iSCSI initiator	10		HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters\NetworkR eadyRetryCount	DEFAULT, REQUIRED to be SET
ErrorRecovery Level	iSCSI initiator	2		HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters\ErrorReco veryLevel	DEFAULT, REQUIRED to be SET
AsyncLogout PauseTimeout	iSCSI initiator	10		HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters\AsyncLog outPauseTimeout	DEFAULT, REQUIRED to be SET

Table 6 Windows Registry settings for TCP/IP

Setting	Scope/ device affected	Default value per MS TCPIP	Shared host value	Setting details	Required/ option
TCPAckFrequency (DelayACK): Configurable on Windows 2012 R2 using Microsoft PowerShell™ applets. Refer to appendix B.	TCP/IP	Not present 2	1	HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces\<Interface GUID>\TCPAckFrequency	RECOMMENDED



Setting	Scope/ device affected	Default value per MS TCPIP	Shared host value	Setting details	Required/ option
TCPNoDelay (Nagle's): Configurable on Windows 2012 R2 using Microsoft PowerShell™ applets. Refer to appendix B.	TCP/IP	Not present	1 (disabled)	HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces\<SAN interface GUID>\TcpNoDelay	REQUIRED
WindowSize Scaling	TCP/IP	time stamp + window scaling enabled (3)			REQUIRED

## B.2 Windows Server 2008 R2 NIC best practices

This section provides recommendations based on the performance results and analysis detailed in the document, [Windows Server 2008 R2 NIC Optimization and Best Practices with EqualLogic SAN](#), which also provides a complete list of tested options and default values, as well as instructions on making configuration changes to the storage host. Non-default settings are recommended only when a significant change in performance for one or more workloads is possible, or when a setting is known to provide benefit during network congestion or heavy processor utilization. Only the non-default settings are listed in this section. These settings were tested in the shared Windows host environment, which resulted in nominal to minor performance increases under random and sequential workloads.

## B.2.1 BCM57810 NDIS mode recommended configuration

Based on the performance results and analysis for each workload, the following NIC and OS configuration changes are recommended.

Table 7 Broadcom BCM57810 NDIS mode recommended adapter configuration

Setting	Default value	Recommended value
Flow Control	Auto	Rx & Tx Enabled
Jumbo packet	1514	9614
Receive buffers	0=Auto	3000
TCP Connection Offload	Disabled	Enabled *
Transmit Buffers	0=Auto	5000

\* **Note:** There may be instances with particular combinations of host NICs, SC and PS Series arrays and other copy parameters used in a migration where one of these platforms does not have good performance with TCP Offloads enabled. Therefore, this feature may need to be **disabled** to improve performance.

Table 8 Broadcom BCM57810 NDIS mode recommended Windows Server 2008 R2 TCP configuration

Setting	Default value	Recommended value
Chimney Offload State	Automatic	Enabled

## B.2.2 Intel X520 recommended configuration

Based on the performance results and analysis for each workload, the following NIC configuration changes are recommended.

Table 9 Windows registry settings for iSCSI initiator

Setting	Default value	Recommended value
Jumbo packet	Disabled	Enabled
Receive Buffers	512	4096
Transmit Buffers	512	16384

## B.3 Windows 2012 R2 NIC best practices

This section provides recommendations based on the performance results and analysis detailed in the document, [Windows Server 2012 NIC Optimization and Best Practices with EqualLogic SAN](#), which also provides a complete list of tested options and default values, as well as instructions on making configuration changes to the storage host. Non-default settings are recommended only when a significant performance for one or more workloads is possible, or when a setting is known to provide benefit during network congestion or heavy CPU utilization. Only the non-default settings are listed in this section. These settings were tested in the shared Windows host environment, which resulted in nominal to minor performance increases under random and sequential workloads.

### B.3.1 Broadcom BCM57810 NDIS mode recommended configuration

Based on the performance results and analysis for each workload, the following NIC and OS configuration changes are recommended.

Table 10 Broadcom BCM57810 NDIS mode recommended adapter configuration

Setting	Default value	Recommended value
Jumbo packet	1514	9614
Receive buffers	0=Auto	3000
TCP Connection Offload	Disabled	Enabled *
Transmit Buffers	0=Auto	5000

\* **Note:** There may be instances with particular combinations of host NICs, SC and PS Series arrays and other copy parameters used in a migration where one of these platforms does not have good performance with TCP Offloads enabled. Therefore, this feature may need to be **disabled** to improve performance.

Table 11 Broadcom BCM57810 NDIS mode recommended Windows Server 2012 TCP configuration

Setting	Default value	Recommended value
Chimney Offload State	Disabled	Enabled

### B.3.2 Intel X520 recommended configuration

Based on the performance results and analysis for each workload in the document, [Windows Server 2012 NIC Optimization and Best Practices with EqualLogic SAN](#), the following NIC configuration changes are recommended.

Table 12 Intel X520 recommended adapter configuration

Setting	Default value	Recommended value
---------	---------------	-------------------

Jumbo packet	Disabled	Enabled
Receive Buffers	512	4096
Transmit Buffers	512	16384

## B.4 NIC-specific settings summary

For a shared PS Series and SC Series Windows host, it is recommended that Jumbo frames and flow control be enabled for both the Broadcom 57810 and the Intel® X520. If not using the Broadcom iSCSI Offload Engine, receive and transmit buffers should also be maximized.

When using the Broadcom BCM57810, TCP Offload Engine should be considered for its ability to decrease processor utilization and also to lower retransmission rates in congested networks.

Appendix D in the document, [Windows Server 2012 and Windows Server 2012 R2 NIC Optimization and Best Practices with Dell PS Series](#), contains detailed instructions on how to make recommended changes within the profiles.

**Note:** There may be instances with particular combinations of host NICs, SC and PS Series arrays and other copy parameters used in a migration where one of these platforms does not have good performance with TCP Offloads enabled. Therefore, this feature may need to be **disabled** to improve performance.

## B.5 iSCSI-specific settings

**Note:** The settings that follow apply to all Windows servers configured to use iSCSI.

**Enable RFC1323 time stamps (TCP high performance extensions)** to prevent sequence number wrap under high load, known as Prevention Against Wrapped Sequence (PAWS) or something similar. High load iSCSI connections are prone to this issue, particularly at 10GbE.

```
netsh int tcp set global timestamps=enabled
```

**Disable Nagle's algorithm:** To disable delayed ACK and Nagle's algorithm, create the following entries for each SAN interface subkey in the Windows Server registry:

```
HKEY_LOCAL_MACHINE \ SYSTEM \ CurrentControlSet \ Services \ Tcpip \ Parameters  
 \ Interfaces \ <SAN interface GUID>
```

Entries:

```
TcpAckFrequency  
TcpNoDelay
```

Value type:

```
REG_DWORD, number
```

Value to disable:

1

## B.6 Windows 2012 Server, R2, and Hyper-V

Microsoft implements network transport layer profiles and filters and that allow you to customize and change TCP settings within Windows 2012 Server, Windows Server 2012 R2, and Hyper-V. All modifications of the profiles and filters are done using [Cmdlets specific for TCP/IP in PowerShell](#).

These are the two key command applets used to check and modify the DatacenterCustom TCP profile parameters: [Get-NetTCPSetting](#) and [Set-NetTCPSetting](#).

These are the two key command applets we will use to check, create, and associate the network transport filters with the DatacenterCustom TCP profile: [Get-NetTransportFilter](#) and [New-NetTransportFilter](#).

## C Additional resources

### C.1 Technical support and resources

[Dell.com/support](http://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.

### C.2 Related resources

See the following referenced or recommended Dell EMC publications:

- [\*Dell PS Series Configuration Guide\*](#)
- [\*Dell Storage Compatibility Matrix\*](#)
- [\*Switch Configuration Guides for PS Series or SC Series\*](#)
- [\*Dell Storage SC Series: Microsoft Multipath I/O Best Practices\*](#)