

Dell EMC PowerStore

Host Configuration Guide

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Additional Resources.....	6
Chapter 1: Introduction.....	7
Purpose.....	7
Chapter 2: Best Practices for Storage Connectivity.....	8
General SAN Guidelines.....	8
Recommended Zoning Configurations.....	8
Using LUN 0 with PowerStore.....	8
Fibre Channel SAN Guidelines.....	9
Recommended Configuration Values Summary.....	9
Fibre Channel SAN Configuration Limitations.....	9
Minimal Zoning Configuration for Fibre Channel.....	9
Recommended Zoning Configuration for Fibre Channel.....	10
iSCSI SAN Guidelines.....	13
Recommended Configuration Values Summary.....	13
Minimal Subnetting Configuration for iSCSI.....	13
Recommended Configuration for iSCSI.....	14
Chapter 3: Host Configuration for VMware vSphere ESXi.....	18
Chapter Scope.....	18
Recommended Configuration Values Summary.....	18
Fibre Channel Configuration.....	19
Pre-Requisites.....	19
NVMe over Fibre Channel Configuration on ESXi Hosts.....	20
iSCSI Configuration.....	21
Pre-Requisites.....	21
Software iSCSI Adapter Configuration.....	21
Jumbo Frames.....	22
Delayed ACK.....	22
Login Timeout.....	22
No-Op Interval.....	23
vStorage API for System Integration (VAAI) Settings.....	23
Confirming that VAAI is Enabled on the ESXi Host.....	23
Setting the Maximum I/O.....	24
Confirming UNMAP Priority.....	24
Configuring VMware vSphere with PowerStore Storage in a Multiple Cluster Configuration.....	24
Multipathing Software Configuration.....	25
Configuring Native Multipathing.....	25
Post-Configuration Steps.....	28
AppsOn: Virtual Machine Compute and Storage Collocation Rules for PowerStore X Clusters.....	28
Disk Formatting.....	28
Presenting PowerStore Volumes to the ESXi Host.....	29
Creating a File System.....	29

VMware Paravirtual SCSI Controllers.....	29
Virtual Disk Provisioning.....	29
Virtual Machine Guest OS Settings.....	30
Chapter 4: Host Configuration for Microsoft Windows.....	31
Recommended Configuration Values Summary.....	31
Fibre Channel Configuration.....	32
Pre-Requisites.....	32
iSCSI Configuration.....	32
Pre-Requisites.....	32
Multipathing Software Configuration.....	32
Configuring Native Multipathing Using Microsoft Multipath I/O (MPIO).....	33
Post-Configuration Steps - Using the PowerStore system.....	34
Presenting PowerStore Volumes to the Windows Host.....	34
Creating a File System.....	34
Chapter 5: Host Configuration for Linux.....	35
Recommended Configuration Values Summary.....	35
Fibre Channel Configuration.....	36
Pre-Requisites.....	36
NVMe-FC Configuration on Linux Hosts.....	36
iSCSI Configuration.....	37
Pre-Requisites.....	37
Configuring the PowerStore Cluster Disk Device with iSCSI.....	37
Configuring the PowerStore Cluster Disk Device with iSCSI - Single Network Subnet Support.....	38
Multipathing Software Configuration.....	39
Configuring Linux Native Multipathing.....	39
Configuring Device Mapper Multipathing (DM-multipathing) with NVMe-FC.....	40
Configuring Oracle ASM.....	41
Post-Configuration Steps - Using the PowerStore system.....	42
Presenting PowerStore Cluster Volumes to the Linux Host.....	42
Partition Alignment in Linux.....	42
Creating a File System.....	42
Chapter 6: Host Configuration for AIX.....	44
Recommended Configuration Values Summary.....	44
Fibre Channel Configuration.....	45
Pre-Requisites.....	45
Queue Depth.....	45
Fast I/O Failure for Fibre Channel Devices.....	46
Dynamic Tracking.....	46
Fibre Channel Adapter Device Driver Maximum I/O Size.....	47
Dell EMC AIX ODM Installation.....	47
Dell EMC AIX ODM Installation Requirements.....	47
Chapter 7: Host Configuration for Solaris.....	49
Recommended Configuration Values Summary.....	49
Fibre Channel Configuration.....	50
Pre-Requisites.....	50

Queue Depth.....	51
Solaris Host Parameter Settings.....	51
Configuring Solaris native multipathing.....	51
Host storage tuning parameters.....	52
Post configuration steps - using the PowerStore system.....	55
Partition alignment in Solaris.....	55
Chapter 8: Host Configuration for HP-UX.....	56
Recommended Configuration Values Summary.....	56
Fibre Channel Configuration.....	57
Pre-Requisites.....	57
HP-UX Host Parameter Settings.....	57
Maximum Transfer Length.....	57
Multipathing Software Configuration.....	58
Configuring Native Multipathing Using HP-UX Multipath I/O (MPIO).....	58
Post-Configuration Steps - Using the PowerStore System.....	58
Presenting PowerStore Volumes to the HP-UX Host.....	58
Creating a file system.....	59
Appendix A: Considerations for Boot from SAN with PowerStore.....	60
Consideration for Boot from SAN with PowerStore.....	60

As part of an improvement effort, revisions of the software and hardware are periodically released. Some functions that are described in this document are not supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information about product features. Contact your service provider if a product does not function properly or does not function as described in this document.

Where to get help

Support, product, and licensing information can be obtained as follows:

- **Product information**

For product and feature documentation or release notes, go to the PowerStore Documentation page at <https://www.dell.com/powerstoredocs>.

- **Troubleshooting**

For information about products, software updates, licensing, and service, go to <https://www.dell.com/support> and locate the appropriate product support page.

- **Technical support**

For technical support and service requests, go to <https://www.dell.com/support> and locate the **Service Requests** page. To open a service request, you must have a valid support agreement. Contact your Sales Representative for details about obtaining a valid support agreement or to answer any questions about your account.

Introduction

Topics:

- [Purpose](#)

Purpose

This document provides guidelines and best practices on attaching and configuring external hosts to PowerStore systems, or in conjunction with other storage systems. It includes information on topics such as multipathing, zoning, and timeouts. This document may also include references to issues found in the field and notify you on known issues.

Regarding ESXi hosts, this document provides guidelines only for configuring ESXi hosts that are connected externally to PowerStore. For configuring an internal ESXi host in PowerStore X model, refer to the *PowerStore Virtualization Guide*.

For further host connectivity best practices in conjunction to other Dell EMC storage systems, also refer to the *E-Lab Host Connectivity Guides*. For details, refer to the *E-Lab Interoperability Navigator* at <https://elabnavigator.emc.com>.

Best Practices for Storage Connectivity

This chapter contains the following topics:

Topics:

- [General SAN Guidelines](#)
- [Fibre Channel SAN Guidelines](#)
- [iSCSI SAN Guidelines](#)

General SAN Guidelines

This section provides general guidelines for storage connectivity.

NOTE: This document describes mainly the storage-specific recommendations for PowerStore. It is recommended to always consult with the OS documentation for the up-to-date guidelines specific for the used operating system.

NOTE: In hosts running a hypervisor, such as VMware ESXi, Microsoft Hyper-V or any clustering software, it is important to ensure that the logical unit numbers of PowerStore volumes are consistent across all hosts in the hypervisor cluster. Inconsistent LUNs may affect operations such as VM online migration or VM power-up.

Recommended Zoning Configurations

Consider these recommendations for zoning configuration:

- The optimal number of paths depends on the operating system and server information. To avoid multipathing performance degradation, do not use more than 8 paths per device per host. It is recommended to use 4 paths.
- Keep a consistent link speed and duplex across all paths to the PowerStore cluster per single host or a cluster of hosts.
- With a multi-appliance cluster, it is highly advised to zone the host to as many appliances as possible, in order to achieve best load distribution across the cluster. Make sure to keep the minimum/optimal zoning recommendations for each appliance.

NOTE: A multi-appliance cluster is not designed to provide better resiliency, but to provide better load balance. In order to perform volume migration between appliances, a host must be zoned to both appliances.

Using LUN 0 with PowerStore

The PowerStore system exposes an SCD (Storage Array Controller Device) by default. The device is exposed with a LUN ID #0. A user may choose to override the SCD with a real storage device (e.g. volume or clone) by setting the device LUN ID to #0. Doing so may require the user to manually force a host re-scan to discover that device. For instructions on forcing a host re-scan, refer to the OS vendor documentation.

Fibre Channel SAN Guidelines

This section describes the best practices for attaching hosts to a PowerStore cluster in a highly available resilient and optimal Fibre Channel SAN.


Recommended Configuration Values Summary

The following table summarizes the recommended configuration values related to Fibre Channel SAN.

Validation	Impact	Severity	Refer to Section
Use two separate fabrics.	Redundancy	Mandatory	Recommended Zoning Configuration for Fibre Channel
Each host should be zoned to both nodes of each appliance.	Redundancy	Mandatory	Minimal Zoning Configuration for Fibre Channel
Maximum number of paths per appliance per volume per host: 8	Performance	Warning	Recommended Zoning Configuration
Recommended number of paths per volume per host: 4	Performance	Warning	Recommended Zoning Configuration
Link speed should be consistent across all paths to the PowerStore cluster per single host or a cluster of hosts.	Performance	Warning	Recommended Zoning Configuration
Balance the hosts between the nodes of the appliance to provide a distributed load across all target ports.	Performance	Recommended	Recommended Zoning Configuration
Maximum ISL Hops: 2	Performance	Recommended	Recommended Zoning Configuration for Fibre Channel

Fibre Channel SAN Configuration Limitations

Review these limitations when configuring hosts with PowerStore storage using either SCSI Fibre Channel (SCSI-FC) or NVMe over Fibre Channel (NVMe-FC):


 **NOTE:** NVMe-FC is supported with PowerStore OS version 2 or later.


- Direct-attach storage (DAS) is supported for SCSI-FC.
- Direct attach storage is not supported for NVMe-FC.
- NVMe-FC requires NPIV to be enabled at the switch level (NPIV is enabled by default on PowerStore FC ports).
- If NPIV is disabled on the switch and the administrator wants to enable NPIV on the port for NVMe-FC to work, it is required to disable and then reenable the port at the switch.

Minimal Zoning Configuration for Fibre Channel

To prevent a host path failure due to a single node failure, make sure to maintain redundancy.

- A PowerStore appliance contains two nodes.
- A host must be zoned at minimum with one path to each node for redundancy.

 **NOTE:** The diagrams throughout this chapter illustrate possible implementations of these guidelines. Other possible implementations are not illustrated.




1. Host
2. PowerStore Appliance
3. Fibre Channel Switch
4. Node

Recommended Zoning Configuration for Fibre Channel

Consider the following recommendations when setting up a Fibre Channel SAN infrastructure.


- You can zone the host to 1-4 appliances. It is recommended to zone the host to as many appliances as possible to allow volume migration to and from all appliances.
- Use two separate fabrics. Each fabric should be on a different physical FC switch for resiliency.
- Balance the hosts between the two nodes of the appliance.
- The PowerStore cluster can be shipped with various extension modules for Fibre Channel. If your PowerStore cluster contains more than one extension I/O module per node, distribute the zoning among all I/O modules for highest availability and performance.
- Use single initiator zoning scheme: Utilize single-initiator per multiple-target (1:many) zoning scheme when configuring zoning with a PowerStore cluster.
- Host I/O latency can be severely affected by FC SAN congestion. Minimize the use of ISLs by placing the host and storage ports on the same physical switch. When this is not possible, ensure that there is sufficient ISL bandwidth and that both the Host and PowerStore cluster interfaces are separated by no more than 2 ISL hops.
- For more information on zoning best practices, refer to the *Networked Storage Concepts and Protocols techbook*: https://support.emc.com/docu33913_Networked-Storage-Concepts-and-Protocols-TechBook.pdf
- PowerStore supports direct attachment of hosts to the appliances.

The following diagram describes a simple Fibre Channel connectivity with a **single** PowerStore appliance:




1. Host
2. PowerStore Appliance
3. FC SAN Switch
4. Node

The following diagram describes a simple Fibre Channel connectivity with **two (2)** PowerStore appliances:




1. Host
2. PowerStore Appliance
3. FC SAN Switch
4. Node

The following diagram describes a simple Fibre Channel connectivity with **three (3)** PowerStore appliances:




- 1. Host
- 2. PowerStore Appliance
- 3. FC SAN Switch
- 4. Node

The following diagram describes a simple Fibre Channel connectivity with **four (4)** PowerStore appliances:



- 1. Host
- 2. PowerStore Appliance

3. FC SAN Switch
4. Node


 **NOTE:** Refer to your Fibre Channel switch user manual for implementation instructions.

iSCSI SAN Guidelines

This section details the best practices for attaching hosts to a PowerStore cluster in a highly-available, resilient and optimal iSCSI SAN.

Recommended Configuration Values Summary

The following table summarizes the recommended variables related to iSCSI SAN:

Validation	Impact	Severity	Refer to Section
Each host should be connected to both nodes of each appliance.	Redundancy	Mandatory	Minimal Subnetting Configuration for iSCSI
Maximum number of paths per volume per host: 8	Performance	Warning	Recommended Zoning Configuration
Recommended number of paths per volume per host: 4	Performance	Warning	Recommended Zoning Configuration
Link speed should be consistent across all paths to the PowerStore cluster.	Performance	Warning	Recommended Zoning Configuration
Duplex setting should be consistent across all paths to the PowerStore cluster per single host or a cluster of hosts.	Performance	Warning	Recommended Zoning Configuration
If Jumbo Frames are required for iSCSI traffic or management traffic (Vvol), make sure that all ports (servers, switches and system) are configured with the same MTU value.	Performance	Warning	Recommended Configuration for iSCSI
Enable the TCP Offloading Engine (TOE) on the host iSCSI interfaces.	Performance	Warning	Recommended Configuration for iSCSI
Balance the hosts across the target ports of the appliances to provide a distributed load across all target ports.	Performance	Recommended	Recommended Zoning Configuration
Use dedicated NICs or iSCSI HBAs for PowerStore cluster iSCSI connection.  NOTE: Avoid partitioning the iSCSI interface.	Performance	Recommended	Recommended Configuration for iSCSI
With PowerStore OS version 2 (or later), up to 32 network subnets are supported for the iSCSI target portals, but only up to 8 network subnets per physical port. With PowerStore OS version 1, only a single network subnet is supported.	Performance	Normal	Recommended Configuration for iSCSI

Minimal Subnetting Configuration for iSCSI

To prevent host path failure due to a single node failure, make sure to maintain redundancy.

- A PowerStore Appliance contains two nodes.

- A host must be connected at minimum with one path to each node for redundancy.

Recommended Configuration for iSCSI


Consider the following recommendations when setting up an iSCSI SAN infrastructure.

- External hosts can be attached via iSCSI to a PowerStore cluster via either the embedded 4-port card or via a SLIC:
 - Hosts connected via the first two ports of the 4-port card, will be connected using ToR switches (also used for PowerStore internal communication). With this configuration, it is recommended to use a dedicated VLAN, and if not possible, use a separate subnet/network for iSCSI storage.
 - For hosts connected using any other iSCSI port (that is, not the first two ports), use either dedicated iSCSI switches or a dedicated VLAN for iSCSI storage.
- The PowerStore cluster can be shipped with various extension modules for iSCSI. If your PowerStore cluster contains more than one extension I/O module per node, distribute the connections among all I/O modules for highest availability and performance.
- You can connect the host to 1-4 appliances. It is recommended to connect the host to as many appliances as possible to allow volume migration to and from all appliances.
- iSCSI switch recommendations
 - Use nonblocking switches
 - Use enterprise grade switch
 - Utilize at minimum 10GbE interfaces
- It is recommended to use dedicated NICs or iSCSI HBAs for PowerStore cluster iSCSI and not to partition the iSCSI interface (that is, disable NIC Partitioning - NPAR).
- Enable the TCP Offloading Engine (TOE) on the host iSCSI interfaces, to offload the TCP packet encapsulation from the host's CPU to the NIC or iSCSI HBA, and free up CPU cycles.
- When using PowerStore iSCSI, it is recommended to use interfaces individually rather than using NIC Teaming (Link Aggregation), to combine multiple interfaces into a single virtual interface.
- With PowerStore OS version 2 (or later), up to 32 subnets are supported for the iSCSI target portals, but only up to 8 network subnets are supported per physical port. With PowerStore OS version 1, only a single network subnet is supported.
- If Jumbo Frames are required for iSCSI traffic or management traffic (Vvol), make sure that all ports (servers, switches and system) are configured with the same MTU value.

NOTE: VMware requires setting Jumbo Frames at the virtual switch and VMKERNEL level.

NOTE: Refer to your iSCSI switch user manual for instructions on the implementations.

The following diagram describes a simple iSCSI connectivity with a **single** PowerStore appliance:




1. Host

2. PowerStore Appliance
3. ToR/iSCSI Switch
4. Node

NOTE: For detailed information on connecting the PowerStore appliance to the ToR/iSCSI switch, refer to the *PowerStore Network Planning Guide* and the *Network Configuration Guide for Dell PowerSwitch Series*.


The following diagram describes a simple iSCSI connectivity with **two (2)** PowerStore appliances:



1. Host
2. PowerStore Appliance
3. ToR/iSCSI Switch
4. Node

NOTE: For detailed information on connecting the PowerStore appliance to the ToR/iSCSI switch, refer to the *PowerStore Network Planning Guide* and the *Network Configuration Guide for Dell PowerSwitch Series*.

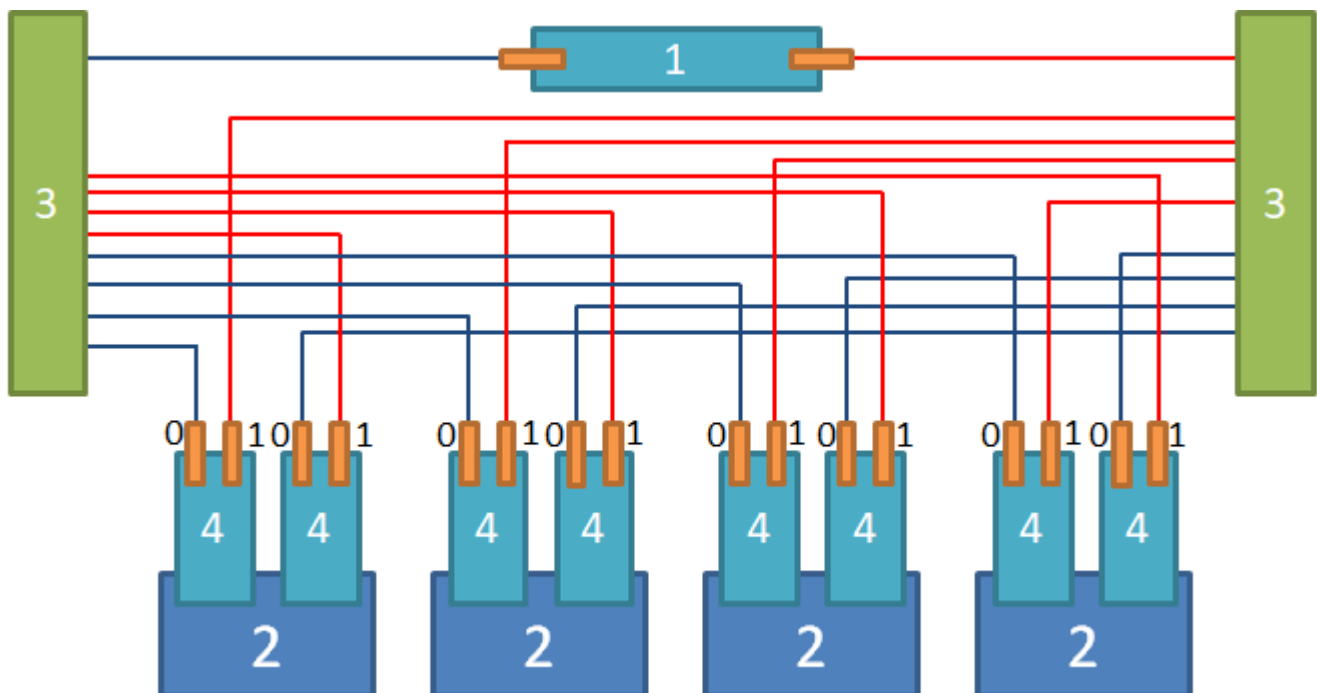
The following diagram describes a simple iSCSI connectivity with **three (3)** PowerStore appliances:



- 1. Host
- 2. PowerStore Appliance
- 3. ToR/iSCSI Switch
- 4. Node

NOTE: For detailed information on connecting the PowerStore appliance to the ToR/iSCSI switch, refer to the *PowerStore Network Planning Guide* and the *Network Configuration Guide for Dell PowerSwitch Series*.


The following diagram describes a simple iSCSI connectivity with **four (4)** PowerStore appliances:





- 1. Host
- 2. PowerStore Appliance

3. ToR/iSCSI Switch

4. Node

 **NOTE:** For detailed information on connecting the PowerStore appliance to the ToR/iSCSI switch, refer to the *PowerStore Network Planning Guide* and the *Network Configuration Guide for Dell PowerSwitch Series*.

 **NOTE:** Make sure to connect port 0 of each node to a different iSCSI switch.

 **NOTE:** Refer to your iSCSI switch user manual for implementation instructions.

Host Configuration for VMware vSphere ESXi

This chapter contains the following topics:

Topics:


- [Chapter Scope](#)
- [Recommended Configuration Values Summary](#)
- [Fibre Channel Configuration](#)
- [iSCSI Configuration](#)
- [vStorage API for System Integration \(VAAI\) Settings](#)
- [Setting the Maximum I/O](#)
- [Confirming UNMAP Priority](#)
- [Configuring VMware vSphere with PowerStore Storage in a Multiple Cluster Configuration](#)
- [Multipathing Software Configuration](#)
- [Post-Configuration Steps](#)

Chapter Scope

This chapter provides guidelines only for configuring ESXi hosts that are connected externally to PowerStore. For configuring an internal ESXi host in PowerStore, refer to the *PowerStore Virtualization Guide* document.

Recommended Configuration Values Summary

The following table summarizes all used and recommended variables and their values when configuring hosts for VMware vSphere.

 **NOTE:** Unless indicated otherwise, use the default parameters values.

Validation	Impact	Severity	Refer to Section
To clarify the above note for using default parameter settings unless stated otherwise in this chapter, make sure that the following are set per the default OS setting: <ul style="list-style-type: none"> • LUN and HBA queue depth • HBA timeout 	Stability & Performance	Recommended	For further details, refer to OS and HBA documentation.
Specify ESXi as the operating system for each defined host.	Serviceability	Mandatory	Presenting PowerStore Volumes to the ESXi Host
ESXi configuration: VAAI should be enabled, therefore: <ul style="list-style-type: none"> • <code>DataMover.HardwareAcceleratedMove=1</code> • <code>DataMover.HardwareAcceleratedInit=1</code> • <code>VMFS3.HardwareAcceleratedLocking=1</code> 	Performance	Mandatory	Confirming that VAAI is Enabled on the ESXi Host
ESXi configuration: <code>Disk.DiskMaxIOSize = 1024</code>	Stability & Performance	Mandatory	Setting Maximum I/O

Validation	Impact	Severity	Refer to Section
ESXi configuration: Keep the UNMAP priority for the host at the lowest possible value (default value for ESXi 6.5).	Stability & Performance	Mandatory	Confirming UNMAP Priority
Path selection policy: VMW_PSP_RR	Stability & Performance	Mandatory	Configuring vSphere Native Multipathing
Alignment: Guest OS virtual machines should be aligned.	Storage efficiency & Performance	Warning	Disk Formatting
iSCSI configuration: Configure end-to-end Jumbo Frames.	Performance	Recommended	Jumbo Frames
iSCSI configuration: Disable Delayed ACK on ESXi.	Stability	Recommended	Delayed ACK
iSCSI configuration: Adjust LoginTimeout to 30.	Stability	Recommended	Login Timeout
iSCSI configuration: Adjust NoopInterval to 5.	Stability	Recommended	No-Op Interval
Path switching: Switch for every I/O.	Performance	Recommended	Configuring vSphere Native Multipathing
Virtual Disk Provisioning: Use thin provisioned virtual disks.	Performance	Recommended	Virtual Machine Formatting
Virtual machine configuration: Configure virtual machines with Paravirtualized SCSI controllers.	Stability & Performance	Recommended	VMware Paravirtual SCSI controllers
RDM volumes: In Guest OS, span RDM volumes used by the virtual machine across SCSI controllers.	Performance	Recommended	Virtual Machine Guest OS Settings

NOTE: For information on virtualization and Virtual Volumes, refer to the *PowerStore Virtualization Guide* and the *Best Practices Guide*.

NOTE: As noted in Dell EMC KB article 000126731 (<https://www.dell.com/support/kbdoc/en-il/000126731>), when using vSphere v6.7 there is a known issue relating to VMFS deadlock. To resolve, install the latest vSphere version (v6.7 P02) that includes a fix for this issue. For further details, refer to fix PR 2512739 in the release notes of this vSphere version (<https://docs.vmware.com/en/VMware-vSphere/6.7/rn/esxi670-202004002.html>).

Fibre Channel Configuration

This section describes the recommended configuration that should be applied when attaching hosts to PowerStore cluster using Fibre Channel .

NOTE: This section applies only to Fibre Channel. If you are using only iSCSI, go to [iSCSI HBA Configuration](#).

NOTE: Review [Fibre Channel SAN Guidelines](#) before you proceed.


Pre-Requisites

Before installing HBAs on an ESXi host, the following pre-requisites should be met:

- Refer to the E-Lab Interoperability Navigator (<https://elabnavigator.emc.com>) for supported FC HBA models and drivers.
- Verify that all HBAs have supported driver, firmware and BIOS versions.
- Verify that all HBAs BIOS settings are configured according to E-Lab recommendations.

NVMe over Fibre Channel Configuration on ESXi Hosts

For details on NVMe over Fibre Channel (NVMe-FC) configuration with ESXi hosts, see the *VMware vSphere Storage* document for the vSphere version running on the ESXi hosts (the *About VMware NVMe storage* section).

 **NOTE:** NVMe-FC on an ESXi hosts connected to PowerStore is currently not supported with Raw Device Mapping files (RDMs) and VMware vSphere Virtual volumes (vVols).

Setting the ESXi Host NVMe Qualified Name

Prerequisites

When configuring NVMe-FC on an ESXi host that is connected to PowerStore, ensure to set the host subsystem NVMe Qualified Name (Host NQN) to a valid name according to spec restrictions. For NVMe-FC, the host NQN should be generated based on the WWN.

Steps

1. Connect to the ESXi host as `root`.
2. Run the following `esxcli` command for the Host NQN to be automatically generated according to spec during the ESXi host reboot:
`esxcli system module parameters set -m vmknvme -p vmknvme_hostnqn_format=0`
3. Reboot the ESXi host.
4. Run the `esxcli nvme info get` command to confirm that the Host NQN was generated correctly.

```
# esxcli nvme info get
Host NQN:
nqn.2014-08.org.nvmexpress:uuid:5ffefdd7-0a51-3700-0b16-001e6792303a
```

Setting up NVMe HBAs

For further details on HBA setup, see NVMe HBA documentation.

Setting Up Marvell NVMe HBAs

Follow these steps to setup Marvell (QLE) NVMe HBA for ESXi:

Steps

1. Connect to the ESXi host as `root`.
2. Run the following `esxcli` command:
`esxcli system module parameters set -p ql2xnvmesupport=1 -m qlnativefc`
3. Reboot the ESXi host.

Setting Up LPe3200x NVMe HPA

Follow these steps to setup LPe3200x NVMe HPA for ESXi:

Steps

1. Connect to the ESXi host as `root`.
2. Run the following `esxcli` command:
`esxcli system module parameters set -m lpfc -p lpfc_enable_fc4_type=3`
3. Reboot the ESXi host.

iSCSI Configuration

This section describes the recommended configuration that should be applied when attaching hosts to a PowerStore cluster using iSCSI.

NOTE: This section applies only for iSCSI. If you are using only Fibre Channel with vSphere and PowerStore, go to [Fibre Channel HBA Configuration](#).

NOTE: Be sure to review [iSCSI SAN Guidelines](#) before you proceed.

Pre-Requisites

Before configuring iSCSI, the following pre-requisites should be met:

- Follow the VMware recommendations for installation and setup of the appropriate NIC/iSCSI HBA for your system.
- It is recommended to install the latest driver version (patch), as described in the VMware support site for each specific NIC/iSCSI HBA.
- Refer to the E-Lab interoperability Navigator (<https://elabnavigator.emc.com>) for supported NIC/iSCSI HBA models and drivers.

Software iSCSI Adapter Configuration

This topic provides a high-level procedure for configuring the iSCSI adapter on ESXi hosts connected to PowerStore.

About this task


For further details on this topic, refer to the VMware vSphere documentation for the ESXi version installed on the ESXi hosts. As noted in the [iSCSI SAN Guidelines](#) section, PowerStore system supports only a single subnet for all iSCSI communication.

Steps

1. Create a single vSwitch comprised of two uplink physical ports.
2. Create two VMkernel ports on the same subnet as the storage cluster iSCSI portals (the communication must not be routable).

Example:

- iSCSI-A-port0 1.1.1.1/24
- iSCSI-A-port1 1.1.1.2/24
- iSCSI-B-port0 1.1.1.3/24
- iSCSI-B-port1 1.1.1.4/24
- vmk1 1.1.1.10/24
- vmk2 1.1.1.11/24



3. Make sure that both VMkernel interfaces are attached to the same vSwitch.
4. Override the default Network Policy for iSCSI. For details refer to VMware vSphere documentation. For example, with ESXi 6.5, refer to <https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.storage.doc/GUID-9C90F3F6-6095-427A-B20C-D46531E39D32.html>
5. Configure port binding for each VMkernel interface as described in the VMware vSphere documentation. For example, with ESXi 6.5, refer to <https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.storage.doc/GUID-D9B862DF-476A-4BCB-8CA5-DE6DB2A1A981.html>

NOTE: Configuring port bindings will cause all traffic to flow from all bound interfaces. Without port binding, only a single session will be used whereas with port binding, four sessions will be used (assuming two VMkernel ports and two target iSCSI portals). For more information, refer to <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-4C19E34E-764C-4069-9D9F-D0F779F2A96C.html>

Jumbo Frames

Configure end-to-end Jumbo Frames for optimal performance.

When using iSCSI with ESXi hosts and PowerStore, it is recommended to configure end-to-end Jumbo Frames (MTU=9000) for optimal performance. Ethernet frames will be larger than the standard frame size of 1500 bytes (for IPv4) or 1280 bytes (for IPv6).

NOTE: When using Jumbo Frames, make sure that all ports (ESXi server, switch/es and storage) are configured with the correct MTU value. With VMware ESXi, the correct MTU size must be set on the virtual switch as well.

For further details on configuring Jumbo Frames with iSCSI on ESXi, refer to VMware KB article 1007654 on the VMware website (<https://kb.vmware.com/kb/1007654>).

Delayed ACK

For optimal traffic, it is recommended to disable Delayed ACK on ESXi.

For optimal iSCSI traffic between the ESXi hosts and PowerStore, especially during periods of network congestion, it is recommended to disable Delayed ACK on ESXi. By disabling Delayed ACK, the ESXi host would send an ACK acknowledgment segment for every received data segment (rather than delaying the sending of ACK acknowledgment segments, while receiving a stream of TCP data segments).

For further details on the Delayed ACK parameter and how to disable it using the vSphere Client, refer to VMware KB article 1002598 on the VMware website (<https://kb.vmware.com/kb/1002598>).

NOTE: The recommended method for configuring the Delayed ACK setting is per discovered iSCSI target. This allows to disable Delayed ACK only for PowerStore iSCSI targets.

Login Timeout

Follow these steps to set the iSCSI login timeout.

About this task

When establishing an iSCSI session between the initiator and target, the login timeout setting controls how long the ESXi host attempts to log in to the iSCSI target before failing the login and retrying. The default setting for `LoginTimeout` is 5. For example, by default an iSCSI session ceases retries after 20 seconds (5 times the `LoginRetryMax` setting, which is set by default to 4).

To optimize the iSCSI session behavior with PowerStore and to better handle periods of network disruptions, it is recommended to adjust `LoginTimeout` to 30.

The following steps describe how to adjust `LoginTimeout`, using command-line.

Steps

1. Connect to the ESXi host as `root`.
2. Run the following command:

```
esxcli iscsi adapter param set -A adapter_name -k LoginTimeout -v value_in_sec
```

Example

For example:

```
esxcli iscsi adapter param set -A vmhba64 -k LoginTimeout -v 30
```

No-Op Interval

Follow these steps to set the iSCSI No-Op interval.

About this task

The noop iSCSI settings (*NoopInterval* and *NoopTimeout*) are used to determine if a path is dead, when it is not the active path. iSCSI will passively discover if this path is dead using *NoopTimeout*. This test is carried out on non-active paths every *NoopInterval*, and if no response is received by *NoopTimeout*, the path is marked as DEAD.

The default setting for *NoopInterval* is 10. To optimize the iSCSI session behavior with PowerStore, it is recommended to adjust *NoopInterval* to 5. This would trigger an iSCSI path failover following a network disconnection sooner, prior to a command timeout.

The following steps describe how to adjust *NoopInterval* using command-line:

Steps

1. Connect to the ESXi host as `root`.
2. Run the following command:

```
esxcli iscsi adapter param set -A adapter_name -k NoopInterval -v value_in_sec
```

Example

```
esxcli iscsi adapter param set -A vmhba64 -k NoopInterval -v 5
```

vStorage API for System Integration (VAAI) Settings

PowerStore Storage cluster fully supports VAAI. VAAI must be enabled on the ESXi host before using PowerStore.

VAAI is an API that offloads operations such as virtual machine provisioning, storage cloning, and space reclamation to storage clusters that supports VAAI. PowerStore clusters fully support VAAI.

To ensure optimal performance of PowerStore storage from vSphere, VAAI must be enabled on the ESXi host before using PowerStore storage from vSphere. Failing to do so may expose the PowerStore cluster to the risk of datastores becoming inaccessible to the host.

This section describes the necessary settings for configuring VAAI for PowerStore storage.

Confirming that VAAI is Enabled on the ESXi Host


Follow these instructions to confirm that VAAI is enabled on the ESXi host.

About this task

When using vSphere ESXi version 6.5 and above, VAAI is enabled by default. Before using the PowerStore storage, confirm that VAAI features are enabled on the ESXi host.

Steps

1. Verify that the following parameters are enabled (that is, set to 1):
 - `DataMover.HardwareAcceleratedMove`
 - `DataMover.HardwareAcceleratedInit`
 - `VMFS3.HardwareAcceleratedLocking`
2. If any of the above parameters are not enabled, click the **Edit** icon and click **OK** to adjust them.

 **NOTE:** These settings will enable ATS-only on supported VMFS Datastores, as noted in VMware KB# 1021976 (<https://kb.vmware.com/s/article/1021976>)


Setting the Maximum I/O


Follow this guideline to set the maximum I/O request size for storage devices.

`Disk.DiskMaxIOSize` determines the maximum I/O request size passed to storage devices. With PowerStore, it is required to change this parameter from 32767 (the default setting of 32MB) to 1024 (1MB).

Example: Setting `Disk.DiskMaxIOSize` to 1024 (1MB)

```
esxcli system settings advanced set -o "/Disk/DiskMaxIOSize" --int-value 1024
```


 **NOTE:** When setting `Disk.DiskMaxIOSize` to 1MB on ESXi hosts connected to arrays other than PowerStore, performance on large I/Os may be impacted.


 **NOTE:** For details on a possible Windows EFI boot issue with PowerStore storage (in case the above maximum I/O block size setting adjustment cannot be done), refer to VMware KB article 2137402 on the VMware website (<https://kb.vmware.com/s/article/2137402>). This KB article also lists all ESXi versions that include a fix for this issue. With these versions the `Disk.DiskMaxIOSize` setting adjustment to 1024 (1MB) is not required. In such cases, to avoid a potential host-side DU issue when upgrading the ESXi version on the host, make sure to refer to this KB and confirm that the target ESXi version also includes a fix for this Windows EFI boot issue.

Confirming UNMAP Priority

This topic provides steps for setting UNMAP priority on a DataStore.

Prerequisites

 **NOTE:** Provisioning Virtual Disks with UNMAP set to a non-default priority on a DataStore provisioned on PowerStore may result in an increased amount of write I/Os to the storage subsystem. It is therefore highly recommended to keep UNMAP set to a default priority.

 **NOTE:** Refer to Dell EMC KB article SLN320477 (<https://www.dell.com/support/article/sln320477>) for further unmap-related recommendations when doing VMFS (Virtual Machine File System) bootstorm or failover with VMware SRM (Site Recovery Manager) on VMFS datastores from ESXi hosts connected to PowerStore.

To set UNMAP priority on a datastore:

Steps

1. Access vSphere Client and locate the **ESX Storage Configuration** panel.
2. Go to the **Storage** tab.
3. Right-click the affected DataStore.
4. Click **Edit Space Reclamation**.
5. Set the value (with ESX 6.7 or later, set to 100MB/s, with ESX 6.5, set to Low).
6. Click **OK**.

Configuring VMware vSphere with PowerStore Storage in a Multiple Cluster Configuration

Use the listed recommended values when multiple clusters are connected to vSphere.

The following table lists the recommended vSphere settings when multiple storage clusters are connected to vSphere (in addition to PowerStore). Follow these recommendations instead of other recommendations included in this chapter.

For reference, this table also includes the corresponding recommendations for settings when vSphere is connected to PowerStore storage only.

Setting	Scope/Granularity	Multi-Storage Setting	PowerStore Only Setting
UCS FC Adapter Policy	Per vHBA	default	default
Cisco nfnic lun_queue_depth_per_path	Global	default (32)	default (32)
Disk.SchedNumReqOutstanding	LUN	default	default
Disk.SchedQuantum	Global	default	default
Disk.DiskMaxIOSize	Global	1MB	1MB
XCOPY (/DataMover/MaxHWTransferSize)	Global	default (4MB)	default (4MB)
XCOPY (Claim Rule)	N/A	No Guidance	No Guidance
vCenter Concurrent Clones (config.vpxd.ResourceManager.maxCostPerHos)	vCenter	default (8)	default (8)

- UCS FC Adapter Policy - The total number of I/O requests that can be outstanding on a per-virtual Host Bus Adapter (vHBA) in UCS.
- Cisco nfnic lun_queue_depth_per_path - Cisco nfnic driver setting to set the LUN queue depth per path. The default value for this setting is 32 (recommended). For details on Cisco nfnic settings, refer to the Cisco nfnic driver documentation on the Cisco website.
- DiskSchedNumReqOutstanding - The total number of outstanding commands permitted from all virtual machines collectively on the host to a LUN. For details, refer to VMware vSphere documentation.
- Disk.SchedQuantum - The maximum number of consecutive "sequential" I/Os allowed from one VM before forcing a switch to another VM. For details, refer to VMware vSphere documentation.
- Disk.DiskMaxIOSize - The maximum I/O size ESX allows before splitting I/O requests. For details, refer to [Setting the Maximum I/O](#)
- XCOPY (/DataMover/MaxHWTransferSize) - The maximum number of blocks used for XCOPY operations. For details, refer to VMware vSphere documentation.
- vCenter Concurrent Clones (config.vpxd.ResourceManager.maxCostPerHost) - The maximum number of concurrent full clone operations allowed (the default value is 8). For details, refer to VMware vSphere documentation.

Multipathing Software Configuration

Configuring Native Multipathing

PowerStore supports the Native Multipathing Plugin (NMP).

This section describes the procedure required for configuring native multipathing for PowerStore volumes.

For best performance, it is recommended to do as follows:




- Set the NMP Round Robin path selection policy on PowerStore volumes presented to the ESXi host.
- Set the NMP Round Robin path switching frequency to PowerStore volumes from the default value (1000 I/O packets) to 1.

These settings ensure optimal distribution and availability of load between I/O paths to the PowerStore storage.

Configuring NMP Round Robin as the Default Pathing Policy for All PowerStore Volumes

Follow this procedure to configure NMP Round Robin as the default pathing policy for all PowerStore volumes using the ESXi command line.


About this task

-  **NOTE:** As of VMware ESXi version 6.7, Patch Release ESXi670-201912001, the satp rule presented in this procedure is already integrated into the ESXi kernel.
-  **NOTE:** Use this method when no PowerStore volume is presented to the host. PowerStore volumes already presented to the host are not affected by this procedure (unless they are unmapped from the host).
-  **NOTE:** With ESXi 6.7 hosts connected to PowerStore, it is recommended to disable `action_OnRetryErrors`. For details on this ESXi parameter, refer to VMware KB# 67006 (<https://kb.vmware.com/s/article/67006>).

Steps


1. Open an SSH session to the host as `root`.
2. Run the following command to configure the default pathing policy for newly defined PowerStore volumes to Round Robin with path switching after each I/O packet:

```
esxcli storage nmp satp rule add -c tpgs_on -e "PowerStore" -M PowerStore -P VMW_PSP_RR -O iops=1 -s VMW_SATP_ALUA -t vendor -V DelleMC -o disable_action_OnRetryErrors
```

 **NOTE:** Use the `disable_action_OnRetryErrors` parameter only with ESXi 6.7 hosts.

This command also sets the NMP Round Robin path switching frequency for newly defined PowerStore volumes to switch every I/O.


Next steps

-  **NOTE:** Using this method does not impact any non-PowerStore volume presented to the ESXi host.

Configuring NMP Round Robin on a PowerStore Volume Already Presented to the ESXi Host

Follow this procedure to configure NMP Round Robin on a PowerStore volume that is already presented to the ESXi host, using ESXi command line:

About this task

-  **NOTE:** Use this method only for PowerStore volumes already presented to the host. For volumes not yet presented to the host, refer to [Configuring vSphere NMP Round Robin as the default pathing policy for all volumes](#).

Steps

1. Open an SSH session to the host as `root`.
2. Run the following command to obtain the NAA of PowerStore LUNs presented to the ESXi host:

```
#esxcli storage nmp path list | grep DelleMC -B1
```

The following example demonstrates issuing the `esxcli storage nmp path list` command to obtain the NAA of all PowerStore LUNs presented to the ESXi host:

```
#esxcli storage nmp path list | grep DelleMC -B1
Device: naa.68ccf09800e8fa24ea37a1bc49d9f6b8
Device Display Name: DellEMC Fibre Channel Disk
(naa.68ccf09800e8fa24ea37a1bc49d9f6b8)
--
Device: naa.68ccf098003a54f16d2eddc3217da922
Device Display Name: DellEMC Fibre Channel Disk
```

```
(naa.68ccf098003a54f16d2eddc3217da922)
--
Device: naa.68ccf09000000000c9f6dlacdale4567
Device Display Name: DellEMC Fibre Channel RAID Ctlr
(naa.68ccf09000000000c9f6dlacdale4567)
--
Device: naa.68ccf098009c1cf3bfe0748a9183681a
Device Display Name: DellEMC Fibre Channel Disk
(naa.68ccf098009c1cf3bfe0748a9183681a)
--
Device: naa.68ccf09000000000c9f6dlacdale4567
Device Display Name: DellEMC Fibre Channel RAID Ctlr
(naa.68ccf09000000000c9f6dlacdale4567)
--
Device: naa.68ccf098009c1cf3bfe0748a9183681a
Device Display Name: DellEMC Fibre Channel Disk
(naa.68ccf098009c1cf3bfe0748a9183681a)
--
Device: naa.68ccf09800e8fa24ea37a1bc49d9f6b8
Device Display Name: DellEMC Fibre Channel Disk
(naa.68ccf09800e8fa24ea37a1bc49d9f6b8)
--
Device: naa.68ccf098003a54f16d2eddc3217da922
Device Display Name: DellEMC Fibre Channel Disk
(naa.68ccf098003a54f16d2eddc3217da922)
```

3. Run the following command to modify the path selection policy on the PowerStore volume to Round Robin:

```
#esxcli storage nmp device set --device="<naa_id>" --psp=VMW_PSP_RR
```

For example:

```
#esxcli storage nmp device set --device="naa.68ccf098003f1461569ea4750e9dac50" --
psp=VMW_PSP_RR
```

4. Run the following command to set the NMP Round Robin path switching frequency on PowerStore volumes from the default value (1000 I/O packets) to 1:

```
#esxcli storage nmp psp roundrobin deviceconfig set --device="<naa_id>" --iops=1 --
type=iops
```

For example:

```
#esxcli storage nmp psp roundrobin deviceconfig set --
device="naa.68ccf098003f1461569ea4750e9dac50" --iops=1 --type=iops
```

5. Run the following command to validate that changes were applied to all PowerStore LUNs:

```
#esxcli storage nmp device list | grep -B1 -A4 DellEMC
```

Each listed PowerStore LUN should have the following NMP settings:


- Path Selection Policy: VMW_PSP_RR
- Path Selection Policy Device Config: policy=rr, iops=1

The following example demonstrates issuing the `esxcli storage nmp device list` command to validate that changes were applied to all PowerStore LUNs:

```
#esxcli storage nmp device list |grep -B1 -A4 DellEMC
naa.68ccf09000000000c9f6dlacdale4567
Device Display Name: DellEMC Fibre Channel RAID Ctlr
(naa.68ccf09000000000c9f6dlacdale4567)
Storage Array Type: VMW_SATP_DEFAULT_AA
Storage Array Type Device Config: {action_OnRetryErrors=off}
Path Selection Policy: VMW_PSP_FIXED
Path Selection Policy Device Config:
{preferred=vmhba2:C0:T0:L0;current=vmhba2:C0:T0:L0}
--
naa.68ccf098003a54f16d2eddc3217da922
Device Display Name: DellEMC Fibre Channel Disk
(naa.68ccf098003a54f16d2eddc3217da922)
Storage Array Type: VMW_SATP_ALUA
Storage Array Type Device Config:
{implicit_support=on; explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off; {TPG_id=1,TPG_state=ANO}}
Path Selection Policy: VMW_PSP_RR
Path Selection Policy Device Config: {policy=rr,iops=1,bytes=10485760,useANO=0;
```

```
lastPathIndex=0: NumIOsPending=0,numBytesPending=0}
--
naa.68ccf09800e8fa24ea37a1bc49d9f6b8
Device Display Name: DellEMC Fibre Channel Disk
(naa.68ccf09800e8fa24ea37a1bc49d9f6b8)
Storage Array Type: VMW_SATP_ALUA
Storage Array Type Device Config:
{implicit_support=on; explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off; {TPG_id=1,TPG_state=AO}}
Path Selection Policy: VMW_PSP_RR
```

Next steps

 **NOTE:** Using this method does not impact any non-PowerStore volumes presented to the ESXi host.

For details, refer to VMware KB articles 1017760 and 2069356 on the VMware website (<https://kb.vmware.com/kb/1017760/> and <https://kb.vmware.com/kb/2069356/> respectively).

Post-Configuration Steps


When host configuration is completed, you can use the PowerStore storage from the host.

You can create, present, and manage volumes accessed from the host via PowerStore Manager, CLI, or REST API. Refer to the *PowerStore Manager Online Help*, *CLI Reference Guide*, or *REST API Reference Guide* for additional information.

The Dell EMC Virtual Storage Integrator (VSI) version 8.4 and later plug-in can be used in to provision from within Virtual Machine File System (VMFS) datastores and Raw Device Mapping volumes on PowerStore. Furthermore, the Dell EMC VSI Storage Viewer version 8.4 and later plug-in extends the vSphere Client to facilitate the discovery and identification of PowerStore storage devices allocated to VMware ESXi hosts and virtual machines.

For further information on using these two vSphere Client plug-ins, refer to the *VSI Unified Storage Management Product Guide* and the *VSI Storage Viewer Product Guide*.

AppsOn: Virtual Machine Compute and Storage Collocation Rules for PowerStore X Clusters

 **NOTE:** The following is applicable only to PowerStore X multi-appliance clusters with operating system version 2.0 (or later).


To ensure compute and storage resources collection for optimal VM performance, you can use the predefined VM and host rules in vCenter server.

To tie a user VM to a host group, add that VM to the predefined VM group in vCenter server.

For more details, see *PowerStore Virtualization Infrastructure Guide*.

Disk Formatting

When creating volumes in PowerStore for a vSphere ESXi host, the following considerations should be made:

- Disk logical block size - The only logical block (LB) size supported by vSphere ESXi for presenting volumes is 512 bytes.
 **NOTE:** For details on formatting a newly created volume, refer to the *PowerStore Configuring Volumes* guide that matches the version running on your PowerStore cluster.
- Disk alignment - Unaligned disk partitions may substantially impact I/O to the disk.

With vSphere ESXi, datastores and virtual disks are aligned by default as they are created. Therefore, no further action is required to align these in ESXi.

With virtual machine disk partitions within the virtual disk, alignment is determined by the guest OS. For virtual machines that are not aligned, consider using tools such as UBERAlign to realign the disk partitions as required.

Presenting PowerStore Volumes to the ESXi Host

Specify ESXi as the operating system when presenting PowerStore volumes to the ESXi host.

NOTE: Using data reduction and /or encryption software on the host side will affect the PowerStore cluster data reduction.

NOTE: When using iSCSI software initiator with ESXi and PowerStore storage, it is recommended to use only lower case characters in the IQN to correctly present the PowerStore volumes to ESXi. For more details, refer to VMware KB article 2017582 on the VMware website (https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2017582).

When adding host groups and hosts to allow ESXi hosts to access PowerStore volumes, specify **ESXi** as the operating system for newly-created hosts.

Creating a File System

It is recommended to create the file system using its default block size.

NOTE: File system configuration and management is out of the scope of this document.

It is recommended to create the file system using its default block size (using a non-default block size may lead to unexpected behavior). Refer to your operating system and file system documentation.

VMware Paravirtual SCSI Controllers

Configure virtual machines with paravirtual SCSI controllers to achieve higher throughput and lower CPU usage.

For optimal resource utilization of virtual machines with PowerStore, it is recommended to configure virtual machines with paravirtualized SCSI controllers. VMware paravirtual SCSI controllers are high performance storage controllers that can provide higher throughput and lower CPU usage. These controllers are best suited for high performance storage environments.

For further details on configuring virtual machines with paravirtualized SCSI controllers, refer to the *vSphere Virtual Machine Administration Guide* in the VMware vSphere documentation.

Virtual Disk Provisioning

Follow these recommendations for provisioning virtual disks on the PowerStore cluster.

For optimal space utilization with vSphere ESXi 6.x and above, it is recommended to provision virtual disks on the PowerStore cluster, using Thin Provisioning.

In Thin Provisioning format, in-guest space reclamation is available, provided the following minimum requirements are fulfilled:

- ESXi 6.x
- Thin virtual disks
- VM hardware version 11
- EnableBlockDelete set to 1
- Guest OS support of UNMAP

NOTE: Refer to the corresponding guest OS chapter within this document for instructions on how to efficiently create a file system.

NOTE: SCSI-3 Persistent Reservations (SCSI3-PRs) on a virtual disk (VMDK) level is currently not supported with PowerStore storage.

Virtual Machine Guest OS Settings

This section details the recommended settings and considerations for virtual machines guest OS.

- LUN Queue Depth - For optimal virtual machine operation, configure the virtual machine guest OS to use the default queue depth of the virtual SCSI controller. For details on adjusting the guest OS LUN queue depth, refer to VMware KB# 2053145 on the VMware website (<https://kb.vmware.com/kb/2053145>).
- RDM volumes in guest OS - Span RDM volumes, used by the virtual machine, across SCSI controllers to prevent a bottleneck on a single SCSI controller.
- RDM volumes in guest OS used for Microsoft Cluster (MSCS) - ESXi hosts with visibility to RDM volumes, used by Microsoft Cluster (MSCS) may take a long time to start or to perform LUN rescan.

For the required settings on the RDM volumes, refer to VMware KB# 1016106 (https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1016106).

Host Configuration for Microsoft Windows


This chapter contains the following topics:

Topics:

- [Recommended Configuration Values Summary](#)
- [Fibre Channel Configuration](#)
- [iSCSI Configuration](#)
- [Multipathing Software Configuration](#)
- [Post-Configuration Steps - Using the PowerStore system](#)

Recommended Configuration Values Summary

The following table summarizes all used variables and their values when configuring hosts for Microsoft Windows.

 **NOTE:** Unless indicated otherwise, use the default parameters values.

Validation	Impact	Severity	Refer to Section
To clarify the above note for using default parameter settings unless stated otherwise in this chapter, make sure that the following are set per the default OS setting: <ul style="list-style-type: none"> • LUN and HBA queue depth • HBA timeout 	Stability & Performance	Recommended	For further details, refer to OS and HBA documentation.
Specify Windows as the operating system for each defined host.	Serviceability	Mandatory	Presenting PowerStore Volumes to the Windows Host
Load balancing: Use Round Robin or Least Queue Depth for Microsoft Native Microsoft Multipath I/O (MPIO) with Windows Server 2012/R2 and above.	Performance	Warning	Configuring Native Multipathing using Microsoft Multipath I/O (MPIO)
Temporarily disable UNMAP during file systems creation: <ul style="list-style-type: none"> • To temporarily disable UNMAP on the host (prior to file system creation): <pre>fsutil behavior set DisableDeleteNotify 1</pre> • To re-enable UNMAP on the host (after file system creation): 	Performance	Recommended	Creating a File System

Validation	Impact	Severity	Refer to Section
fsutil behavior set DisableDeleteNotify 0			

Fibre Channel Configuration

This section describes the recommended configuration that should be applied when attaching hosts to PowerStore cluster using Fibre Channel.

NOTE: This section applies only to FC. If you are using only iSCSI with Windows, go to [iSCSI HBA Configuration](#).

NOTE: Before you proceed, review [Fibre Channel SAN Guidelines](#).

Pre-Requisites

This section describes the pre-requisites for FC HBA configuration.

- Refer to the E-Lab Interoperability Navigator (<https://elabnavigator.emc.com>) for supported FC HBA models and drivers.
- Verify all HBAs are at the supported driver, firmware and BIOS versions.
- Verify all HBAs BIOS settings are configured according to E-Lab recommendations. Follow the procedures in one of the following documents according to the FC HBA type:
 - For Qlogic HBAs, refer to *Dell EMC Host Connectivity with Qlogic Fibre Channel and iSCSI HBAs and Converged Network Adapters (CNAs) for the Windows Environment*.
 - For Emulex HBAs, refer to *Dell EMC Host Connectivity with Emulex Fibre Channel and iSCSI HBAs and Converged Network Adapters (CNAs) for the Windows Environment*.
 - For Cisco UCS fNIC HBAs, refer to the *Cisco UCS Virtual Interface Card Drivers for Windows Installation Guide* for complete driver installation instructions .

iSCSI Configuration

This section describes the recommended configuration that should be applied when attaching hosts to PowerStorecluster using iSCSI.

NOTE: This section applies only to iSCSI. If you are using only Fibre Channel with Windows and PowerStore, go to [Fibre Channel HBA Configuration](#). Make sure you review [iSCSI SAN Guidelines](#) before you proceed.

Pre-Requisites

This section describes the pre-requisites when attaching hosts to PowerStore cluster using iSCSI.

Follow the Microsoft recommendations for installation and setup of the appropriate NIC/iSCSI HBA for your system. It is recommended to install the latest NIC/iSCSI HBA driver version (patch), as described in the support documentation for each specific NIC/iSCSI HBA.

Refer to the E-Lab Interoperability Navigator (<https://elabnavigator.emc.com>) for supported NIC/iSCSI HBA models and drivers.

Multipathing Software Configuration

This topic introduces Multipathing Software Configuration for Microsoft Windows

PowerStore supports native multipathing using Microsoft Native Microsoft Multipath I/O (MPIO) with Windows Server 2012/R2 and above, or multipathing using PowerPath.

Configuring Native Multipathing Using Microsoft Multipath I/O (MPIO)

This topic describes configuring native multipathing using Microsoft Multipath I/O (MPIO).

For optimal operation with PowerStore storage, configure the Round-Robin (RR) policy or the Least Queue Depth policy for MPIO for devices presented from PowerStore. Using these policies, I/O operations are balanced across all available paths.

To configure the native multipathing, using Microsoft Multipath I/O (MPIO), you should perform the following task:

[Enabling MPIO on the Windows Host and Configuring MPIO for PowerStore Volumes Presented to the Host](#)

Enabling MPIO on the Windows Host and Configuring MPIO for PowerStore Volumes Presented to the Host

This topic describes enabling and configuring MPIO on the Windows host.

About this task

Before configuring the native multipathing, you should enable MPIO on the server by adding the MPIO feature to Windows.

Steps

1. Open PowerShell on the host.
2. Run the following commands to install MPIO if it is not already installed:

```
Get-WindowsOptionalFeature -Online -FeatureName MultiPathIO
Enable-WindowsOptionalFeature -Online -FeatureName MultiPathIO
```
3. Run the following command to set vid/pid:

```
New-MSDSMSupportedHW -VendorId DellEMC -ProductId PowerStore
```
4. Run one of the following commands to set RoundRobin failover policy or Least Queue Depth failover policy, respectively:
 - Round-Robin

```
Set-MSDSMGlobalDefaultLoadBalancePolicy -Policy RR
```
 - Least Queue Depth

```
Set-MSDSMGlobalDefaultLoadBalancePolicy -Policy LQD
```
5. Run the following commands to set MPIO timeout values:

```
Set-MPIOSetting -NewPathVerificationState Enabled
Set-MPIOSetting -NewPathVerificationPeriod 30
Set-MPIOSetting -NewPDORemovePeriod 20
Set-MPIOSetting -NewRetryCount 3
Set-MPIOSetting -NewRetryInterval 3
Set-MPIOSetting -custompathrecovery enabled
Set-MPIOSetting -newpathrecoveryinterval 10
Set-MPIOSetting -NewDiskTimeout 30
```
6. To verify MPIO settings on the host, run the following command:

```
Get-MPIOSetting
```

Post-Configuration Steps - Using the PowerStore system

This topic describes the post-configuration steps using the PowerStore system.

After the host configuration is completed, you can use the PowerStore storage from the host.

You can create, present, and manage volumes accessed from the host via PowerStore Manager, CLI, or REST API. Refer to the *PowerStore Manager Online Help*, *CLI Reference Guide*, or *REST API Reference Guide* for additional information.

Presenting PowerStore Volumes to the Windows Host

This topic discusses presenting PowerStore Volumes to the Windows host.

When adding host groups and hosts to allow Windows hosts to access PowerStore volumes, specify **Windows** as the operating system for the newly-created hosts.

NOTE: Setting the host's operating system is required for optimal interoperability and stability of the host with PowerStore storage. You can adjust the setting while the host is online and connected to the PowerStore cluster with no I/O impact.

NOTE: Refer to the *PowerStore Configuring Volumes* Guide for additional information.

Creating a File System

This topic discusses creating a file system.

File system configuration and management are out of the scope of this document.

NOTE: Some file systems may require you to properly-align the file system on the PowerStore volume. It is recommended to use specified tools to optimally match your host with application requirements.

NOTE: Creating a file system formatting with UNMAP enabled on a host connected to PowerStore may result in an increased amount of write I/Os to the storage subsystem. It is highly recommended to disable UNMAP during file system creation.

To disable UNMAP during file system creation:

1. Open a Windows CMD window on the host.
2. Run the following `fsutil` command to temporarily disable UNMAP on the host (prior to file system creation):

```
fsutil behavior set DisableDeleteNotify 1
```

3. Once file system creation is complete, re-enable UNMAP by running the following command:

```
fsutil behavior set DisableDeleteNotify 0
```

NOTE: To verify the current setting of the file system, run the following `fsutil` command:

```
fsutil behavior query DisableDeleteNotify
```

- `DisableDeleteNotify=0` - Indicates that the 'Trim and Unmap' feature is on (enabled).
- `DisableDeleteNotify=1` - Indicates that the 'Trim and Unmap' feature is off (disabled).

Host Configuration for Linux


This chapter contains the following topics:

Topics:

- [Recommended Configuration Values Summary](#)
- [Fibre Channel Configuration](#)
- [iSCSI Configuration](#)
- [Multipathing Software Configuration](#)
- [Post-Configuration Steps - Using the PowerStore system](#)

Recommended Configuration Values Summary

The following table summarizes all used and recommended variables and their values when configuring hosts for Linux.


 **NOTE:** Unless indicated otherwise, use the default parameters values.


Validation	Impact	Severity	Refer to Section
To clarify the above note for using default parameter settings unless stated otherwise in this chapter, make sure that the following are set per the default OS setting: <ul style="list-style-type: none"> • LUN and HBA queue depth • HBA timeout 	Stability & Performance	Recommended	For further details, refer to OS and HBA documentation.
Specify Linux as the operating system for each defined host.	Serviceability	Mandatory	Presenting PowerStore Volumes to the Linux Host
Native Multipathing: To configure the PowerStore disk device, modify <code>/etc/multipath.conf</code> file: <ul style="list-style-type: none"> • <code>vendor</code> - "DellEMC" • <code>product</code> - "PowerStore" • <code>path_selector</code> - "queue-length 0" • <code>path_grouping_policy</code> - "group_by_prio" • <code>path_checker</code> - "tur" • <code>detect_prio</code> - "yes" • <code>failback</code> - "immediate" • <code>no_path_retry</code> - "3" • <code>rr_min_io_rq</code> - "1" • <code>fast_io_fail_tmo</code> - "15" • <code>max_sectors_kb</code> - "1024" 	Performance	Recommended	Configuring Linux Native Multipathing
iSCSI Configuration: To configure the PowerStore disk device, modify <code>/etc/iscsi/iscsid.conf</code> file:	Performance	Recommended	Configuring PowerStore Cluster Disk Device with iSCSI

Validation	Impact	Severity	Refer to Section
<ul style="list-style-type: none"> Set <code>node.session.timeo.replacement_timeout = 15</code> 			
Temporarily disable UNMAP during file system creation. <ul style="list-style-type: none"> When creating a file system using the <code>mke2fs</code> command - Use the <code>"-E nodiscard"</code> parameter When creating a file system using the <code>mkfs.xfs</code> command - Use the <code>"-K"</code> parameter 	Performance	Recommended	Creating a File System

Fibre Channel Configuration

This section describes the recommended configuration that should be applied when attaching hosts to PowerStore cluster using Fibre Channel.

 **NOTE:** This section applies only to FC. If you are using only iSCSI with Linux, go to [iSCSI HBA Configuration](#).

 **NOTE:** Before proceeding, review [Fibre Channel SAN Guidelines](#).

Pre-Requisites

When attaching a host to PowerStore cluster using Fibre Channel, the following pre-requisites need to be met:

- Refer to the E-Lab Interoperability Navigator (<https://elabnavigator.emc.com>) for supported Fibre Channel HBA models and drivers.
- Verify all HBAs are at the supported driver, firmware and BIOS versions.
- Verify all HBAs BIOS settings are configured according to E-Lab recommendations.

NVMe-FC Configuration on Linux Hosts

For details on NVMe over Fibre Channel (NVMe -FC) configuration with Linux hosts, see the Linux host documentation.

For NVMe-FC configuration on Red Hat Enterprise Linux, see Red Hat KB 4706181 (<https://access.redhat.com/solutions/4706181>).

Setting Up NVMe HBAs

For further details on HBA setup, see NVMe HBA documentation.

Setting Up Qlogic NVMe HBA

Follow these steps to setup a Qlogic NVMe HBA.

Steps

- Access the Linux host as `root`.
- Edit the `/etc/modprobe.d/qla2xxx.conf` configuration file with the following data:


```
options qla2xxx ql2xextended_error_logging=1 ql2xfwloadbin=2 ql2xnvmeenable=1
```

Setting Up Emulex NVMe HBA

Follow these steps to setup an Emulex NVMe HBA.

Steps

1. Access the Linux host as `root`.
2. Edit the `/etc/modprobe.d/lpfc.conf` configuration file with the following data:

```
options lpfc lpfc_lun_queue_depth=128 lpfc_sg_seg_cnt=256 lpfc_max_luns=65535  
lpfc_enable_fc4_type=3
```

iSCSI Configuration

This section provides an introduction to the recommended configuration to be applied when attaching hosts to PowerStore cluster using iSCSI.

NOTE: This section applies only to iSCSI. If you are using only Fibre Channel with Linux and PowerStore, go to [Fibre Channel Configuration](#).

NOTE: Before proceeding, review the [iSCSI SAN Guidelines](#) section.

Pre-Requisites

Refer to the E-Lab Interoperability Navigator (<https://elabnavigator.emc.com>) for supported NIC/iSCSI HBA models and drivers.

Configuring the PowerStore Cluster Disk Device with iSCSI

To configure the PowerStore disk device, modify the configuration file with the following parameters.

The configuration file for the multipath daemon is `multipath.conf`. It is used to overwrite the built-in configuration table of the multipath daemon. When iSCSI is used with PowerStore, the `iscsi.conf` file is used to overwrite iSCSI specific multipathing related settings.

To configure the PowerStore disk device, modify the `/etc/iscsi/iscsid.conf` file, using the following parameters:

NOTE: The example below is based on RedHat. For details on configuring PowerStore disk device with iSCSI, refer to the specific instructions of your operating system.

Parameter	Description	Value
<code>node.session.timeo.replacement_timeout</code>	Specifies the number of seconds the iSCSI layer waits for a timed-out path/session to re-establish before failing any commands on that path/session. The default value is 120.	15

Using these settings prevents commands from being split by the iSCSI initiator and enables instantaneous mapping from the host to the volume.

To apply the adjusted `iscsid.conf` settings, run the following command on the Linux host:

```
service iscsi restart
```

NOTE: If a previous iSCSI target is discovered on the Linux host, delete the iSCSI DB and rerun the iSCSI target discovery procedure with the `iscsid.conf` setting described above.

Configuring the PowerStore Cluster Disk Device with iSCSI - Single Network Subnet Support

NOTE: This subsection is applicable only to a cluster running PowerStore OS version 1, or in case you are using only a single network subnet for the iSCSI portals.

In PowerStore OS version 1, only a single network subnet is supported for the iSCSI target portals.

By design, on various Linux distributions, only two network interfaces can be configured on the same network subnet. For details, see RedHat KB 30564 (<https://access.redhat.com/solutions/30564>) and RedHat KB 53031 (<https://access.redhat.com/solutions/53031>).

In light of this limitation, use one of the following solutions to make both network interfaces accessible with hosts that are connected to PowerStore storage:

- [Policy-Based Routing](#) (recommended)
- [Bonding/Teaming](#)
- [Disable Reverse Path Filtering](#)

Policy-Based Routing

This topic outlines policy-based routing as a solution to the single network subnet limitation (recommended solution).

This solution is based on adding routing tables and rules, binding source IP address for each route, and adding those as default gateways for each network interface.

Using this solution, a routing table is defined for each interface, thus the default routing table is redundant for those interfaces.

For additional technical information on Policy-Based Routing, refer to <https://access.redhat.com/solutions/30564>.

Bonding/Teaming

This topic outlines bonding/teaming as a solution to the single network subnet limitation.

NOTE: The following section does not apply to hosts directly attached to the PowerStore appliances.

This solution is based on the Bond and Network teaming configuration.

- Bond - Binding multiple network interfaces into a single bonded channel enables them to act as one virtual interface. That way, only a single network address is defined and the said limitation does not apply.

For additional technical information on configuring network bond on Red Hat Enterprise Linux 7, refer to https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/networking_guide/ch-configure_network_bonding.

- Network Teaming - With Red Hat Enterprise Linux 7, Network Teaming is offered as a new implementation of the bonding concept. The existing bonding driver is unaffected.

Network Teaming is offered as an alternative and does not replace bonding in Red Hat Enterprise Linux 7. For additional technical information on configuring Network Teaming, refer to https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/networking_guide/ch-configure_network_teaming.

For a comparison between Bonding and Network Teaming implementations, refer to https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/networking_guide/sec-comparison_of_network_teaming_to_bonding.

Disabling Reverse Path Filtering

This topic outlines disabling reverse path filtering as a solution to the single network subnet limitation.

RHEL 6 and above are configured by default to apply Strict Reverse Path Forwarding filtering recommended in RFC 3704 - Ingress Filtering for Multihomed Networks.

NOTE: Before making this change, refer to the Root Cause section of this article to understand what it does and review alternative solutions as explained in <https://access.redhat.com/solutions/53031>.

Setting the Reverse Path Filtering to 2 (loose) on the relevant network interfaces makes them both accessible and routable.

To apply this change, add the following lines to `/etc/sysctl.conf`:

```
net.ipv4.conf.eth2.rp_filter = 2
```

```
net.ipv4.conf.eth3.rp_filter = 2
```

NOTE: In this example, `eth2` and `eth3` are the network interfaces used for iSCSI. Make sure to modify to the relevant interfaces.

To reload the configuration:

```
sysctl -p
```

To view the current Reverse Path Filtering configuration on the system:

```
sysctl -ar "\.rp_filter"
```

```
net.ipv4.conf.all.rp_filter = 1
net.ipv4.conf.default.rp_filter = 1
net.ipv4.conf.eth0.rp_filter = 1
net.ipv4.conf.eth1.rp_filter = 1
net.ipv4.conf.eth2.rp_filter = 2
net.ipv4.conf.eth3.rp_filter = 2
net.ipv4.conf.lo.rp_filter = 0
```

Multipathing Software Configuration

Configuring Linux Native Multipathing


For a PowerStore cluster to function properly with Linux hosts, configure the multipath settings file `/etc/multipath.conf`:

NOTE: If the host is connected to a cluster other than PowerStore, the configuration file may include additional devices.

NOTE: If the `multipath.conf` file includes a blacklist section, it should come before the devices section. For additional information, refer to the *Importing External Storage to PowerStore Guide* document.

```
devices {
device
{
vendor DelliEMC
product PowerStore
path_selector "queue-length 0"
path_grouping_policy group_by_prio
path_checker tur
detect_prio yes
failback immediate
no_path_retry 3
rr_min_io_rq 1
fast_io_fail_tmo 15
max_sectors_kb 1024 //only for RHEL 6.9 (or later 6.x versions) and RHEL 7.4 (or later
7.x versions)
} //other devices}
}
```

Parameter	Description	Value
vendor	Specifies the vendor name	DelliEMC
product	The below configuration applies only to PowerStore volumes.	PowerStore
path_selector	Sends the next group of I/Os to the path with the least number of outstanding I/O requests	queue-length 0

Parameter	Description	Value
path_grouping_policy	Specifies the default path grouping policy to apply to PowerStore volumes <ul style="list-style-type: none"> • Paths are grouped by priorities assigned by the cluster. • A higher priority (50) is set as Active/Optimized. Lower priority (10) is set as Active/Non-Optimized. 	group_by_prio
path_checker	Specifies TEST UNIT READY as the default method used to determine the state of the paths.	tur
detect_prio	If set to yes, multipath will try to detect whether the device supports ALUA. If so, the device will automatically use the alua prioritizer. Otherwise, the prioritizer will be selected as usual. Default value is no.	yes
failback	Manages the path group failback. Immediate refers to immediate failback to the highest priority path group that contains active paths.	Immediate
no_path_retry	Specifies the number of times the system should attempt to use a failed path before disabling queuing.	3
rr_min_io_rq	Specifies the number of I/O requests to route to a path before switching to the next path in the same path group, using request-based device-mapper-multipath. This setting should be used on systems running current kernels. On systems running kernels older than 2.6.31, use rr_min_io. Default value is 1.	1
fast_io_fail_tmo	Specifies the number of seconds the scsi layer will wait after a problem has been detected on a FC remote port before failing I/O to devices on that remote port. This value should be smaller than dev_loss_tmo. Setting this parameter to off will disable the timeout.	15
max_sectors_kb	Applies to Red Hat Enterprise Linux Release 6.9 (or later 6.x versions) and Red Hat Enterprise Linux Release 7.4 (or later 7.x versions). Sets the max_sectors_kb device queue parameter to the specified value on all underlying paths of a multipath device before the multipath device is first activated. When a multipath device is created, the device inherits the max_sectors_kb value from the path devices. Manually raising this value for the multipath device or lowering it for the path devices can cause multipath to create I/O operations larger than the path devices allow. Using the max_sectors_kb parameter is an easy way to set these values before a multipath device is created on top of the path devices and prevent invalid-sized I/O operations from being passed. If this parameter is not set by the user, the path devices have it set by their device driver, and the multipath device inherits it from the path devices.  NOTE: In PowerStore cluster the maximum I/O size is 1MB. PowerStore does not set an optimal transfer size.	1024

Configuring Device Mapper Multipathing (DM-multipathing) with NVMe-FC

When configuring NVMe-FC on a Linux host that is connected to PowerStore, also configure DM-multipathing to setup multiple I/O paths between the Linux host and the PowerStore array into a single device over NVMe-FC.

When configuring DM-multipathing for PowerStore NVMe-FC devices on the Linux host, configure the multipath settings file `/etc/multipath.conf`:

```
devices {
device
{
    vendor                .*
    product               dell EMC PowerStore
    uid_attribute         ID_WWN
    prio                  ana
    failback              immediate
    path_grouping_policy  "group_by_prio"
    # path_checker        directio
    path_selector          "queue-length 0"
    detect_prio           "yes"
    fast_io_fail_tmo      15
    no_path_retry         3
    rr_min_io_rq          1
} //other devices}
}
```

Configuring Oracle ASM

For proper functioning with Linux hosts using Oracle ASM with PowerStore, configure the Oracle ASM settings file.

Prerequisites

For a PowerStore cluster to function properly with Linux hosts that are using the Oracle ASM volume management software with ASMLib driver, follow these steps to configure the `/etc/sysconfig/oracleasm` settings file:

Steps

1. Modify the following lines in the `/etc/sysconfig/oracleasm` file according to the multipathing used on the host:

```
# ORACLEASM_SCANORDER: Matching patterns to order disk scanning
```

```
ORACLEASM_SCANORDER=""
```

```
# ORACLEASM_SCANEXCLUDE: Matching patterns to exclude disks from scan
```

```
ORACLEASM_SCANEXCLUDE=""
```

- When DM-MPIO multipathing is used on the Linux host, edit these lines as follows:

```
# ORACLEASM_SCANORDER: Matching patterns to order disk scanning
```

```
ORACLEASM_SCANORDER="dm"
```

```
# ORACLEASM_SCANEXCLUDE: Matching patterns to exclude disks from scan
```

```
ORACLEASM_SCANEXCLUDE="sd"
```

- When PowerPath multipathing is used on the Linux host, edit these lines as follows:

```
# ORACLEASM_SCANORDER: Matching patterns to order disk scanning
```

```
ORACLEASM_SCANORDER="emcpower"
```

```
# ORACLEASM_SCANEXCLUDE: Matching patterns to exclude disks from scan
```

```
ORACLEASM_SCANEXCLUDE="sd"
```

2. Shutdown the Oracle instance running on the specific host.
3. Execute the following two commands to restart Oracle ASM:

```
/etc/init.d/oracleasm stop
```

```
/etc/init.d/oracleasm start
```

Post-Configuration Steps - Using the PowerStore system

After the host configuration is completed, you can access the PowerStore system from the host.

You can create, present, and manage volumes accessed from the host via PowerStore Manager, CLI, or REST API. Refer to the *PowerStore Manager Online Help*, *CLI Reference Guide*, or *REST API Reference Guide* for additional information.

Presenting PowerStore Cluster Volumes to the Linux Host

Specify Linux as the operating system when presenting PowerStore cluster volumes to the Linux host.

- When adding host groups and hosts to allow Linux hosts to access PowerStore cluster volumes, specify **Linux** as the operating system for the newly created hosts.
- Setting the operating system of the host is required for optimal interoperability and stability of the host with PowerStore cluster storage. You can adjust the setting while the host is online and connected to the PowerStore cluster with no I/O impact.

Partition Alignment in Linux

When using disk partitions with a Linux host attached to a PowerStore cluster, alignment is recommended. Follow these guidelines to align disk partitions.

To align partitions on PowerStore cluster volumes that are presented to Linux hosts, use the default value (2048). Then, create a partition using the `fdisk` command to ensure that the file system is aligned.

When you perform partition alignment, the logical device (`/dev/mapper/<naa-name>`) should be used rather than the physical device (`/dev/<device>`). When multipathing is not used (for example in a virtual machine), the physical device should be used.

The following example demonstrates using the `fdisk` command to create an aligned partition on a PowerStore cluster volume.

```
[root@lg114 ~]# fdisk -c -u /dev/mapper/ 368ccf098003f1461569ea4750e9dac50
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel
Building a new DOS disklabel with disk identifier 0x12d4e90c Changes will remain
in memory only, until you decide to write them. After that, of course, the previous
content won't be recoverable.
Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)

Command (m for help): p


Disk /dev/mapper/3514f0c5b12a00004: 1649.3 GB, 1649267441664 bytes
255 heads, 63 sectors/track, 200512 cylinders, total 3221225472 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes I/O size
(minimum/optimal): 16384 bytes / 65536 bytes Disk identifier: 0x12d4e90c
    Device Boot      Start         End      Blocks      Id      System
```

In this mode, rather than using "cylinders" for creating partitions, the `fdisk` command uses sectors, which are a direct mapping to the LBA space of the cluster. Thus, to verify that the partition is aligned, simply verify that the starting sector number is a multiple of 16 (16 sectors, at 512 bytes each, is 8 KB). The `fdisk` command defaults to a starting sector for the first partition of 2048, which is divisible by 16, and thus is correctly assigned.

Creating a File System

When creating a file system with PowerStore cluster storage, use its default block size, and disable UNMAP during creation.


It is recommended to create the file system using its default block size (using a non-default block size may lead to unexpected behavior). Refer to your operating system and file system documentation.

 **NOTE:** Creating a file system with UNMAP enabled on a host connected to PowerStore may result in an increased amount of write I/Os to the storage subsystem. It is highly recommended to disable UNMAP during file system creation.

To disable UNMAP during file system creation:

- When creating a file system using the `mke2fs` command - Use the "`-E nodiscard`" parameter.
- When creating a file system using the `mkfs.xfs` command - Use the "`-K`" parameter.

For a more efficient data utilization and better performance, use Ext4 file system with PowerStore cluster storage instead of Ext3. For details on converting to Ext4 file system (from either Ext3 or Ext2), refer to <https://ext4.wiki.kernel.org/index.php/UpgradeToExt4>.

 **NOTE:** File system configuration and management are out of the scope of this document.

Host Configuration for AIX


This chapter contains the following topics:


Topics:

- [Recommended Configuration Values Summary](#)
- [Fibre Channel Configuration](#)
- [Dell EMC AIX ODM Installation](#)

Recommended Configuration Values Summary

The following table summarizes all used and recommended variables and their values when configuring hosts for AIX.

 **NOTE:** Unless indicated otherwise, use the default parameters values.

Validation	Impact	Severity	Refer to Section
To clarify the above note for using default parameter settings unless stated otherwise in this chapter, make sure that the following are set per the default OS setting: <ul style="list-style-type: none"> • LUN and HBA queue depth • HBA timeout 	Stability & Performance	Recommended	For further details, refer to OS and HBA documentation.
Fibre Channel Configuration: No more than eight (8) paths per LUN should be used with an AIX host that is connected to PowerStore.	Stability & Performance	Mandatory	Fibre Channel Configuration
LUN Queue Depth: If I/O throttling is required, the default depth value of 256 should be modified to a lower value.	Performance	Recommended	Queue Depth
HBA FC max I/O size: <code>max_xfer_size</code> should be set to 1 MB.	Performance	Recommended	Fibre Channel Adapter Device Driver Maximum I/O Size
ODM Minimum Version: DellEMC.AIX.6.2.0.1.tar.Z  NOTE: When upgrading the connected PowerStore system to PowerStore OS version 1.0.2.0.5.003 (or later), AIX ODM version 6.2.0.1 is required when using PowerStore.	Stability	Stability	Dell EMC ODM Installation
To enable Fast I/O Failure for all fscsi devices, set the fscsi device attribute set <code>fc_err_recov</code> to <code>fast_fail</code>	Stability and Performance	Warning	Fast I/O Failure for Fibre Channel Devices
To enable dynamic tracking of FC devices, set:	Stability and Performance	Warning	Dynamic Tracking

Validation	Impact	Severity	Refer to Section
dyntrk= yes			

Fibre Channel Configuration

This section describes the recommended configuration that should be applied when attaching AIX hosts to PowerStore cluster using Fibre Channel.

- NOTE:** When using Fibre Channel with PowerStore, the FC Host Bus Adapters (HBA) issues that are described in this section should be addressed for optimal performance.
- NOTE:** In general, no more than eight(8) paths per LUN should be used with an AIX host that is connected to PowerStore. If more paths are needed, an RPQ is required.

Pre-Requisites

Before you install HBAs on an AIX host, the following pre-requisites should be met.

Follow the IBM recommendations for installation and setup of the appropriate HBA for your system. It is recommended to install the latest driver version (patch), as described on the IBM support site for each specific FC HBA.

Refer to the E-Lab Interoperability Navigator (<https://elabnavigator.emc.com>) for supported FC HBA models and drivers.

Queue Depth

Follow these recommendations when setting queue depth.

- NOTE:** Changing queue depth settings is designed for advanced users. Increasing the queue depth may cause the host to overstress other clusters that are connected to the AIX host, resulting in performance degradation while communicating with them. Therefore, especially in mixed environments with multiple cluster types that are connected to the AIX host, compare the PowerStore recommendations for queue depth with those of other platforms before applying them.

Queue depth is the number of SCSI commands (including I/O requests) that a storage device can handle at a given time. Queue depth can be set on either of the following levels:

- Initiator level - HBA queue depth
- LUN level - LUN queue depth

LUN queue depth setting controls the amount of outstanding I/O requests per a single path.

HBA queue depth (also referred to as execution throttle) setting controls the number of outstanding requests per HBA port.

For optimal operation with PowerStore, it is recommended to adjust the HBA queue depth setting of the FC HBA.

The driver module for the card controls LUN queue depth settings at the operating system level. Change the LUN queue depth default value (256) to a lower value only if I/O throttling is required.

Setting the Queue Depth

Follow these steps to set the HBA queue depth.

About this task

To set the HBA queue depth:

Steps

- Run the `chdev` command for each HBA in the AIX host to set the HBA firmware level queue depth:

```
chdev -l fcs# -a num_cmd_elems=2048 -P
```
- Reboot the AIX host to apply the HBA queue depth settings.

Fast I/O Failure for Fibre Channel Devices


This topic describes the Fast I/O Failure feature for FC devices and details the setting recommendations.

AIX supports Fast I/O Failure for Fibre Channel devices after link events in a switched environment.

When the FC adapter driver detects a link event, such as a lost link between a storage device and a switch, it waits for the fabric to stabilize (approximately 15 s). If the device is not on the fabric, the FC adapter driver fails all new I/Os or future retries of the failed I/Os, until the device rejoins the fabric. The fscsi device attribute `fc_err_recov` controls Fast I/O Failure (default value is `delayed_fail`).

It is recommended to enable Fast I/O Failure for FC adapters that are connected to PowerStore storage.

To enable Fast I/O Failure for all fscsi devices, set the `fc_err_recov` attribute to `fast_fail`, as shown in the following example:

 **NOTE:** In the example, the fscsi device instance is `fscsi0`.

```
chdev -l fscsi0 -a fc_err_recov=fast_fail -P
```

Run the following command to verify that the setting was enabled in the ODM:

```
lsattr -El fscsi0
```

 **NOTE:** The `-P` flag only modifies the setting in the ODM and requires a system reboot for the changes to apply.

Fast fail logic is applied when the switch sends a Registered State Change Notification (RSCN) to the adapter driver, indicating a link event with a remote storage device port.

Fast I/O Failure is useful when multipathing software is used. Setting the `fc_err_recov` attribute to `fast_fail` can decrease I/O failure due to link loss between the storage device and switch, by supporting faster failover to alternate paths.

Dynamic Tracking

This topic describes the dynamic tracking logic for FC devices and details the setting recommendations.

Dynamic tracking logic is applied when the adapter driver receives an indication from the switch that a link event with a remote storage device port has occurred.

If dynamic tracking of FC devices is enabled, the FC adapter driver detects when the Fibre Channel N_Port ID of a device changes. The FC adapter driver then reroutes the traffic that is destined for that device to the new address, while the devices are still online.


Events that can cause an N_Port ID to change include:

- Moving a cable that connects a switch to a storage device from one switch port to another.
- Connecting two separate switches using an Inter-Switch Link (ISL).
- Rebooting a switch.

The fscsi device attribute `dyntrk` controls dynamic tracking of FC devices (default value is `no` for non-NPIV configurations).

It is recommended to enable dynamic tracking for PowerStore volumes.

To enable dynamic tracking for FC devices, change all fscsi device attributes to `dyntrk=yes`, as shown in the following example:

 **NOTE:** In the example, the fscsi device instance is `fscsi0`.

```
chdev -l fscsi0 -a dyntrk=yes -P
```

Run the following command to verify that the setting was enabled in the ODM:

```
lsattr -El fscsi0
```

 **NOTE:** The `-P` flag only modifies the setting in the ODM and requires a system reboot for the changes to apply.

Fibre Channel Adapter Device Driver Maximum I/O Size

Set the `max_xfer_size` attribute for optimal AIX host operation over FC with PowerStore.


Prerequisites

The `max_xfer_size` FC HBA adapter device driver attribute for the `fscsi` device controls the maximum I/O size that the adapter device driver can handle. This attribute also controls a memory area the adapter uses for data transfers.

For optimal AIX host operation over FC with PowerStore, perform the following steps:

Steps

1. Run the following command on all FC adapters that are connected to PowerStore:
`chdev -l fcs0 -a max_xfer_size=0x100000 -P`
2. Reboot the AIX host to apply the `max_xfer_size` setting adjustments.

 **NOTE:** For virtualized AIX hosts, make sure to apply the `max_xfer_size` setting adjustments on all LPARs of the host that is connected to PowerStore storage.

Dell EMC AIX ODM Installation

This topic provides an introduction to Dell EMC ODM.

The Object Data Manager (ODM) is a database of system and device configuration information that is integrated into the AIX operating system. Information is stored and maintained as objects with associated characteristics. The Dell EMC ODM support package contains a series of installable filesets. These filesets are used to update the AIX ODM with customized Dell EMC storage device configuration attributes.

Dell EMC AIX ODM Installation Requirements

This section outlines the requirements for Dell EMC AIX ODM installation.

To meet the Dell EMC storage cluster requirements, you must install the correct Dell EMC ODM filesets to support Fibre Channel attachment to the PowerStore cluster.

The minimum ODM and AIX operating system versions that are supported with PowerStore and native MPIO are:

`DellEMC.AIX.6.2.0.1.tar.Z` -> For AIX 7.1, 7.2 and VIOS 3.1.0

PowerStore AIX ODM software package needs to be updated to version 6.2.0.1 prior to running PowerStore OS 1.0.2.0.5.003 (or later) with PowerPath. An RPQ for PowerPath is required for this configuration. PowerStore OS SP2 contains a change to the PowerStore cluster serial number that requires this new ODM package version. A reboot of the connected AIX host is required for this change to take effect.

To install the Dell EMC fileset:

1. Download the correct Dell EMC ODM fileset version and place it in the `/tmp/ODM` directory.
2. Untar the `DellEMC.AIX.6.2.0.1.tar.Z` file, using the following command:

```
uncompress DellEMC.AIX.6.2.0.1.tar.Z
tar -xvf DellEMC.AIX.6.2.0.1.tar.Z
```
3. Run the following command to create a table of contents file:

```
inutoc .
```
4. Run the following command to install the following filesets to support native MPIO:

```
installp -ad . EMC.PowerStore.aix.rte EMC.PowerStore.fcp.MPIO.rte
```

Installation Summary				
Name	Level	Part	Event	Result

EMC.PowerStore.aix.rte	6.2.0.1	USR	APPLY	SUCCESS
EMC.PowerStore.fcp.MPIO.rte	6.2.0.1	USR	APPLY	SUCCESS

5. Run the following command to install the following filesets to support PowerPath (an RPQ for PowerPath is required for this configuration):

```
installp -ad . EMC.PowerStore.aix.rte EMC.PowerStore.fcp.rte
```

Installation Summary				
Name	Level	Part	Event	Result
EMC.PowerStore.aix.rte	6.2.0.1	USR	APPLY	SUCCESS
EMC.PowerStore.fcp.rte	6.2.0.1	USR	APPLY	SUCCESS

Host Configuration for Solaris

This chapter contains the following topics:

Topics:

- [Recommended Configuration Values Summary](#)
- [Fibre Channel Configuration](#)
- [Solaris Host Parameter Settings](#)
- [Post configuration steps - using the PowerStore system](#)

Recommended Configuration Values Summary

The following table summarizes all used and recommended variables and their values when configuring hosts for Solaris Operating System.

NOTE: Unless indicated otherwise, use the default parameters values.

NOTE: Solaris OS can use two types of disk drivers to manage disk storage. The driver type depends on the platform architecture (x86 or SPARC) and the version of Solaris installed on the platform.

All versions of Solaris x86 OS are using SD disk drivers to manage all disk storage.

For SPARC platform versions prior to 11.4, release SSD driver type is used to manage all disk storage.

To simplify configuration and disk storage management, as of SPARC platform version 11.4, both platforms are using SD driver.

If the SPARC system is upgraded to Solaris 11.4 from one of the earlier versions, the system will continue to use SSD driver. All new installations of Solaris 11.4 will be configured to use SD driver for disk management.

Make sure that you update the tuning settings in the correct disk driver configuration file.

Validation	Config File	Impact	Severity	Refer to Section
Set the maximum I/O size to 1 MB: <code>set maxphys = 0x100000</code>	/etc/system	Stability	Mandatory	Updating ssd.conf configuration file Updating sd.conf configuration file
Configure ZFS space reclamation: <code>set zfs:zfs_unmap_ignore_size=256</code> <code>set zfs:zfs_log_unmap_ignore_size=256</code>	/etc/system	Efficiency	Recommended	Updating /etc/system configuration file
Enable Solaris MPxIO multipathing: <code>mpxio-disable="no";</code>	fp.conf	Stability	Recommended	Updating fp.conf configuration file
Fibre Channel path failover tuning: <code>fp_offline_ticker = 20;</code>	fp.conf	Stability	Recommended	Updating fp.conf configuration file
Fibre Channel path failover tuning: <code>fcp_offline_delay = 20;</code>	fp.conf	Stability	Recommended	Updating fp.conf configuration file

Validation	Config File	Impact	Severity	Refer to Section
Maximum I/O size for ssd driver for Solaris 10, 11-11.3 (SPARC) <code>ssd_max_xfer_size=0x100000;</code>	ssd.conf	Stability	Mandatory	Updating ssd.conf configuration file
Maximum I/O size for sd driver for Solaris 11.4 (SPARC) 11.x (x86) <code>sd_max_xfer_size=0x100000;</code>	sd.conf	Stability	Mandatory	Updating sd.conf configuration file
Soaris ssd driver tuning for Solaris 10, 11-11.3 (SPARC) <code>ssd-config-list = "DellEMC PowerStore", "throttle-max:64, physical-block-size:4096, disksort:false, cache-nonvolatile:true";</code>	ssd.conf	Stability	Recommended	Updating ssd.conf configuration file
Soaris sd driver tuning for Solaris 11.4 (SPARC) 11.x (x86) <code>sd-config-list = "DellEMC PowerStore", "throttle-max:64, physical-block-size:4096, disksort:false, cache-nonvolatile:true";</code>	sd.conf	Stability	Recommended	Updating sd.conf configuration file
Solaris MPxIO multi-path driver tuning <code>load-balance="round-robin"; auto-failback="enable";</code>	scsi_vhci.conf	Stability	Mandatory	Updating scsi_vhci.conf configuration file
Solaris MPxIO multi-path driver tuning <code>scsi-vhci-update-pathstate-on-reset = "DellEMC PowerStore", "yes";</code>	scsi_vhci.conf	Stability	Mandatory	Updating scsi_vhci.conf configuration file

Fibre Channel Configuration

This section describes the recommended configuration to apply when attaching a host to thePowerStore cluster using host Fibre Channel HBAs.

NOTE: When using Fibre Channel with PowerStore, the FC HBA (Host Bus Adapters) issues described in this section should be addressed for optimal performance.

NOTE: Before you proceed, review [Fibre Channel SAN Guidelines](#).

Pre-Requisites

Before installing HBAs in a Solaris host, the following pre-requisites should be met:

- Follow Oracle's recommendations for installation and setup of the appropriate HBA for your system.
- It is recommended to install the latest driver version (patch), as described on the Oracle support site for each specific FC HBA.
- Refer to the E-Lab Interoperability Navigator (<https://elabnavigator.emc.com>) for supported FC HBA models and drivers.

Queue Depth

Queue depth is the amount of SCSI commands (including I/O requests) that can be handled by a storage device at a given time. A queue depth can be set on either of the following:

- Initiator level - HBA queue depth
- LUN level - LUN queue depth

The LUN queue depth setting controls the amount of outstanding I/O requests per a single path. The HBA queue depth (also referred to as execution throttle) setting controls the amount of outstanding I/O requests per HBA port.

With PowerStore and Solaris, the HBA queue depth setting should retain its default value, and the initial LUN queue depth setting should be modified to 64. This is a good starting point provided good I/O response times. The specific value can be adjusted based on particular infrastructure configuration, application performance and I/O profile details.


Solaris Host Parameter Settings

This section describes the Solaris host parameter settings required for optimal configuration when using Dell Technologies PowerStore storage.

Configuring Solaris native multipathing

For a PowerStore cluster to properly function with Oracle Solaris hosts, configure the multipath settings as described in the following sections.


 **NOTE:** If the host is connected to a cluster other than PowerStore, the configuration file may include additional devices.

 **NOTE:** Currently, PowerStore clusters are only supported with native Solaris multipathing (MPxIO).

Enable Solaris native multipathing on Solaris 10 and 11.0-11.4 hosts (SPARC and x86)

To enable management of storage LUNS that are presented to the host with MPxIO, use the following command:

```
# stmsboot -e
```

 **NOTE:** the host must be rebooted immediately after the command execution is complete. It is recommended to update all storage-related host configuration files before rebooting.

Updating scsi_vhci.conf configuration file

About this task

The `scsi_vhci.conf` file is used to configure third-party storage multipathing parameters on Solaris 11 hosts, based on SCSI inquiry responses. The host sends SCSI inquiry commands and, based on the returned data, MPxIO driver will activate the corresponding multipathing module. The load balancing and failover policies are also configured, based on the settings in the `scsi_vhci.conf` file.

Steps

1. Run the following command to verify the `scsi_vhci.conf` file location:

```
# ls /etc/driver/drv/
```
2. If the file is not in the expected location, run the following command to copy it from `/kernel/drv`:

```
# cp /kernel/drv/scsi_vhci.conf /etc/driver/drv
```
3. Run the following commands to create a backup copy and modify the file:

```
# cp -p /etc/driver/drv/scsi_vhci.conf /etc/driver/drv/scsi_vhci.conf_ORIG
# vi /etc/driver/drv/scsi_vhci.conf
```

Example

Below are the entries recommended for PowerStore storage.

```
load-balance="round-robin";
auto-failback="enable";
ddi-forceload =
    "misc/scsi_vhci/scsi_vhci f asym sun",
    "misc/scsi_vhci/scsi_vhci f asym lsi",
    "misc/scsi_vhci/scsi_vhci f asym emc",
    "misc/scsi_vhci/scsi_vhci f sym emc",
    "misc/scsi_vhci/scsi_vhci f sym hds",
    "misc/scsi_vhci/scsi_vhci f sym",
    "misc/scsi_vhci/scsi_vhci f tpgs";

scsi-vhci-update-pathstate-on-reset = "DellEMC PowerStore",    "yes";
```

Parameter	Description	Value
load-balance	Specifies the default load balancing policy. Possible values are none, logical-block, round-robin	round-robin
auto-failback	Specifies whether LUN access should be restored on the path restore. Possible values are enable, disable	enable
scsi-vhci-update-pathstate-on-reset	Enable path status update on reset. Possible values are yes, no	yes

Host storage tuning parameters

Configure host storage tuning parameters as described in the following sections.

Updating fp.conf configuration file

About this task

The fp.conf host file is used to control options for Fibre Channel storage. The MPxIO settings in fp.conf file should match the settings in scsi_vhci.conf.

Steps

1. Run the following command to verify the fp.conf file location:
ls /etc/driver/drv/
2. If the file is not in the expected location , run the following command to copy it from /kernel/drv:
cp /kernel/drv/fp.conf /etc/driver/drv
3. Run the following commands to create a backup copy and modify the file:
cp -p /etc/driver/drv/fp.conf /etc/driver/drv/fp.conf_ORIG
vi /etc/driver/drv/fp.conf

Example

Below are the entries recommended for PowerStore storage.

```
mpxio-disable="no";
fp_offline_ticker=20;
```

Parameter	Description	Value
mpxio-disable	Specifies whether MPxIO is disabled. MPxIO can be enabled for Fibre Channel storage or it can be disabled for a particular HBA.	no
fp_offline_ticker	Used to prevent errors from being generated immediately for transient/brief connection interruptions, and should prevent any errors if the connections are restored before the fcp and fp delays expire.	20

Updating /etc/system configuration file

the /etc/system host file is used to control Solaris kernel tuning settings.

Example

```
set maxphys = 0x100000
set zfs:zfs unmap ignore size=256
set zfs:zfs log unmap ignore size=256
```

Parameter	Description	Value
set maxphys	Sets the maximum size of a single I/O request. For PowerStore it must be set to no more than 1MB.	0x100000
set zfs:zfs_unmap_ignore_size	ZFS TRIM settings	256
set zfs:zfs_log_unmap_ignore_size	ZFS TRIM settings	256

Updating fcp.conf configuration file

About this task

The fcp.conf host file is used to control options for Fibre Channel storage.

Steps

1. Run the following command to verify the fcp.conf file location:
ls /etc/driver/drv/
2. If the file is not in the expected location, run the following command to copy it from /kernel/drv:
cp /kernel/drv/fcp.conf /etc/driver/drv
3. Run the following commands to create a backup copy and modify the file:
cp -p /etc/driver/drv/fcp.conf /etc/driver/drv/fcp.conf_ORIG
vi /etc/driver/drv/fcp.conf

Example

Below are the entries recommended for PowerStore storage.

```
fcp_offline_delay = 20;
```

Parameter	Description	Value
fcp_offline_delay	The tuning setting is designed to prevent errors from being	20

Parameter	Description	Value
	generated immediately for transient/brief connection interruptions, and should prevent any errors if the connections are restored before the fcp and fp delays expire.	

Updating `ssd.conf` configuration file (Solaris 10 and 11.0-11.3 SPARC)

About this task

The `ssd.conf` host file is used to control options for SCSI disk storage device.

Steps

1. Run the following command to verify the `ssd.conf` file location:

```
# ls /etc/driver/drv/
```
2. If the file is not in the expected location, run the following command to copy it from `/kernel/drv`:

```
# cp /kernel/drv/ssd.conf /etc/driver/drv
```
3. Run the following commands to create a backup copy and modify the file:

```
# cp -p /etc/driver/drv/ssd.conf /etc/driver/drv/ssd.conf_ORIG
# vi /etc/driver/drv/ssd.conf
```

Example

Below are the entries recommended for PowerStore storage.

```
ssd_max_xfer_size=0x100000;
ssd-config-list = "DellEMC PowerStore", "throttle-max:64, physical-block-size:4096,
disksort:false, cache-nonvolatile:true";
```

Parameter	Description	Value
<code>ssd_max_xfer_size</code>	Restrict SCSI disk driver maximum I/O size to 1MB	0x100000
<code>ssd-config-list</code>	SCSI inquiry storage response, the VPD for PowerStore LUNs	DellEMC PowerStore
<code>throttle-max</code>	Maximum SCSI queue depth setting	64
<code>physical-block-size</code>	Optimal LUN block size in bytes	4096
<code>disksort</code>	SCSI device command optimization	false
<code>cache-nonvolatile</code>	Indicate whether the storage has a nonvolatile cache	true

Updating `sd.conf` configuration file (Solaris 11.x x86 and 11.4 SPARC)

About this task

The `sd.conf` host file is used to control options for SCSI disk storage device.

Steps

1. Run the following command to verify the `sd.conf` file location:

```
#ls /etc/driver/drv/
```
2. If the file is not in the expected location, run the following command to copy it from `/kernel/drv`:

```
# cp /kernel/drv/sd/conf /etc/driver/drv
```

3. Run the following commands to create a backup copy and modify the file:

```
# cp -p /etc/driver/drv/sd.conf /etc/driver/drv/sd.conf_ORIG
```

```
# vi /etc/driver/drv/sd.conf
```

Example

Below are the entries recommended for PowerStore storage.

```
sd_max_xfer_size=0x100000;  
sd-config-list = "DellEMC PowerStore", "throttle-max:64, physical-block-size:4096,  
disksort:false, cache-nonvolatile:true";
```


Parameter	Description	Value
sd_max_xfer_size	Restrict SCSI disk driver maximum I/O size to 1MB	0x100000
sd-config-list	SCSI inquiry storage response, the VPD for PowerStore LUNs	DellEMC PowerStore
throttle-max	Maximum SCSI queue depth setting	64
physical-block-size	Optimal LUN block size in bytes	4096
disksort	SCSI device command optimization	false
cache-nonvolatile	Indicate whether the storage has a nonvolatile cache	true

Post configuration steps - using the PowerStore system

When host configuration is complete, you can access the PowerStore system from the host.

You can create, present, and manage volumes accessed from the host via PowerStore Manager, CLI, or REST API. Refer to the *PowerStore Manager Online Help*, *CLI Reference Guide*, or *REST API Reference Guide* for additional information.

When adding host groups and hosts to allow Solaris hosts to access PowerStore volumes, specify **Solaris** as the operating system for the newly-created hosts.

 **NOTE:** Setting the host's operating system is required for optimal interoperability and stability of the host with PowerStore storage. You can adjust the setting while the host is online and connected to the PowerStore cluster with no I/O impact.

 **NOTE:** Refer to the *PowerStore Configuring Volumes* Guide for additional information.

Partition alignment in Solaris

Use Solaris `format` command to create partitions aligned to 4K on a PowerStore cluster LUNs for use as a raw or UFS devices.

When a PowerStore LUN is added to the ZFS pool, ZFS automatically creates aligned partitions.

Host Configuration for HP-UX


This chapter contains the following topics:

Topics:

- [Recommended Configuration Values Summary](#)
- [Fibre Channel Configuration](#)
- [HP-UX Host Parameter Settings](#)
- [Multipathing Software Configuration](#)
- [Post-Configuration Steps - Using the PowerStore System](#)

Recommended Configuration Values Summary

The following table summarizes all used variables and their values when configuring hosts for HP-UX

 **NOTE:** Unless indicated otherwise, use the default parameters values.

Validation	Impact	Severity	Refer to Section
To clarify the above note for using default parameter settings unless stated otherwise in this chapter, make sure that the following are set per the default OS setting: <ul style="list-style-type: none"> • LUN and HBA queue depth • HBA timeout 	Stability & Performance	Recommended	For further details, refer to OS and HBA documentation.
Specify HP-UX as the operating system for each defined host.	Serviceability	Mandatory	Presenting PowerStore Volumes to the HP-UX Host
Maximum transfer length: Change <code>escsi_maxphys</code> value to 1MB (256 increments of 4KB) <code>scsimgr save_attr -a escsi_maxphys=256</code>	Performance	Mandatory	Maximum Transfer Length
Load balancing: Keep the following HP-UX native multipathing parameters at their default values: <ul style="list-style-type: none"> • <code>load_bal_policy</code> - set to "round_robin" • <code>path_fail_secs</code> - set to 120 seconds 	Performance	Mandatory	Configuring Native Multipathing using HP-UX Multipath (MPIO)
Temporarily disable UNMAP during file systems creation (only when using Veritas Volume Manager):	Performance	Recommended	Creating File System

Validation	Impact	Severity	Refer to Section
<ul style="list-style-type: none"> To temporarily disable UNMAP for the targeted device on the host (prior to file system creation): #vxdisk set reclaim=off "disk name" To re-enable UNMAP for the targeted device on the host (after file system creation): # vxdisk reclaim "disk name" 			

Fibre Channel Configuration

This section describes the recommended configuration that should be applied when attaching hosts to PowerStore cluster using Fibre Channel.

NOTE: This section applies only to Fibre Channel.

NOTE: PowerStore supports only FC-SW FCP connections. GigE iSCSI and FC direct connections from HP-UX initiators to PowerStore target ports are not supported.

NOTE: Before you proceed, review [Fibre Channel SAN Guidelines](#).

Pre-Requisites

This section describes the pre-requisites for FC HBA configuration

- Refer to the E-Lab Interoperability Navigator (<https://elabnavigator.emc.com>) for supported FC HBA models and drivers.
- Verify all HBAs are at the supported driver, firmware and BIOS versions.
- For instructions about installing the FC HBA and upgrading the drivers or the firmware, see HP documentation.

HP-UX Host Parameter Settings

This section describes the HP-UX host parameter settings required for optimal configuration when using Dell Technologies PowerStore storage.

NOTE: PowerStore is supported only with HP-UX version Iliiv3. HP-UX versions Iliiv2 and Iliiv1 are only supported with the volume set addressing method, which currently is not available with PowerStore.

To configure HP-UX Iliiv3 with PowerStore make sure that the following requirements are met:

- PowerStore supports only native MPIO on HP-UX Iliiv3. For further details refer to the <Multipathing Software Configuration> section.
- Maximum request transfer length of PowerStore array is 2048 blks (512 byte blks) for a maximum transfer length of 1MB.

Maximum Transfer Length

HP-UX Iliiv3 implements a default maximum transfer length of 2MB and does not have support for VPD page B0h. A tunable parameter, named `escsi_maxphys`, enables the modification of its FC fcp maximum transfer length. You can configure `escsi_maxphys` as follows:

- Reset to default during host reboot - `scsimgr set_attr -a escsi_maxphys=<value>`

- Persistent through reboot - `scsimgr save_attr -a escsi_maxphys=<value>`

The set value is defined in 4KB increments. To support the PowerStore devices in HP-UX, change the `escsi_maxphys` value to 256 using the following commands:

- 1MB max transfer length, reset to default 2MB during host reboot - `scsimgr set_attr -a escsi_maxphys=256`
- 1MB max transfer length persistent through reboot `scsimgr save_attr -a escsi_maxphys=256`

NOTE: You can configure the `escsi_maxphys` attribute only on a global basis and it will be effective for all FC fcp black devices connected to the host.

Multipathing Software Configuration

This topic introduces Multipathing Software Configuration for HP-UX

PowerStore supports only native multipathing using multipath I/O (MPIO) with HP-UX IliV3.

NOTE: Other multipathing software, such as PowerPath or Veritas Dynamix Multipathing (DMP), are not supported.

Configuring Native Multipathing Using HP-UX Multipath I/O (MPIO)

This topic describes configuring native multipathing using HP-UX Multipath I/O (MPIO).

For optimal operation with PowerStore storage, configure the following HP-UX native multipathing parameters at their default values:

- `load_bal_policy` - I/O load balancing policy: This parameter must be set with the default value of "round-robin" to set the Round-Robin (RR) policy for MPIO for devices presented from PowerStore. Using this policy, I/O operations are balanced across all available paths.
- `path_fail_secs` - Timeout in seconds before declaring a LUN path offline: This parameter must be set with the default value of 120 seconds.

NOTE: The man page of `scsimgr_esdisk(7)` provides a list of parameters related to HP-UX native multipathing.

Post-Configuration Steps - Using the PowerStore System

This topic describes the post-configuration steps using the PowerStore system.

After the host configuration is completed, you can use the PowerStore storage from the host.

You can create, present, and manage volumes accessed from the host via PowerStore Manager, CLI, or REST API. Refer to the *PowerStore Manager Online Help*, *CLI Reference Guide*, or *REST API Reference Guide* for additional information.

Presenting PowerStore Volumes to the HP-UX Host

This topic describes presenting PowerStore volumes to the HP-UX host.

When adding host groups and hosts to allow HP-UX hosts to access PowerStore volumes, specify **HP-UX** as the operating system for the newly-created hosts.

NOTE: Setting the host's operating system is required for optimal interoperability and stability of the host with PowerStore storage. You can adjust the setting while the host is online and connected to the PowerStore cluster with no I/O impact.

NOTE: Refer to the *PowerStore Configuring Volumes Guide* for additional information.

Creating a file system

About this task

File system configuration and management are out of the scope of this document.

NOTE: Some file systems may require you to properly-align the file system on the PowerStore volume. It is recommended to use specified tools to optimally match your host with application requirements.

NOTE: Creating a file system formatting with UNMAP enabled on a host connected to PowerStore may result in an increased amount of write I/Os to the storage subsystem. When possible, it is highly recommended to disable UNMAP during file system creation. Disabling UNMAP can be done when using the Veritas Volume Manager on the HP-UX host. However, when using the HP-UX native volume manager, this recommendation is not applicable as UNMAP is not supported in this case.

To disable UNMAP during file system creation (only when using Veritas Volume Manager):

Steps

1. Access the HP-UX host using SSH as `root`.
2. Run the following `vxdisk` command to temporarily disable UNMAP for the targeted device on the host (prior to file system creation):

```
# vxdisk set reclaim=off "disk name"
```

3. Once file system creation is complete, re-enable UNMAP for the targeted device on the host, by running the following command:

```
# vxdisk reclaim "disk name"
```

NOTE: To verify the current setting of a specific device using its corresponding disk group, run the following `vxprint` command:

```
#vxprint -z -g "disk group name"
```

Example: Using the `vxprint` command to verify the current UNMAP setting of a specific device:

```
# vxprint -z -g testdg
...
dm testdg02      3pardata0_55 -      2031232 -      -      -
sd testdg02-01 -   ENABLED 409600 -   RECLAIM -   -
sd testdg02-02 -   ENABLED 67840  -   RECLAIM -   -
```

Considerations for Boot from SAN with PowerStore

This appendix provides considerations for configuring boot from SAN with PowerStore.

Topics:

- [Consideration for Boot from SAN with PowerStore](#)

Consideration for Boot from SAN with PowerStore

NOTE: See your operating system documentation for general boot from SAN configuration.

NOTE: The current PowerStore version does not support mapping individual LUNs to host under host group.

Follow these guidelines when configuring boot from SAN:

- Before the installation on a PowerStore LUN, the host must have been manually registered on the system and assigned to a Storage Group.
 - At least one LUN must be bound to the host's Volume Group and owned by the node that is connected to the HBA being used for the fabric boot.
 - The lowest-numbered path to the boot LUN must be the active path.
- It is recommended to assign the boot LUN with Host LUN ID 0. During the installation procedure, it is recommended that only one LUN be mapped to a host for ease of use.
 - Once the installation is complete, additional LUNs may be mapped to the host.

Follow these guidelines for clustered boot from SAN mapping:

- For every physical host, there should be one host in the PowerStore manager. Do not create a host group.
- Boot LUNs are mapped only for the specific host.
- Shared LUNs are mapped to all hosts in the cluster. The user must keep the same LUN ID across the cluster for shared LUNs.

When migrating a boot volume, using the PowerStore native migration tool, follow these guidelines:

- Power off the connected host before the migration; after the boot LUN ID is successfully changed, power it back on .
- It is recommended to assign the boot LUN with HOST LUN ID 0.
- When a volume is migrated from one PowerStore appliance to another appliance in the same cluster using the native migration tool, the LUN ID number changes automatically.
- After migrating a boot from SAN volume, the LUN ID number can be changed back to 0.
- Perform the following steps to change the boot LUN ID:
 1. On the PowerStore Manager of the destination appliance, select **Storage > Volumes**.
 2. Click the name of the boot volume and select the **Host Mapping** tab.
 3. Click the checkbox next to the boot volume name and select **MORE ACTIONS > Edit Logical Unit Number**.

NOTE: Changing the LUN ID number is a disruptive operation. The following message is displayed before changing the number: Changing the Logical Unit Number of the host will disrupt its access to the volume until a host-side rescan is performed.